

Human–Computer Interaction Series

Michael Filimowicz · Veronika Tzankova  
*Editors*

New Directions  
in Third Wave  
Human–Computer  
Interaction: Volume 1 -  
Technologies

 Springer

# Human–Computer Interaction Series

## **Editors-in-chief**

Desney Tan  
Microsoft Research, Redmond, WA, USA

Jean Vanderdonckt  
Louvain School of Management, Université catholique de Louvain,  
Louvain-la-Neuve, Belgium

More information about this series at <http://www.springer.com/series/6033>

Michael Filimowicz • Veronika Tzankova  
Editors

New Directions in Third  
Wave Human-Computer  
Interaction:  
Volume 1 - Technologies

 Springer

*Editors*

Michael Filimowicz  
School of Interactive Arts & Technology  
Simon Fraser University  
Surrey, BC, Canada

Veronika Tzankova  
School of Communications  
Simon Fraser University  
Burnaby, BC, Canada

ISSN 1571-5035

ISSN 2524-4477 (electronic)

Human–Computer Interaction Series

ISBN 978-3-319-73355-5

ISBN 978-3-319-73356-2 (eBook)

<https://doi.org/10.1007/978-3-319-73356-2>

Library of Congress Control Number: 2018942204

© Springer International Publishing AG, part of Springer Nature 2018

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by the registered company Springer International Publishing AG part of Springer Nature.

The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Contents

<b>1</b>	<b>Introduction   New Directions in Third Wave HCI</b> . . . . .	<b>1</b>
	Michael Filimowicz and Veronika Tzankova	
<b>2</b>	<b>The Relational Turn: Third Wave HCI and Phenomenology</b> . . . . .	<b>11</b>
	David J. Gunkel	
<b>3</b>	<b>Giving Form to Smart Objects: Exploring Intelligence as an Interaction Design Material</b> . . . . .	<b>25</b>
	Marco C. Rozendaal, Maliheh Ghajargar, Gert Pasman, and Mikael Wiberg	
<b>4</b>	<b>De-instrumentalizing HCI: Social Psychology, Rapport Formation, and Interactions with Artificial Social Agents</b> . . . . .	<b>43</b>
	Ritwik Banerji	
<b>5</b>	<b>Interaction Design for Metacreative Systems</b> . . . . .	<b>67</b>
	Oliver Bown and Andrew R. Brown	
<b>6</b>	<b>3D Printing Technologies: A Third Wave Perspective</b> . . . . .	<b>89</b>
	Deborah Lupton	
<b>7</b>	<b>Denaturalizing 3D Printing’s Value Claims</b> . . . . .	<b>105</b>
	Gabby Resch, Daniel R. Southwick, and Matt Ratto	
<b>8</b>	<b>Physical Computing   When Digital Systems Meet the Real World</b> . . . . .	<b>123</b>
	Alan Dix and Steve Gill	
<b>9</b>	<b>Third-Wave HCI Perspectives on the Internet of Things</b> . . . . .	<b>145</b>
	Tom Jenkins	
<b>10</b>	<b>Inclusion in the Third Wave: Access to Experience</b> . . . . .	<b>163</b>
	Christopher Power, Paul Cairns, and Mark Barlet	

<b>11</b>	<b>Deep Subjectivity and Empathy in Virtual Reality: A Case Study on the Autism TMI Virtual Reality Experience . . . . .</b>	<b>183</b>
	Jonathan Weinel, Stuart Cunningham, and Jennifer Pickles	
<b>12</b>	<b>Sonification and HCI . . . . .</b>	<b>205</b>
	Visda Goudarzi	
<b>13</b>	<b>Media Poetics and Cognition in Colocative Audiovisual Displays . . .</b>	<b>223</b>
	Michael Filimowicz	
<b>14</b>	<b>Language Technology and 3rd Wave HCI: Towards Phatic Communication and Situated Interaction . . . . .</b>	<b>251</b>
	Lars Borin and Jens Edlund	
<b>15</b>	<b>Sensorial Computing . . . . .</b>	<b>265</b>
	Puja Varsani, Ralph Moseley, Simon Jones, Carl James-Reynolds, Eris Chinellato, and Juan Carlos Augusto	

# Contributors

**Juan Carlos Augusto** Middlesex University London, London, UK

**Ritwik Banerji** Department of Music, University of California, Berkeley, CA, USA

**Mark Barlet** Able Gamers Charity, Charles Town, WV, USA

**Lars Borin** University of Gothenburg, Gothenburg, Sweden

**Oliver Bown** University of New South Wales, Sydney, Australia

**Andrew R. Brown** Griffith University, Nathan, Australia

**Paul Cairns** Department of Computer Science, University of York, York, UK  
The AbleGamers Charity, Kearneysville, WV, USA

**Eris Chinellato** Middlesex University London, London, UK

**Stuart Cunningham** Wrexham Glyndŵr University, Wrexham, UK

**Alan Dix** HCI Centre, University of Birmingham, Birmingham, UK

**Jens Edlund** KTH Royal Institute of Technology, Stockholm, Sweden

**Michael Filimowicz** School of Interactive Arts & Technology, Simon Fraser University, Surrey, BC, Canada

**Maliheh Ghajargar** Polytechnic University of Turin, Turin, Italy

**Steve Gill** Cardiff Metropolitan University, Cardiff, Wales

**Visda Goudarzi** Audio Arts and Acoustics, Columbia College Chicago, Chicago, IL, USA

**David J. Gunkel** Northern Illinois University, DeKalb, IL, USA

**Carl James-Reynolds** Middlesex University London, London, UK

**Tom Jenkins** Georgia Institute of Technology, Atlanta, GA, USA



**Simon Jones** Middlesex University London, London, UK

**Deborah Lupton** University of Canberra, Bruce, Australia

**Ralph Moseley** Middlesex University London, London, UK

**Gert Pasman** Delft University of Technology, Delft, The Netherlands

**Jennifer Pickles** University of Southampton, Southampton, UK

**Christopher Power** Department of Computer Science, University of York, York, UK

The AbleGamers Charity, Kearneysville, WV, USA

**Matt Ratto** Faculty of Information, University of Toronto, Toronto, Canada

**Gabby Resch** Faculty of Information, University of Toronto, Toronto, Canada

**Marco C. Rozendaal** Delft University of Technology, Delft, The Netherlands

**Daniel R. Southwick** Faculty of Information, University of Toronto, Toronto, Canada

**Veronika Tzankova** School of Communications, Simon Fraser University, Burnaby, BC, Canada

**Puja Varsani** Middlesex University London, London, UK

**Jonathan Weinel** Aalborg University, Aalborg, Denmark

**Mikael Wiberg** Umeå University, Umeå, Sweden

# About the Editors

**Michael Filimowicz** is Senior Lecturer in the School of Interactive Arts and Technology (SIAT) at Simon Fraser University. He has developed new forms of general purpose multimodal and audiovisual display technology, exploring novel product lines across different application contexts including gaming, immersive exhibitions, process control, command and control, telepresence, and simulation-based training.

**Veronika Tzankova** is a PhD candidate in SIAT, having previously gained her MA in the same program. Her Master's research was in the areas of interactive technologies within the contexts of social appropriation and democratic inclusiveness. Her current doctoral work is in the area of embodied cognition approaches to interaction design, with a focus on contact sports.

# Chapter 1

## Introduction | New Directions in Third Wave HCI



Michael Filimowicz and Veronika Tzankova

**Abstract** *New Directions in 3rd Wave Human-Computer Interaction* explores the diverse interdisciplinary inquiries comprising the forefront of developments in the field of HCI. This wide ranging collection aims at understanding the design, methods and applications of emerging forms of interaction with new technologies and the rich varieties of human knowledge and experiences. All chapters are structured around two major themes presented in two volumes: Volume 1 – Technologies, and Volume 2 – Methodologies.

### 1.1 Waves, Paradigms, and Cultures

*New Directions in 3rd Wave Human-Computer Interaction* explores the diverse interdisciplinary inquiries comprising the forefront of developments in the field of HCI. This wide ranging collection aims at understanding the design, methods and applications of emerging forms of interaction with new technologies and the rich varieties of human knowledge and experiences. All chapters are structured around two major themes presented in two volumes: Volume 1– Technologies, and Volume 2– Methodologies.

These two volumes address the widespread notion that the field of HCI can historically be divided into three ‘waves’ of approaches and application areas. Although there is a consensus on the presence of different ‘waves’, the definition

---

M. Filimowicz (✉)

School of Interactive Arts & Technology, Simon Fraser University, Surrey, BC, Canada  
e-mail: [mfa13@sfu.ca](mailto:mfa13@sfu.ca)

V. Tzankova

School of Communications, Simon Fraser University, Burnaby, BC, Canada  
e-mail: [vkt@sfu.ca](mailto:vkt@sfu.ca)

and understanding of what constitutes these is far from set. Bødker (2015) following Bannon (1986), for example, has defined the sequence and conceptualization of the waves as follows:

- 1st Wave – based on model-driven cognitive science and human factors methods and focusing on strict, formal guidelines grounded in systematic study and testing.
- 2nd Wave – initiated as an extension of the human-technology nexus to include collaborative, mediated, and distributed applications within work settings, involving a higher degree of participation from users of systems.
- 3rd Wave – triggered by the expanding context of HCI far beyond the workplace, an expansion resulting from the increasingly pervasive and ubiquitous essence of computation in everyday life. The 3rd wave places a stronger emphasis on human values, meaning-making, situated knowledge, and experiences.

Grudin (2005) considers the divergent foci on “computer operation, information systems management, and discretionary use” (46) to be the defining features of the three waves of HCI. Grudin emphasizes that these three strands of HCI have not converged and remain relatively autonomous today, and the three frames of inquiry are defined by two research cultures reinforced by differences in scholarly production and activity. A particularly interesting aspect of Grudin’s account is that these three kinds of HCI research run in a parallel historical developments. The first wave, for Grudin, is grounded in engineering psychology and has been initiated with the human-machine interface. Although such machines were not necessarily in the first instance computational, the consideration of Human Factors can be traced back to its earliest origins in Taylor’s scientific management of the early twentieth century. Grudin distinguishes compulsory (e.g. work and war) and discretionary (e.g. home and leisure) use as a critical difference that defines second- and third wave HCI research approaches. HCI oriented toward discretionary uses begins, in his account, with general purpose computers circa 1945, while information management-focused HCI begins in the mid 1970s with sociotechnical and participatory design approaches. Once founded, all three strands of HCI continue to progress in parallel through new developments discussed in journals, academic societies, and conferences which operate in relative cultural isolation from each other to the present day.

Harrison, Tatar, and Sengers (2007) formulate a third understanding of the three waves defined as a difference of ‘paradigms’ in the Kuhnian sense. For these authors, the first wave centered on engineering and human factors (i.e. human-machine “coupling”) and was essentially atheoretic and entirely oriented toward pragmatic design enhancements and solutions, such as helping pilots effectively utilize cockpit instrument panels of increasing complexity. The second wave of HCI was grounded in cognitive science disciplines, where human-computer interaction is understood in terms of information transfer and efficiency of communication between a mind-as-information-processor, and an interface communication with that mind. The third wave is characterized by a growing interest in design that takes into account the full ‘messy’ context of socially situated and embodied action, which introduces humanistic and social science considerations into design research.

These once marginal research agendas have moved toward more central positions in HCI discourse, prompting the notion of a third paradigm. “Participatory design, value-sensitive design, user experience design, ethnomethodology, embodied interaction, interaction design, and critical design” (2) as a grouping are brought together under the heading of the “phenomenological matrix” due to the highlighting of the embodied and socially situated interactor where more than simply efficiency of operation or information transfer is at stake.

While the conceptualization of the historical waves of HCI differ significantly in the details, at a more global level there is a commonality in the sense of a gradual and considerable expansion of HCI’s concerns, methodologies, and application areas. The earliest HCI work was strongly based on the concept of human-machine coupling, which expanded to workplace collaboration as computers came into mainstream professional use. Today, HCI can connect to increasingly more sides of human experience because now there is an app for every almost any aspect of daily life. Despite this clear sense of a commonly understood trajectory in the expansion of HCI’s domains of research and application, there are some tensions to be noted as to how one understands this historical progression. Do the new waves replace the old, or update them? Can one combine waves through hybrid research agendas? Are they complementary to each other? Which wave is ‘the right one’ today? What is meant by a ‘wave’ in the first place? What fourth wave might be on the horizon? These two volumes allow us to explore such general disciplinary questions while also focusing in depth on particular aspects of methodologies and technologies to better understand the range of practices associated with third wave HCI today.

## 1.2 Are the Waves ‘Paradigms?’

In one of the articles noted above that inspired this project, the three waves of HCI are understood as Kuhnian paradigms (Harrison et al. 2007). As compelling as this appears in terms of a general disciplinary taxonomy, careful consideration reveals some conceptual matters of potentially problematic nature. One of the most apparent issues to note is that the three paradigms of physics described by Kuhn (e.g. Aristotle, Newton, Einstein) unfold over thousands of years, whereas HCI paradigm formation seems to emerge and develop within a very short timeframe. Such fast speed of progression triggers new paradigmatic shifts in a matter of decades, producing a historical development several orders of magnitude faster than the sciences studied by Kuhn.

A key aspect of Kuhn’s paradigm theory relates to the idea of incommensurability between paradigms, and the alterations between normal and extraordinary science. The way scientists in the Greek-, Enlightenment-, and contemporary periods understand phenomena (such as force, substance, motion, and acceleration for example) are incommensurate because of the difference in the conceptual and terminological frameworks that describe the underlying phenomena in question. Such frameworks seem to not be translatable into each other. Moreover, paradigm shifts

are ‘revolutions’ in which normal science – which Kuhn conceptualizes as a form of mundane puzzle solving – is shaken up by extraordinary science, which takes up new research agendas in relation to anomalies that have turned up within normal science:

In a given scientific field, long periods of conservative, tradition-bound normal science are punctured by an occasional crisis and, still less frequently, by a revolution. Normal science is highly regimented work under a paradigm. It aims to extend and articulate the paradigm, not to test it, for the paradigm *defines* the research tradition, the scientific life, of a particular discipline and its practitioners.

During a crisis period the usual conservative strictures relax somewhat, and truly innovative ideas and practices may emerge as serious alternatives. The repeated failures of the normal scientists to handle the crisis situation, together with the emergence of a promising new approach, may trigger a revolution.

[T]wo competing paradigms are “incommensurable,” meaning, roughly, that they cannot be measured against the same standard. . . . [I]n the more radical passages of *Structure*, he spoke of paradigm changes as akin to Gestalt perceptual switches, religious conversions, and political revolutions, comparisons he later dropped (Nickles 2002: loc 77, emphasis in original).

In order to more accurately appropriate the notion of paradigms into HCI discourses, we can distinguish between ‘hard’ and ‘soft’ understandings of paradigms. A ‘hard’ notion focuses on common dynamics of generational change and upheaval, revolution, accounting for anomalies, emergence of new exemplars and methods to take up unsolved puzzles. Such ‘hard’ essentialization of a paradigm creates discourses that are incommensurate with each other, where epochal and historical progressions in a discipline confine researchers to ‘living in different worlds.’ This scenario seems to be a poor fit with the three waves of HCI, not least because the waves conduct inquiry into very different phenomena, as opposed to studying the same or similar problems through differing and incompatible conceptual frameworks. As the contributions to these two volumes show, many practitioners develop hybrid approaches and technologies bridging across the discursive terrains of the various waves.

We believe that a ‘softer’ conception of paradigms is better situated to fit the domains of HCI discourse. A ‘soft’ understanding emphasizes communities of inquiry and shared exemplars, held together by a fuzzier logic of ‘family resemblances.’ The three waves under this paradigmatic model approximates families of related approaches, examples, puzzles, problems and solutions. Nersessian (2002) applies Eleanor Rosch’s theories of concept formation to Kuhn’s notion of paradigms to articulate the discourse and practices of research communities:

Most of Kuhn’s work after writing *Structure* centered on issues of what he called the scientific “lexicon,” specifically, on how the language of a scientific community is acquired and how language changes relates to incommensurability.

What one acquires in learning a conceptual structure are not sets of defining characteristics and specifiable rules for the concepts that participate in the problem exemplars comprised by the paradigm. Rather, one acquires sets of “family resemblances” that include both similarities and differences amongst instances.

[R]esearch on categorization in cognitive psychology begun in the early 1970s by the psychologist Eleanor Rosch and her collaborators provides a cognitive underpinning for many of Kuhn's intuitive insights about concept representation and acquisition.

[R]ather than representing concepts by sets of defining criteria, humans represent both natural and artificial concepts by a prototypical example. Category membership is determined by similarity and dissimilarity to the features of the prototype.

Further, concepts show graded structures. That is, some instances of a given concept are better examples of the concept than other instances. (loc 2622).

Nersessian's approach seems to form a better match for the situation of paradigms in contemporary HCI fields. Understood in this 'softer' manner of graded category membership and family resemblance, the difference between HCI discourses and practices takes on a more recognizable outline. These two volumes can thus be understood as a way of organizing the family resemblances of third wave HCI across rich application and methodological domains – at once highly different from each other, yet recognizably belonging together in their distinctive differences from first and second wave approaches.

### 1.3 Theoretic Integration

In preparing for a CHI 2015 panel on transdisciplinary design [...], I was asked if a fourth wave is coming. My best answer is that HCI is in the middle of a chaos of multiplicity in terms of technologies, use situations, methods, and concepts. Hopefully something lies beyond that horizon, but for now, I'll leave it to others to identify it (Bødker 2015).

While it is not the direct goal of these volumes to point the way to a fourth wave, it is possible to see some paths emerging for what this might look like – especially if we note the global commonalities in the distinctions between the waves, and take a softer or fuzzier family resemblance stance toward category membership of such vast research terrains. Niklas Luhmann's systems theory could serve as a basis for a more integrationist positioning amongst the divergent academic cultures and exemplary problem-types of contemporary HCI.

Luhmann's systems theory transcends Mind-Body dualisms (and by extension, traditional subjective/objective dichotomies) by introducing a third term – Communication – into the conceptual mix. Appropriating Varela and Maturana's concept of autopoiesis, Luhmann understood (1) Mind, (2) Body, and (3) Communication as separate and distinctive autopoietic systems in structural couplings to each other and to their environments. Two minds in close physical proximity, for instance, are operationally closed to each other – this is demonstrated by the absence of telepathic effects. A third autopoietic system – that of Communication – is needed in order to achieve information transfer between them. For Luhmann, minds don't communicate, only communication communicates (this is a function of its being an operationally closed, autopoietic system). He considered his theory to be a 'super theory' because it included itself in itself, as a theory of making distinctions generally, conceiving of communication, cognition, and bodies as systems that

are always making self-other distinctions between their own operational closure and their environments.

Luhmann's systems conception aligns strongly with the three waves of human-computer interaction as he understands technologies to be in the environment of living systems. Within such conceptualization, the first wave of ergonomics-oriented approaches corresponds to structural couplings of technology to the Body, while second wave information processing models address the cognitive capacities of Mind. The third wave's central focus on meaning-making completes the mapping to Communication as its own autopoietic process. We will ground this somewhat abstract discussion in a concrete example by referring to Veronika Tzankova's current research in interactive sports technologies. Her research involves the development of new technical systems for horse riders to improve their overall performance in this contact sport. The successful operation of such systems involves considerations at several levels that closely correspond to the three waves associated with HCI. First, the system requires to be physically constructed which engages technical and ergonomic concerns – such as physical design of the equipment, posture of rider, and kinesiological characteristics of the horse. This level corresponds to problem conceptualization characteristic of first wave HCI. Second, the design of the system should take into account cognitivist considerations – e.g. not distracting the rider through misallocation of limited attentional resources – problematics essential to second wave HCI. Last, the system should effectively communicate to the rider by providing feedback that makes sense – facilitating interspecies communication between technology, horse, and rider through embodied interactions. This level of 'meaning making' is a distinct theme of third wave HCI. A system such as this – especially coming from a sports context where all three levels are vital to the safety and security of the sportsperson engaged – exemplifies the growing necessity of a research agenda that integrates all three HCI waves through discursive and practical variations based upon Luhmann's three autopoietic systems categories of Body, Mind, and Communication.

It is not just that Luhmann's theory logically maps to the typologies of the three HCI waves as explicated by others, but rather actually provides the only model available amongst the major theorists for imagining a possible convergence of all HCI discourses and practices. Instead of academic tribes of subspecialists narrowly concerned with their own local and preferred exemplars and problem-solution spaces, the Mind-Body-Communication matrix could point to a fourth wave of 'integrationist' agendas that at this point we can offer as a speculative gesture on our part. This goes somewhat further than Bødker's discussion "When second wave HCI meets third wave challenges" (2006) by suggesting that even first wave HCI might have potential for reintegration with the new domains and methods presented by the third wave.

While not usually grouped together as a set of related intellectual movements, systems theory shares a common origin with phenomenology and pragmatism in the development of new concepts in an attempt of transgressing Enlightenment binary positions. Just as thinkers like Husserl, Merleau-Ponty, James, and Dewey sought a way out of the traditional Empiricism vs. Idealism philosophical impasses,



Bertalanffy's original formulation of general systems theory served as a way of moving beyond Determinism and Vitalism as explanatory frameworks for understanding organized wholes of self-interacting elements. Taken together, phenomenology, systems theory, and pragmatism can be understood broadly as 'third way' approaches that move beyond reductive-causal concepts on the one side, and ideal-spiritual explanations on the other, within an all-encompassing consideration of subjective, objective, and praxeological phenomena.

## 1.4 Trading Zones and Interactional Expertise

The third wave has generated perhaps the greatest expansion in the disciplinary interactions of HCI with other fields, and can be broadly understood as a trading zone with humanist and social science theories:

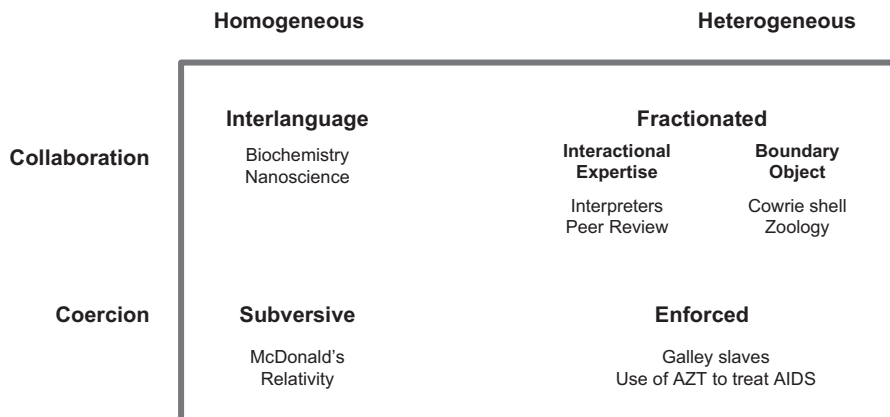
Two groups can agree on rules of exchange even if they ascribe utterly different significance to the objects being exchanged; they may even disagree on the meaning of the exchange process itself. Nonetheless, the trading partners can hammer out a local coordination, despite vast global differences. In an even more sophisticated way, cultures in interaction frequently establish contact languages, systems of discourse that can vary from the most function-specific jargons, through semi-specific pidgins, to full-fledged creoles rich enough to support activities as complex as poetry and metalinguistic reflection. (Galison 1997: 783)

A trading zone can gradually become a new area of expertise, facilitated by interactional expertise and involving negotiations over boundary objects (objects represented in different ways by different participants). (Gorman 2010)

Third wave HCI has proposed in a sense a 'double condition' of negotiating trading zone inquiry with other areas of HCI research, together with scholarly domains far beyond HCI. Collins et al. (2010) have modelled trading zone inquiry into quadrants defined by the axes Homogeneity-Heterogeneity and Collaboration-Coercion as follows (Fig. 1.1):

Interlanguage trading zones operate by developing new cultural tools, subversive trading zones operate by imposing one culture on another, while enforced trading zones operate with almost no cultural interchange. The final type of trading zone, which occupies the top right-hand area of the table, involves fractions of cultures as the medium of interchange. There are two kinds of fractionated trading zones: boundary object trading zones, which are mediated by material culture largely in the absence of linguistic interchange, and interactional expertise trading zones, which are mediated by language largely in the absence of the material (loc 169).

HCI clearly has both sides covered in the Fractionated quadrant, being a research practice typically organized around the development of new technical designs, while also being a subject of academic discourse. Where the material culture aspect is perhaps most foregrounded is in the appropriation of forms such as artworks, critique, or various communications media, where content and connotation considerations may take on as much interest as usability. Since third wave HCI has as an orienting feature a concern with meaning making, entertainment, aesthetic



**Fig. 1.1** Collins et al. general model of trading zones (As presented in Gorman 2010: Fig. 2.1)

experiences, culture forming, style trends, or rhetorical arguments, for instance, can take on an increasing role in investigations and research output. Considering these volumes as a whole in relation to the general model of trading zones, third wave HCI seems ‘squarely’ (no pun intended!) in the Collaboration-Heterogeneous quadrant. Our framing of Luhmann’s ‘super theory’ as a method for integrating all HCI waves could in trading zone terms be understood as a convergence toward the top left Interlanguage quadrant. This quadrant is also understood as the end-phase of trading zone development:

Thus biochemistry, though it grew up as a trading zone, is now just a new homogenous cultural location in which trades happen. When they reach their end points, all the examples in the left-hand areas slip off the table in the westerly direction, as it were (loc 210).

It will remain to be seen of course whether HCI continues along its current path of increasing divergence and plurality of approaches, or whether new lines of convergence may start to draw the different strands together. Regardless of the course of development, the understanding of future trends necessarily depends on our thorough understanding of current affairs. Thus, the objective of *New Directions in 3rd Wave HCI* is to position present and emerging trends shaping the field of human-computer interaction both in terms of (1) technological dynamics (Volume I), and (2) systemic practices of study (Volume 2). As most individuals interact with technology routinely for extensive periods of time (File and Ryan 2014), it is important to understand the experiential dimension of HCI as a source of knowledge and design.

To address these issues, *Volume 1 – Technologies* focuses on the conceptualization and documentation of contemporary third wave HCI. It presents key developments at the leading edge of human computer interactions by providing reflective insights on the theoretic and practical conceptualization, valuation, and development of contemporary technologies. By doing so, this compilation of essays serves as a resource for understanding human-computer interaction through a multiplicity

of interdisciplinary perspectives that can facilitate the systematic epistemological shaping (and reshaping) of technological design and production practices. The combination of perspectives from the humanities and social sciences emphasize the importance of human and experiential dimensions within HCI and contribute to the better conceptualization of the challenges and opportunities that arise as a result of the rapid development and impact of technological progress. Transcending the task-orientedness characteristic of earlier HCI research, *Volume 1: Technologies* covers areas related to artificial intelligence, machine learning, metacreation, 3D printing, critical making, sensorial computing, physical computing, the internet of things, virtual reality, multimodal display, sonification and language technologies, within a frame of experiential inquiry. Drawing on the vast interdisciplinary expertise of the contributors, this volume investigates the experiential and expressive dimension essential to the positive progress of the field of HCI.

Designed to introduce the central themes of research design approaches, *Volume 2 – Methodologies* focuses on latest practices and conceptualizations of the systematic study of HCI. The volume introduces new methodological approaches – often situated in practical case-studies – that integrate human and experiential inquiry within the study of human-computer interactions. Its objective is to identify and address methodological challenges specific to third wave HCI and to propose research approaches embedded within phenomenological, experiential, and expressive modes of investigation. We also hope that the systematization of ‘third wave’ approaches to the study of HCI can serve further as a platform that invites ideas and ‘ways of knowing’ from different epistemological domains into ongoing design practices and applications. This volume integrates diverse research methods, ideas, and perspectives with the aim to highlight and integrate relevant – but often segregated – expertise from the arts, design, social sciences, and the humanities. The application of methodological approaches specific to the particularities of third wave HCI is essential to the development of new, effective, usable *and meaningful* technologies. *Volume 2: Methodologies* covers methodological approaches grounded in autoethnography, empathy-based design, crowdsourcing, psychometrics, user engagement, speculative design, peripheral practices, somatics, embodied cognition and transdisciplinarity. In addition to facilitating inquiry into the design of new technologies, this survey of approaches aims to encourage researchers and designers of technology to critically examine the gamut of processes involved in the production of contemporary technologies.

## References

- Bannon L (1986) From human factors to human actors: the role of psychology and human-computer interaction studies in system design. In: Kyng M, Greenbaum J (eds) *Design at work: cooperative design of computer systems*. Erlbaum, Hillsdale, pp 25–44
- Bødker S (2006) When second wave HCI meets third wave challenges. Keynote paper NordiCHI’06. In: *Proceedings of the 4th Nordic conference on human-computer interaction: changing roles NordiCHI 2006*. ACM, Oslo

- Bødker S (2015) Third-Wave HCI 10 years later— participating and sharing. *Interactions* Sept-Oct 2015. <https://doi.org/10.1145/2804405>
- Collins H, Evans R, Gorman ME (2010) Trading zones and interactional expertise. In: Gorman ME (ed) *Trading zones and interactional expertise: creating new kinds of collaboration*. MIT Press, Cambridge MA
- File T, Ryan C (2014) Computer and Internet use in the United States: 2013, American community survey reports. <https://www.census.gov/history/pdf/2013computeruse.pdf>. Accessed 15 Jan 2018
- Galison P (1997) *Image & logic: A material culture of microphysics*. The University of Chicago Press, Chicago
- Gorman ME (2010) *Trading zones and interactional expertise: creating new kinds of collaboration*. MIT Press, Cambridge, MA. Amazon Kindle edition. Retrieved from: <https://mitpress.mit.edu/books/trading-zones-and-interactional-expertise>
- Grudin J (2005) Three faces of human-computer interaction. *IEEE Ann Hist Comput* 27(4):46–62
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: *Proceedings of the conference on human factors in computing systems (CHI 2007)*. San Jose, pp 1–18
- Nersessian NJ (2002) Kuhn, conceptual change, and cognitive science. In: Thomas Kuhn (*Contemporary philosophy in focus*). Cambridge University Press, Cambridge. Amazon Kindle edition
- Nickles T (2002) Introduction. In: Thomas Kuhn (*Contemporary philosophy in focus*). Cambridge University Press, Cambridge. Amazon Kindle edition

# Chapter 2

## The Relational Turn: Third Wave HCI and Phenomenology



David J. Gunkel

**Abstract** Third wave HCI (Human Computer Interaction) proposes an innovative method for framing human computer interactions by putting emphasis on the terms and conditions of the interactive relationship prior to determinations concerning the human subject and its computational object. As promising as this “relational turn” appears to be, there are important theoretical, epistemological, and axiological challenges that remain and need to be addressed. This chapter takes up and investigates a number of these open questions regarding third wave HCI. It begins by briefly reconsidering the three waves or paradigms of HCI research and demonstrating how what appears last in the numbered sequence, the third wave, is actually older and “more original” than it initially appears to be. It then examines the opportunities and challenges of the phenomenological commitment that is operationalized in third wave HCI. And it concludes by identifying and outlining the consequences of this innovation for current and future research efforts.

Standard methods for conducting HCI research typically assume and operationalize a subject/object dichotomy. Formulated in this way, the characteristics and features of interaction are considered to be a subsequent result of the two interacting components: the human subject and the computational object. Third wave HCI proposes to flip the script on this transaction, putting emphasis on the interaction before and in advance of determinations concerning the subject and object of the relationship. As promising as this change in perspective might sound for altering the way we address and investigate human/machine interactions, there remains important theoretical, epistemological, and axiological challenges to this modification that need to be addressed and formalized.

The task of the following is pursue these open questions in an effort to account for the full potential and consequences of third wave HCI. This objective will be

---

D. J. Gunkel (✉)  
Northern Illinois University, DeKalb, IL, USA  
e-mail: [dgunkel@niu.edu](mailto:dgunkel@niu.edu)

pursued by capitalizing on innovations already available in other fields, especially the phenomenological tradition in continental philosophy and recent developments in moral epistemology, which have already gone through and experienced a similar “relational turn” (Coeckelbergh 2012; Gunkel 2012). Toward this end, the analysis that follows will proceed by way of three steps or movements. The first will briefly reconsider the three waves or paradigms of HCI research and demonstrate how what appears last in the numbered sequence, the third wave, is actually older and “more original” than it initially appears to be. The second section will then examine the opportunities and challenges of the phenomenological commitment that is operationalized in third wave HCI. And the third and final section will identify and outline the consequences of this innovation for current and future research.

## 2.1 Paradigms of HCI Research

Scholars currently recognize three intellectual waves of HCI research. Steve Harrison, Phoebe Sengers and Deborah Tatar (2007), following the influential work of Thomas Kuhn (1996), have characterized these as three unique paradigms. “Central to each paradigm in HCI,” they argue, “is a different metaphor of interaction. Each such metaphor introduces ‘centers’ and ‘margins’ that drive choices about what methods are appropriate for studying and designing interaction and for how knowledge claims about interaction can be validated” (Harrison et al. 2007: 4). This adaptation of Kuhn’s theory (and it is an adaptation insofar as Kuhn’s text does not utilize the terms “center” and “margin”) is then utilized to identify and explicate three different and competing approaches to the study of HCI.

### 2.1.1 *Three Waves*

First wave HCI, which consists of “an amalgam of engineering and human factors, saw interaction as a form of man-machine [SIC] coupling in ways inspired by industrial engineering and ergonomics. The goal of work in this paradigm, then, is to optimize the fit between humans and machines; the questions to be answered focus on identifying problems in coupling and developing pragmatic solutions to them” (Harrison et al. 2007: 4). First wave HCI, therefore, is about human control of computational mechanisms and is concerned with the best way to design input and output affordances to facilitate effective human/machine couplings. Although Harrison, Sengers and Tatar do not state it explicitly, this paradigm is informed by first wave cybernetics, where the controlling issue was control (Wiener 1996).

The second paradigm of HCI research shifts focus from questions of efficient control and ergonomics to computational capability and information processing and transmittal. The second wave, as Harrison, Sengers and Tatar (2007: 4) explain, “is organized around a central metaphor of mind and computer as symmetric, coupled

information processors. At the center is a set of information processing phenomena or issues in computers and users such as ‘how does information get in’, ‘what transformations does it undergo’, ‘how does it go out again,’ ‘how can it be communicated efficiently’ etc.” This second intellectual wave of HCI research focuses attention not on matters of control but on the flow of information into and out of the device and the transformations in data that occur by way of this process. It is, therefore, concerned with “communication” as characterized by Claude Shannon and Warren Weaver (1963) in *The Mathematical Theory of Communication*, namely, how information gets into the device, how it is processed, and how the output is generated and conveyed to the human user by way of various interface applications and features.

Third wave HCI introduces another alteration in focus or what Harrison, Sengers and Tatar describe as a movement to the center of items that had been (in terms of the two previous paradigms) considered marginal. “We are now in a position to define the 3rd paradigm more precisely. It contains a variety of perspectives and approaches whose central metaphor is interaction as phenomenologically situated. The goal for interaction is to support situated action in the world, and the questions that arise revolve around how to complement formalized, computational representations and actions with the rich, complex, and messy situations at hand around them” (Harrison et al. 2007: 9). This significant shift in perspective can be characterized as a kind of inversion of the other two paradigms insofar as it is concerned not with the capabilities or operations of the two interacting components—the human user and the computational artifact—but with the phenomenon of the relationship that is situated between them. Third wave HCI, therefore, emphasizes the terms, conditions, and situation of the interaction and not (at least not primarily) the subject and object of the relationship. In this way, “relations are prior to the things related” (Callicott 1989: 110), instituting what other theorists have called a “relational turn” (Coeckelbergh 2012: 49).

### 2.1.2 *Third Wave Avant La Lettre*

From a third wave perspective, HCI research is framed in such a way that the central matter of concern is the situation and characteristics of the interaction and not the ontological capabilities or features of the two elements that comprise (or are presumed to comprise) the terms of the relationship. Although this shift in focus is presented as a more recent innovation in the lineage and evolution of HCI research (it is the third item in a sequence of intellectual developments or waves), it is a viewpoint that is already available and operationalized in Alan Turing’s agenda-setting paper on artificial intelligence. Or to put it another way, the Turing Test, or what Turing himself calls the “game of imitation,” is third wave HCI *avant la lettre*.

Although Turing begins this essay by proposing to consider the question “Can machines think?” he immediately recognizes persistent and seemingly irresolvable terminological difficulties with the question itself. “I propose,” Turing (1999: 37) writes, “to consider the question, ‘Can machines think?’ This should begin with

definitions of the meaning of the terms ‘machine’ and ‘think.’ The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words ‘machine’ and ‘think’ are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, ‘Can machines think?’ is to be sought in a statistical survey such as a Gallup poll. But this is absurd.” In response to this difficulty—a semantic problem with the very words that would be employed to articulate the question to begin with—Turing proposes to pursue an alternative line of inquiry: “Instead of attempting such a definition,” Turing (1999: 37) continues, “I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words. The new form of the problem can be described in terms of a game which we call the ‘imitation game.’ It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman.” This determination, as Turing explains, is to be made by way of a sequence of questions and answers. The interrogator (C) asks participants A and B various things, and based on their responses tries to discern whether the respondent is a man or a woman. “In order that tone of voice may not help the interrogator,” Turing (1999: 37–38) further stipulates, “the answers should be written, or better still, typewritten. The ideal arrangement is to have a teleprinter communicating between the two rooms.”

In this way, the initial arrangement of the “game of imitation” is, as Turing describes it, predicated on a kind of computer-mediated communication (CMC). The interrogator interacts with two unknown participants via a form of synchronous computer-mediated interaction that we now routinely call “chat.” Because the exchange takes place via text messages routed through the instrumentality of a machine, the interrogator cannot see or otherwise perceive the identity of the two interlocutors and must, therefore, ascertain gender based on responses that are supplied to questions like “Will X please tell me the length of his/her hair” (Turing 1999: 37). Consequently, the identity of the interlocutors is something that is hidden from view and only able to be ascertained by way of the messages that come to be exchanged.

Turing then takes his thought experiment one step further. “We can now ask the question, ‘What will happen when a machine takes the part of A in this game?’ Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, ‘Can machines think?’” (Turing 1999: 38). In other words, if the man (A) in the game of imitation is replaced with a computing machine, would this device be able to respond to questions and “pass” as another person, effectively fooling the interrogator into thinking that it was just another human interlocutor? It is this question, according to Turing, that replaces the initial and unfortunately ambiguous inquiry “Can machines think?” Consequently, if a computer does in fact become capable of successfully simulating a human being, of either gender, in communicative exchange with a human interrogator to such an extent that the



interrogator cannot tell whether s/he is interacting with a machine or another human individual, then that machine would, Turing concludes, need to be considered “intelligent.”

For Turing’s game of imitation, what is of principal importance is what actually transpires in the communicative interaction. The test, therefore, is not an evaluation of the internal capabilities of the interactants per se but of the communicative behavior evidenced in and by the interaction, and human-grade interpersonal conversational interaction in particular. Furthermore, the game of imitation is not really concerned with what kind of information is provided by the interlocutors but with whether the performance of the conversational interaction was believable or not as judged by a human interrogator.

## 2.2 Phenomenology

What distinguishes third wave HCI, therefore, is an epistemological shift from efforts to determine “what something is” to “how it appears to be.” This is precisely why third wave HCI can be described as phenomenological. The concept of phenomenology develops from an important epistemological pivot in modern philosophy. Beginning with Immanuel Kant’s *Critique of Pure Reason* (at least with this work, but there are ways that this entire tradition can and has been traced all the way back to Plato, if not beyond), there is recognition that what something is in-itself needs to be distinguished from how it appears to us, finite human beings, by way of interactions with our senses. Phenomenology, as Peter-Paul Verbeek (2011: 15) explains, names “a philosophical movement that seeks to analyze the relations between human beings and their world rather than as a method for describing reality.” Although different brands of phenomenological thinking (e.g. Hegel, Husserl, Heidegger, Merleau-Ponty, etc.) approach this effort in significantly different ways, the basic structure remains in play—namely, the apparent separation between the knowing subject and the object of knowledge and the need to account for (if not remediate) this seemingly irreducible difference. And despite the fact that the vocabulary is different, this is precisely what Turing had focused on: how different objects (either another human individual or a computer) appear to function in conversational interactions with a human subject, irrespective of what they actually are, which is, according to the stipulations of the game of imitation, always and already hidden from direct view.

### 2.2.1 Epistemological Complications

There is the one important epistemological complication with this procedure, and that complication is already evident in Turing’s parlor game. The Turing test derives a determination of intelligence from the simulation of behavior. It therefore makes

a decision concerning what is from how it appears to be. This is precisely what is targeted and critiqued by the philosopher John Searle in his Chinese Room thought experiment.

Imagine a native English speaker who knows no Chinese locked in a room full of boxes of Chinese symbols (a database) together with a book of instructions for manipulating the symbols (the program). Imagine that people outside the room send in other Chinese symbols which, unknown to the person in the room, are questions in Chinese (the input). And imagine that by following the instructions in the program the man in the room is able to pass out Chinese symbols which are correct answers to the questions (the output). The program enables the person in the room to pass the Turing Test for understanding Chinese but he does not understand a word of Chinese (Searle 1999: 115).

The point of Searle's imaginative albeit somewhat ethnocentric illustration ("ethnocentric" insofar as Chinese has always constituted the "other" of European philosophy since at least the time of Leibniz) is quite simple—simulation is not the real thing. "The Turing test," as Searle (1999: 115) concludes, "fails to distinguish real mental capacities from simulations of those capacities. Simulation is not duplication." In other words, merely shifting verbal symbols around in a way that looks like linguistic understanding is not really an understanding of the language. A computer, as Terry Winograd (1990: 187) explains, does not really understand the linguistic tokens it processes; it merely "manipulates symbols without respect to their interpretation." Or, as Searle (1984: 34) characterizes it, computers have syntax, a method of symbol manipulation, but they do not have semantics.

The important question is whether this kind of simulation is a useful social fiction, i.e. a kind of "game" that has its utility (as it does for Turing), or whether it is an inherently deceptive practice that should be tightly controlled, if not actively constrained? In response to this question, there have been two kinds of answers. For Sherry Turkle, this pretense is a significant problem: "I find people willing to seriously consider robots not only as pets but as potential friends, confidants, and even romantic partners. We don't seem to care what their artificial intelligences 'know' or 'understand' of the human moments we might 'share' with them...The performance of connection seems connection enough" (Turkle 2011: 9). According to Turkle's diagnosis, users of emerging technology are in danger of substituting the technological interface for the genuine face-to-face encounters we used to have with other human beings. "Technology," she explains, "is seductive when what it offers meets our human vulnerabilities. And as it turns out, we are very vulnerable indeed. We are lonely but fearful of intimacy. Digital connections and the sociable robot may offer the illusion of companionship without the demands of friendship" (Turkle 2011: 1).

In an effort to restrict or at least protect users from this apparently dangerous form of deception, the "Principles of Robotics" (Boden et al. 2017: 127) stipulates the need for transparency. "Robots are manufactured artefacts. They should not be designed in a deceptive way to exploit vulnerable users; instead their machine nature should be transparent." In stating this, the authors of the principles are not dogmatic absolutists. They recognize that there may be instances where the appearance of intelligence is part of the game. But they are clear in their specification that

users always, and from the very beginning, have a right to know that this is a game: “Although it is permissible and even sometimes desirable for a robot to sometimes give the impression of real intelligence, anyone who owns or interacts with a robot should be able to find out what it really is and perhaps what it was really manufactured to do” (Boden et al. 2017: 127). What the authors of the principles find objectionable is not simulation (or “deception,” which has negative overtones and denotations) per se but unacknowledged simulation, where the user does not explicitly consent to playing the game.

At the other end of the spectrum, there are other voices, like those of Kate Darling and Tony Prescott, who argue that this proclivity is not necessarily a dangerous ruse that needs to be avoided at all costs but the very condition for possibility of social interaction. “Looking at state of the art technology,” Darling (2012: 1) points out, “our robots are nowhere close to the intelligence and complexity of humans or animals, nor will they reach this stage in the near future. And yet, while it seems far-fetched for a robot’s legal status to differ from that of a toaster, there is already a notable difference in how we interact with certain types of robotic objects.” This happens, Darling continues, principally due to our tendencies to anthropomorphize things by projecting into them cognitive capabilities, emotions, and motivations that do not necessarily exist. Socially interactive robots, in particular, are intentionally designed to leverage and manipulate this predilection. “Social robots,” Darling (2012: 1) explains, “play off of this tendency by mimicking cues that we automatically associate with certain states of mind or feelings. Even in today’s primitive form, this can elicit emotional reactions from people that are similar, for instance, to how we react to animals and to each other.” In other words, how something appears to be—how it operates and acts in real social situations and circumstances—might be more important than what it actually is (or has been assumed to be). For this reason, “we should,” as Prescott (2017: 146) concludes formulating a kind of phenomenological maxim, “take into account how people see robots, for instance, that they may feel themselves as having meaningful and valuable relationships with robots, or they may see robots as having important internal states, such as the capacity to suffer, despite them not having such capacities.”

### ***2.2.2 From Phenomenology to Postphenomenology***

Enabling this debate is a difference between the assumed (im)possibility of a final revelation, where the simulation could be compared to and evaluated against what actually is. This is, in fact, a crucial component of both Turing’s and Searle’s thought experiments. For Turing, the game of imitation is organized around a final exhibition and dramatic revelation. In order for the game to be concluded and for the results to be obtained, the interlocutor needs to be able to look behind the interface in order to see who or what had been doing the talking, i.e. another human person or a computer. For Searle, the point of the Chinese room demonstration—namely, that the simulation of an understanding of the language is not really an

understanding of the language—is only possible insofar as we have privileged access to and can observe the inner workings of the room itself. Without this knowledge, one cannot evaluate the difference that separates simulation from the real thing.

The epistemological necessity and importance of this final reveal is evident in another parlor game, *To Tell the Truth* (Gunkel 2010). This TV game show, which ran intermittently on several U.S. television networks since its premier in the mid-1950's, featured a set of four celebrity panelists who were confronted with a group of three individuals or challengers. Each member of this trio claimed to be a particular individual who had some unusual background, notable life experience, or unique occupation. The panel was charged with interrogating the three challengers and deciding, based on the responses to their questions, which one of the three was actually the person s/he purported to be—who, in effect, was telling the truth. In this exchange, two of the challengers engaged in deliberate deception, answering the questions of the panel by pretending to be someone they were not, while the remaining individual told the truth. The “moment of truth” came at the game's conclusion, when the program's host asked the pivotal question, “Will the real [insert name of the person] please stand up?” at which time one of the three challengers stood. In doing so, this one individual revealed him/herself as the real thing and exposed, by comparison, the other two to be false pretenders and imposters.

The final revelation, therefore, is a component that is necessary for resolving these phenomenological games. But this is where things also get complicated. First, there are situations where the conclusive revelation simply cannot take place for technical reasons. Consider the following example from the early days of online interaction on the Internet. In January of 1996, *Wired* magazine published a rather surprising interview with their self-proclaimed “patron saint,” Marshall McLuhan. The interview was surprising, because at the time it was conducted, McLuhan had been deceased for over a decade. Here's how it happened, as explained in the article's introduction: “About a year ago, someone calling himself Marshall McLuhan began posting anonymously on a popular mailing list called Zone (zone@wired.com). Gary Wolf began a correspondence with the poster via a chain of anonymous remailers” (Wolf 1996: 1). So with whom (or what) was Wolf interacting? Was this “virtual McLuhan” the ghost of Marshall McLuhan, an imposter engaging in a little role playing, or an automated chatter bot programmed with, as Wolf (1996: 1) described it, “an eerie command of McLuhan's life and inimitable perspective”? Technically there was no way to answer this question. The interviewer was limited to what had appeared online and, because the exchange took place through the instrumentality of anonymous remailers, was unable to get behind the screen to ascertain the real thing as such. In the face of this dilemma, *Wired* did something that was, from the perspective of accepted journalistic practices, either “embarrassingly wrongheaded and pretentious” (Morrison 2006: 5) or incredibly innovative and inventive. Instead of writing off the whole affair as ultimately unverifiable, the editors decided to publish the interview as is, leaving the question about the true status of the real thing-in-itself open ended and unresolved. This approach recognizes the inaccessibility of the thing as it is in-itself and the need to tarry with and

make decisions based on nothing more than what appears in and by the interaction.

Second, there are circumstances where the revelation simply does not make a difference; where the appearance trumps knowledge of what actually is. This phenomenon had been initially demonstrated and theorized by Byron Reeves and Clifford Nass (1996) in the computer as social actor (CASA) studies. “Computers, in the way that they communicate, instruct, and take turns interacting, are close enough to human that they encourage social responses. The encouragement necessary for such a reaction need not be much. As long as there are some behaviors that suggest a social presence, people will respond accordingly... Consequently, any medium that is close enough will get human treatment, even though people know it’s foolish and even though they likely will deny it afterwards” (Reeves and Nass 1996: 22). The CASA model, which was developed in response to numerous experiments with human subjects, describes how users of computers, irrespective of the actual intelligence possessed (or not) by the machine, tend to respond to the technology as another socially aware and interactive subject. In other words, even when experienced users know quite well that they are engaged with using a machine, they make, what Reeves and Nass (1996: 22) call, the “conservative error” and tend to respond to it in ways that afford this other thing social standing on par with another human individual. Consequently, in order for something to be recognized and treated as another social actor, “it is not necessary,” as Reeves and Nass (1996: 28) conclude, “to have artificial intelligence” strictly speaking. All that is needed is that they appear to be “close enough” to encourage some kind of social response.

This behavior is not limited to sophisticated social robots that are designed to elicit this kind of response. We appear to be able to do it with just about any old mechanism that has some kind of social presence, like the very industrial-looking EOD (Explosive Ordnance Disposal) robots that are being utilized on the battlefield. As Peter W. Singer (2009: 338) and Joel Garreau (2007) have reported, soldiers form surprisingly close personal bonds with their units’ EODs, giving them names, awarding them battlefield promotions, risking their own lives to protect that of the robot, and even mourning their death. This happens, Singer explains, as a product of the way the mechanism is situated within the unit and the role that it plays in battlefield operations. And it happens in direct opposition to accurate data concerning the actual facts of the device in question: They are just dumb technologies that feel nothing.

Third, there is a more sophisticated and empirically grounded articulation of phenomenology that can respond to and explain these results. Though it is rarely identified with the philosophical traditions of phenomenology, this is something that is already in play with the other minds problem. “How does one determine,” as Paul Churchland (1999: 67) characterized it, “whether something other than oneself—an alien creature, a sophisticated robot, a socially active computer, or even another human—is really a thinking, feeling, conscious being; rather than, for example, an unconscious automaton whose behavior arises from something other than genuine mental states?” This problem, at least in its modern form, is often attributed to Rene Descartes, who argues that he can only be certain of his own

mind—*cogito ergo sum*—but cannot be so sure about the mental state of the other entities he sees on the street and who interact with him. Because the knowing subject cannot ascertain—not with the kind of certitude that is often required for empirical knowledge—whether another entity possesses or does not possess a conscious mind, all that can be done is to interact with it and derive an assumption about “the mind of the other” from an experience of the interaction.

Like Turing’s game of imitation, who or what the other actually is in-itself is information that is hidden from view. We are only able to make a conjecture based on the interactive behavior that is evident and observable. The temporal sequence involved with this inference is important and noteworthy. In interactions with other entities (whether human, computer, or otherwise), we infer the presence of various cognitive capabilities based on the externally observable behaviors they exhibit. In other words, we project a consciousness into the other. But then we reverse the direction of the vector, making an assumption that the derived result—the projection of conscious thinking into or onto the other—had been the original cause of the externally observed behaviors. Slavoj Žižek (2008a: 209) identifies the curious temporality of this operation—whereby an effect is posited as the original cause of that from which it is derived—with the neologism “retroactively (presup)posited.” This formulation provides for a more radical mode of phenomenology, something that Verbeek, following Don Ihde, calls postphenomenology.

The postphenomenological approach makes it possible to move beyond the modernist subject-object dichotomy in two distinct ways. First of all, Ihde shows the necessity of thinking in terms of human-technology associations rather than approaching human subjects and technological objects as separate entities... Second, human-world relationships should not be seen as relations between preexisting subjects who perceive and act upon a preexisting world of objects, but rather as sites where both the objectivity of the world and the subjectivity of those who are experiencing it and existing in it are constituted. What the world “is” and what subjects “are” arise from the interplay (Verbeek 2011: 15).

## 2.3 Conclusions

HCI research consists and can be organized in terms of three different paradigms or intellectual waves. Each paradigm focuses research efforts on certain questions and problematics while pushing to the margin other issues and concerns that do not fit that particular frame of reference. In the end, therefore, the one question that remains to be answered appears to be this: Which paradigm is (more) correct? This question, however, is already a problem. Its mode of inquiry is formulated in terms of a particular frame of reference (or paradigm) that operationalize a set of epistemological commitments that are (or at least should be) already in question. In response to this problem—this question concerning the question—we can take note of three important consequences by way of conclusion.

### 2.3.1 *Competing Paradigms*

In the face of competing paradigms, the trick is not a matter of selecting one or the other and staking a claim to it, but of learning how to recognize which paradigm has been operationalized and how it simultaneously enables and forecloses what can be asked about and investigated. For instance, it is now common for users to say “thank you” to their digital assistants and speech dialogue systems (SDS), like Amazon’s Echo/Alexa, Google Home, and Apple’s Siri. Each HCI paradigm frames a different way of conceptualizing and evaluating this phenomenon. From a first wave perspective, saying “thank you” to an SDS does not appear to have any noticeable impact on the control of the device. When looked at through the lens of first wave HCI, this expression of gratitude could be criticized as unnecessary, superfluous, or both. From a second wave perspective, saying “thank you” to a computational mechanism does not seem to provide any additional input that could be processed by the SDS object. It would be a kind of social “noise” that is ultimately unimportant to the exchange and processing of information. From a third wave perspective, however, one can begin to perceive how this seemingly superfluous and noisy performance is part and parcel of the social milieu. Following what had been discovered in the CASA studies, human users extend social standing to computers not because they are (or can be known to be) intelligent and conscious beings, but because they occupy a social role and function. There is a significant co-creation of social presence in the simple act of saying thank you to Alexa or Siri, and third wave HCI allows for us to see how this functions, why it is important, and what impact it has on human sociality.

### 2.3.2 *Speculative Science*

Following from this, HCI needs to become a “speculative science.” For a phenomenological theorist, like G. W. F. Hegel (1969), “speculative” is not, as is often the case in colloquial usage, a pejorative term meaning groundless consideration or idle review of something that is often inconclusive and indeterminate. Instead, Hegel understands and utilizes the word “speculative” in its strict etymological sense, which is derived from the Latin noun *speculum*, meaning mirror or reflector. “Speculative,” therefore, designates a form of self-reflective knowing. According to Slavoj Žižek, it designates an epistemology that explicitly recognizes the way that what comes to be known is always and already conditioned by the situation or condition of knowing. “At the level of positive knowledge,” Žižek (2008b: 3) writes, “it is, of course, never possible to (be sure that we have) attain(ed) the truth—one can only endlessly approach it, because language is ultimately self-referential, there is no way to draw a definitive line of separation between sophism, sophistic exercises, and Truth itself (this is Plato’s problem). Lacan’s wager is here the Pascalian one: the wager of Truth. But how? Not by running after ‘objective’ truth, but by holding

onto the truth about the position from which one speaks.” The strategic advantage of this particular approach (an approach that Verbeek and Ihde would call “postphenomenological”) is not that it provides one with privileged and immediate access to the real thing in its raw or naked state but that it continually conceptualizes the place from which one claims to know anything and submits to investigation the particular position that is occupied by any knowledge-claim whatsoever.

### 2.3.3 *Social and Ethical Consequences*

Finally there are social and moral consequences to this way of thinking and conducting research. Once it is recognized that knowledge production is the product of epistemological paradigms and that there are competing paradigms that frame different ways of knowing, one might be tempted to ask which one or ones are correct or true. Typically responses to this question pull in two opposite and ultimately unsatisfactory directions—democratism and totalitarianism. “Both liberal-political democracy and ‘totalitarianism.’” Žižek (2002: 176) writes, “foreclose a politics of truth. Democracy, of course, is the reign of sophists: there are only opinions; any reference by a political agent to some ultimate truth is denounced as ‘totalitarian.’ What ‘totalitarianism’ regimes impose, however, is also a mere semblance of truth: an arbitrary Teaching whose function is simply to legitimize the pragmatic decisions of the Rulers.” Both democratism and totalitarianism attempt to respond to and take responsibility for competing paradigms. One does so by saying that anything that appears to anyone is acceptable and true; the other by making what is ultimately an arbitrary decision and imposing a form of orthodoxy.

What is important here is not what makes these two extreme positions different. What is important is what they share in common. Both democratization and totalitarianism are devised in an effort to contend with the perceived threat of relativism—“the claim that no universally valid beliefs or values exist” (Ess 1996: 204). But as I have argued elsewhere (Gunkel 2010, 2012) “relative,” which has an entirely different pedigree in a discipline like physics, need not be construed negatively and decried, as Žižek (2006: 281) has often done, as the epitome of postmodern multiculturalism run amok. Robert Scott (1967), for instance, understands “relativism” to be a positive rather than negative term: “Relativism, supposedly, means a standardless society, or at least a maze of differing standards, and thus a cacophony of disparate, and likely selfish, interests. Rather than a standardless society, which is the same as saying no society at all, relativism indicates circumstances in which standards have to be established cooperatively and renewed repeatedly” (Scott 1967: 264).

Charles Ess (2009: 21) calls this alternative “ethical pluralism.” “Pluralism stands as a third possibility—one that is something of a middle ground between absolutism and relativism... Ethical pluralism requires us to think in a ‘both/and’ sort of way, as it conjoins both shared norms and their diverse interpretations and applications in different cultures, times, and places” (Ess 2009: 21–22). Likewise



Luciano Floridi (2013: 32) advocates a “pluralism without endorsing relativism,” calling this “middle ground” relationalism: “When I criticize a position as relativistic, or when I object to relativism, I do not mean to equate such positions to non-absolutist, as if there were only two alternatives, e.g. as if either moral values were absolute or relative, or truths were either absolute or relative. The method of abstraction enables one to avoid exactly such a false dichotomy, by showing that subjectivist position, for example, need not be relativistic, but only relational” (ibid.). Like Žižek, Floridi recognizes that truth can be neither totalitarian nor completely democratized such that “anything goes.” It is always formulated and operationalized from a particular position of “enunciation” (Žižek’s Lacanian inspired terminology) or what Floridi calls “level of abstraction,” which is dynamic and alterable. The task of responsible research, therefore, is to learn how to take responsibility for these necessary alterations in perspective and their social and moral consequences.

## References

- Boden M, Bryson J, Caldwell D, Dautenhahn K et al (2017) Principles of robotics: regulating robots in the real world. *Connect Sci* 29(2):124–129. <https://doi.org/10.1080/09540091.2016.1271400>
- Callicott, JB (1989) *In defense of the land ethic: essays in environmental philosophy*. State University of New York Press, Albany
- Churchland PM (1999) *Matter and consciousness*. MIT Press, Cambridge, MA
- Coeckelbergh M (2012) *Growing moral relations: critique of moral status ascription*. Palgrave Macmillan, New York
- Darling K (2012) Extending legal protection to social robots. *IEEE Spectrum*. <http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/extending-legal-protection-to-social-robots>. Accessed 8 Dec 2017
- Ess C (1996) The political computer: democracy, CMC, and Habermas. In: Ess C (ed) *Philosophical perspectives on computer-mediated communication*. SUNY Press, Albany, pp 196–230
- Ess C (2009) *Digital media ethics*. Polity, Cambridge
- Floridi L (2013) *The ethics of information*. Oxford University Press, Oxford
- Garreau J (2007) Bots on the ground: in the field of battle (or even above it), robots are a soldier’s best friend. *Washington Post*. <http://www.washingtonpost.com/wp-dyn/content/article/2007/05/05/AR2007050501009.html>. Accessed 8 Dec 2017
- Gunkel DJ (2010) The real problem: avatars, metaphysics and online social interaction. *New Media Soc* 12(1):127–141. <https://doi.org/10.1177/1461444809341443>
- Gunkel DJ (2012) *The machine question: critical perspectives on ai, robots and ethics*. MIT Press, Cambridge, MA
- Harrison S, Sengers D, Tatar D (2007) The three paradigms of HCI. *CHI 2007*, 28 April–3 May, San Jose, CA
- Hegel GWF(1969) *Science of logic* (trans: Miller AV). Humanities Press International, Atlantic Highlands
- Kuhn TS (1996) *The structure of scientific revolutions*. University of Chicago Press, Chicago
- Morrison JC (2006) Marshall McLuhan: no prophet without honor. *AmeriQuests*. <http://ejournals.library.vanderbilt.edu/ojs/index.php/ameriquests/issue/view/4/showToc>. Accessed 8 Dec 2017
- Prescott TJ (2017) Robots are not just tools. *Connect Sci* 29(2):142–149. <https://doi.org/10.1080/09540091.2017.1279125>

- Reeves B, Nass C (1996) *The media equation: how people treat computers, television, and new media like real people and places*. Cambridge University Press, Cambridge, MA
- Scott RL (1967) On viewing rhetoric as epistemic. *Cent States Speech J* 18(1):9–17. <https://doi.org/10.1080/10510976709362856>
- Searle J (1984) *Minds, brains and science*. Harvard University Press, Cambridge, MA
- Searle J (1999) The Chinese room. In: Wilson R, Keil F (eds) *The MIT encyclopedia of the cognitive sciences*. MIT Press, Cambridge, MA, pp 115–116
- Shannon CE, Weaver W (1963) *The mathematical theory of communication*. Illinois Press, Urbana
- Singer P (2009) *Wired for war: the robotics revolution and conflict in the twenty-first century*. Penguin Books, New York
- Turing A (1999) Computing machinery and intelligence. In: Meyer P (ed) *Computer media and communication: a reader*. Oxford University Press, Oxford
- Turkle S (2011) *Alone together: why we expect more from technology and less from each other*. Basic Books, New York
- Verbeek P (2011) *Moralizing technology: understanding and designing the morality of things*. University of Chicago Press, Chicago
- Wiener N (1996) *Cybernetics: or control and communication in the animal and the machine*. MIT Press, Cambridge, MA
- Winograd T (1990) Thinking machines: can there be? Are we? In: Partridge D, Wilks Y (eds) *The foundations of artificial intelligence: a sourcebook*. Cambridge University Press, Cambridge, MA
- Wolf G (1996) Channeling McLuhan. *Wired* 4(1). [http://www.wired.com/wired/archive/4.01/channeling\\_pr.html](http://www.wired.com/wired/archive/4.01/channeling_pr.html). Accessed 8 Dec 2017
- Žižek S (2002) *Revolution at the gates*. Verso, London
- Žižek S (2006) *The parallax view*. The MIT Press, Cambridge, MA
- Žižek S (2008a) For they know not what they do: enjoyment as a political factor. Verso, London
- Žižek S (2008b) *In defense of lost causes*. Verso, London

# Chapter 3

## Giving Form to Smart Objects: Exploring Intelligence as an Interaction Design Material



Marco C. Rozendaal, Maliheh Ghajargar, Gert Pasman, and Mikael Wiberg

**Abstract** Artificial intelligence (AI) has recently been highlighted as a design material in the HCI community. This acknowledgement is a call for interaction designers to consider intelligence as a resource for design. While this view is valid and well-grounded, it brings with it a need to better understand how intelligence as a design material can be used in formgiving practices. This chapter seeks to address this need by suggesting a new approach that integrates AI in the designer's toolkit. This approach considers intelligence as being part of, and expressed through, an object's character, hereby integrating artificial intelligence into a product's form. We describe and discuss this approach by presenting and reflecting on our experiences in a design course where students were asked to give form to intelligent everyday objects in three iterative design cycles. We discuss the implications of our approach and findings within the frame of third wave HCI.

### 3.1 Introduction

The movement that defines new usages and physical settings for computing technologies through intermixing and broadening the boundaries of theories, concepts and design methods, is often referred to as 'third wave HCI' (Bødker 2006). It has been defined as a paradigm that focuses mostly on interaction as a phenomenon, that is going on between humans and machines and that is embodied and

---

M. C. Rozendaal (✉) · G. Pasman  
Delft University of Technology, Delft, The Netherlands  
e-mail: [M.C.Rozendaal@tudelft.nl](mailto:M.C.Rozendaal@tudelft.nl)

M. Ghajargar  
Polytechnic University of Turin, Turin, Italy

M. Wiberg  
Umeå University, Umeå, Sweden

phenomenologically situated (Dourish 2004; Harrison et al. 2007). Third wave HCI is evolving through an approach that encourages user participation in conceptual thinking for designing shared interaction of artifacts (Bødker 2015). Further, it has been outlined that there is a need for methodological investigations in third wave HCI, focusing on artifact ecologies and their shared usage (Bødker and Klokmoose 2012; Bødker 2015).

Some main concerns that have emerged within third wave HCI – such as the value, the meaning, and the place of interaction – have historically played important roles in formgiving practices in design. For instance, Harrison et al. (2007) highlight an explicit focus on value-based design, whereas the evaluation of a ‘good’ design is based on the user and the particular context of use. This focus shifts the evaluation of design of computing artifacts from universally valid approaches to more participatory and value-sensitive design approaches. Further, in third wave HCI, knowledge is subjective, bringing new perspectives into practice. This is in contrast to first and second wave HCI that focus on objective knowledge and its general applicability. We see great potential in studying formgiving practices in third wave HCI, and in particular on how to give form to smart objects.

The emergence of smart objects – being mundane yet intelligent products – impose new challenges for interaction design (Rozendaal 2016). In many products that are labeled as being smart or intelligent, the material expression of this quality appears to be somewhat unattended. For example, while voice-controlled smart speakers, such as the Amazon Echo (2017), the Google Home (2017), and the Apple Homepod (2017) are capable of performing actions that qualify as intelligent, hardly anything in their material form is expressing this capacity. As a consequence, a product’s intelligence might be experienced to be ‘stuck on’ as an added feature rather than being integrated into a product’s form. Understanding and developing intelligence as expressed in the form of smart objects is important to tackle the challenges related to our interactions with them. This resonates well with what Krippendorff highlights in his definition of product semantics in relation to human interfaces (1989, 2006). It is about our experiences and interactions with objects and environments that renders things understandable, meaningful, transparent, alive and usable (1989: 4–5).

Within the last decade, materiality and material aspects of computational objects have received increased attention within the HCI and design communities. Under the umbrella of the *material turn* in HCI (Wiberg 2014, 2018; Wiberg et al. 2013; Robles and Wiberg 2010; Vallgård and Redström 2007; Rosner 2012; Dourish and Mazmanian 2011), this branch of research spans from exploring formgiving of computational objects and materials, to new theoretical understandings and approaches to material interaction, and to practical oriented and craft-based works. Further, ‘intelligence’ itself has most recently been highlighted as a design material in HCI (Holmquist 2017; Dove et al. 2017) stating that in the near future, intelligence will become available as a resource to be used by non-experts. Therefore, the opportunities and limitations of such material must be carefully considered in the design process, similar to any other design material.

The core aim of this paper is to contribute to these streams of developments in HCI by advocating how intelligence can be used as an interaction design material,

particularly in formgiving practices. In doing so, we will provide both a theoretical and empirical grounding on how to give form to interactive and smart objects using intelligence as a resource.

We first review related work that centers around understanding Artificial Intelligence (AI) in the context of HCI as it spans from the 60s to contemporary views on smart objects in the Internet of Things (IoT). Then, we will review the history of formgiving practices and its materials in Product Design and HCI. Subsequently, we will report on a design course, where students were asked to give form to smart everyday objects to empower people in specific problematic situations. By reflecting on the design course, we suggest a new methodology that considers intelligence as being part of, and expressed through, an object's character. The paper concludes by discussing lessons learned from our design course as well as discussing implications for formgiving practices in HCI.

## 3.2 Artificial Intelligence in HCI

Since the 60s, intelligence has started to become a matter of concern in HCI. Due to their increasing capabilities and autonomy, interactions with computers changed, as they transformed from being 'tools' to becoming 'agents'. Generally speaking, agents are entities that can function autonomously without human intervention, and that can interact with other agents (including humans). They are further able to log, sense and collect data, perceive their environment, and react to changes in a timely fashion (Wooldridge and Jennings 1995).

Based on the emerging trend at that time, Licklider (1960) proposed a new perspective on human-computer interaction – that he labelled *human-computer symbiosis* – to better understand the new opportunities and challenges when interacting with computers as agents. Licklider highlighted the changing role of computers, as not only performing automated tasks predefined by humans, but also taking part in task-formulation and in collaboratively carrying our complex tasks *as partners*.

Later, in the 90s, when AI was developed in the context of virtual assistants, people continued to think about what it would be like when agents will be able to help us do things and augment our abilities (Norman 1994). The main challenges foreseen by Norman were related to the *social ability* of such agents concerning the ways they interact with other agents and humans. Some of the specific issues he mentioned were people's overblown expectations of what agents can do, and feelings of losing control in relation to the automated actions exhibited by agents. In particular, he addressed issues of privacy, which would become more prominent as agents access personal data and act on this information without human intervention.

Contemporary scholars in HCI have revisited the classical work of Licklider (Farooq and Grudin 2016; Jacucci et al. 2014), proposing that new studies need to be conducted to tackle the issue of increased automation and independence of computers in computer use (Jacucci et al. 2014). Further, this development urges the

HCI community to develop new approaches to the design and evaluation of interactive systems (Farooq and Grudin 2016), thus touching upon many of the issues brought forth by Licklider and Norman.

Now in 2017, we find ourselves in a situation where intelligence moves from AI and virtual assistants, to intelligence that is embedded in smart objects as part of the IoT. Here the new question becomes how intelligent sensing and actuating capabilities can be meaningfully connected to human activities, allowing the object to have awareness and interactivity, which require new technological architectures (Allmendinger and Lombreglia 2005; Kortuem et al. 2010; Sabou 2010). Further, the sustained presence of smart objects in everyday life also requires sensitivity in designing their smartness in accord to the materiality and sociability of everyday life (Giaccardi 2015; Janlert and Stolterman 2017).

### 3.3 Materials in Formgiving Practices

Product design has a long tradition of the engagement with and manipulation of materials (Pevsner 1991; Raizman 2004). At a specific moment in a design process, a transition takes place from a design being an abstract functional description to becoming a concrete design manifestation (Ramduny-Ellis et al. 2010). Muller identifies this moment as the beginning of the *form-creation phase*, “that starts at the moment that any conceptualization about the material form emerges, and ends when a definitive design is established” (Muller 2001). In this phase, the ideas of the designer are externalized, explored interactively and represented tentatively in a visual form using a variety of media (McKim 1980). The process is not solely concentrated on determining the material conditions for the fulfillment of the function of the product, but is also aimed at establishing the product’s desired semantic and aesthetic qualities in relationship to its intended experience, meaning and use (e.g. Krippendorff 2006; Bergström et al. 2010; Vallgård 2014).

In product design, a design process involves the use and manipulations of physical materials such as paper, wood, clay, foam or plastics for making mock-ups in the early stages of a design project, and later stages of prototyping may involve electronics such as Arduino’s<sup>tm</sup> and a variety of sensors and actuators. However, these design materials can also include digital ones and media (graphics, sounds, images), as well as data in different types and formats (Crilly et al. 2009). Design materials are often made available to designers in toolboxes, kits, or in other ways that allow designers to *sketch* by easily changing, combining and reconfiguring them (Tholander et al. 2012; Verplank 2009).

Formgiving has been also defined as the ability to identify form attributes that convey sensorial and interactive experiences (Smets et al. 1994). Under the umbrella of *aesthetics of interaction*, formgiving has further evolved in HCI (Petersen et al. 2004; Djajadiningrat et al. 2002) and has been applied in interaction design and HCI strongly influenced by Gibson’s theory of ecological perception and later by prag-

matist philosophy and phenomenology (Ross and Wensveen 2010). Although shaping such materials helps to gain an extensive knowledge and understanding about their behaviors and potentials, current interactive artifacts increasingly include digital materials and information technology, so the form of objects is no longer driven by the technologies within them (Evans and Sommerville 2007) and traditional modes of product's expression of meaning may no longer apply (e.g. Vihma 1995).

Researchers working on the *material turn* in HCI seek to understand the material basis of interactive artifacts in order approach digital and physical aspects of interactive artifacts as a whole (Wiberg 2014, 2018). This branch of research spans from exploring formgiving of computational objects and materials, to new theoretical understandings and approaches to material interaction, and to more practical oriented and craft-based works. Robles and Wiberg (2010) propose *texture* as a concept that unites these physical and digital elements of interactive artifacts from an experience point of view. The term texture denotes how these aspects come together forming a composition of cultural and aesthetic forms of interaction, where tangible and digital aspects are integrated.

Similarly, the concept of *computational composites* as proposed by Vallgård and Redström (2007) considers the interdependency of interactive artifacts as experienced wholes by better understanding three constituting elements: physical form, temporal form and the 'form' of interaction. While a physical form is perceptible through its three dimensional and tangible shape, a temporal form is a pattern of the state-changes that a computer can produce over time. The 'form' of interaction, is a concept that is related to interaction as a dynamically experienced interplay between a human and an artifact, also referred to as the 'interaction gestalt' (Lim et al. 2007). All these three forms are interdependent and should therefore be designed in tandem.

### 3.4 Artificial Intelligence as a Material

The ongoing development and increased possibilities of computational technologies do call on designers to add another design resource and material in their toolkits, which is *intelligence* (Holmqvist 2017; Dove et al. 2017). Designing with materials usually requires particular skills and knowledge that enable designers to work with the material and explore its opportunities and limitations. For example, a study carried out by Dove et al. (2017) presents some of the difficulties that designers face in dealing with Machine Learning (ML) in their practice. These involved having an accurate understanding of the nature and use of ML, how to prototype it, and how to purposefully design with it.

Holmqvist (2017) talks about the challenges involved when integrating AI in the design of User-Interfaces (UI's) and mentions the need to design UI-components that clearly communicate how control is shared between humans and AI, and how


designers could deal with the new learning capabilities of their designs. This recent view on intelligence as a design material discusses some of the required skills and knowledge. For instance, developing an appropriate technical understanding of AI and ML, and ways to program it. While we share this idea, our position in this paper seeks to go beyond understanding and applying AI as technical instrumentation, but rather to understand how intelligence can meaningfully be dealt with in formgiving practices in interaction design and HCI. The question then becomes, what are the critical material and communicative aspects of intelligence when fully integrating it into the product's form *as a whole*?

Building on this understanding of intelligence as a material specifically in relation to formgiving practices, we propose an approach that builds on the work of Janlert and Stolterman on *the character of things* (1997). Janlert and Stolterman introduced the notion of character as an entry point to understand computer artifacts as meaningful wholes. The character of an artifact is defined as the unity of its multiple characteristics, that involves the sustained impressions of aspects of the artifact's function, appearance and manner of behaving, aggregated over time in a complete and coherent way. The character of an artifact helps users to interpret its behavior. For example, the simple movement of stretching out one's arm can be interpreted as either offensive or welcoming, depending on whether the person is perceived to have a hostile or friendly character (Janlert and Stolterman 1997). Thus, actions (i.e. behaviors) of a computer artifact are explained in an *interpretive frame* provided by its character.

Janlert and Stolterman's work also links to notions of *anthropomorphism* and *animism* in understanding the character of computer artifacts by referring to the work of Reeves and Nass (1996) on people's inherent ability to treat computers as social actors and by referring to Brenda Laurel's work on the perceived animism of computer technology (1997). We therefore propose that the overall perceived form of an object, (i.e., its character), should be able to convey information about its function and usage, and also about its intelligence.

We further suggest that in order to give form to the character of smart objects we need to consider their *intent*, *materiality*, and *agency*. Pixar's well-known animated lamp 'Junior' (2017) is a good example to clarify these considerations. The *intent* of a smart object is considered to be shaped by its purpose as a product but also by its motive as an agent. We see Junior as an everyday lamp that acts according to its own beliefs and desires. The *materiality* of a smart object is concerned with its physical and digital structures that enables its expressiveness as an agent and also communicates its functional, aesthetical and symbolic characteristics. For example, Junior's structure, materials and components provide the lamp its ability to adjust light in a specific way. However, these elements are also carefully animated to express its inner life. Finally, the *agency* of a smart object defines the possibilities to interact in the world, both as a tool and as an actor. The hybrid character of Junior (i.e., being a lamp and an agent), makes Junior act in a story in a particular way.





**Fig. 3.1** Process structure of the design process: Conceptual phase (left), Embodiment phase (middle), and Enactment phase (right)

## 3.5 Design Course

To illustrate how smart objects are given form by considering intelligence as part of an object's character, we will describe and reflect on a design course called Interactive Formgiving, a Master elective within the Industrial Design program at Delft University of Technology. As a design brief for the course, students had to develop smart objects as being both everyday products and intelligent agents. Starting from this design brief, the students went through three design cycles (or phases), in which the relations between *intent*, *materiality* and *agency* of the object were explored and defined on increasing levels of detail and integration (Fig. 3.1). The process involved (1) a Conceptual phase, in which the user, situation and object were identified, the character of the object was defined and a first exploration of its behavior was conducted; (2) an Embodiment phase, in which the expressions and behaviors of the everyday object was explored and articulated through its inherent structure and materials; (3) an Enactment phase, in which the collaboration between user and object in the problematic situation was acted out and experienced in context.

To help the reader understand the type of products that resulted from the course, we will first describe three of the designs and highlight their purpose and target group, functionality, and context of use. This is followed by a detailed description of each of the three phases of the design process and an in-depth discussion of how the activities in each phase informed the design of these objects through the articulation of their *intent*, *materiality* and *agency*.

### 3.5.1 Design Cases

The smart objects that were developed had the purpose to empower people in problematic situations. As both being everyday products and intelligent agents, these designs are described as having their own motives and personality. In their functioning, these objects directly and actively intervene in activities that people consider problematic. Below three of these designs are described in more detail. See Fig. 3.2 for a visual impression.



**Fig. 3.2** Image showing Pat the social backpack (left), Harry the power drill (middle) and Waggle the shoe (right)

### 3.5.1.1 Pat the Social Backpack

Pat is a smart backpack specially designed for teenagers that tend to be somewhat insecure and thus have some trouble connecting to others at school. Pat, however, is full of confidence and eager to make new friends, which in Pat's case means connecting to other backpacks. As Pat cannot move around by itself, it needs the support of the teenager. In this way, Pat will empower teenagers in making contact with other people more easily as well: providing them with a 'free-ride'. Pat encourages the teenager to approach others by gently nudging him or her forward and rewarding the teenager when doing so, by gently squeezing the shoulders. However, when the teenager doesn't approach, Pat will show disappointment by sliding of the shoulders of its owner, thereby enticing the teenager to regain control by firmly pulling the backpack back up again. This action reinstalls a sense of confidence in the teenager, thus encouraging him or her to successfully make the next approach.

### 3.5.1.2 Harry the Power Drill

Harry is a smart power drill designed for construction workers that tend to work improperly or even recklessly, which can ultimately lead to broken equipment and accidents. Harry is a true professional who wants to be used properly and doesn't like the idea of overheating or being broken. It is eager to participate with humans when used properly, but becomes reluctant to drill when the drilling angle is slightly off or the drilling pressure is too high. Harry will express this by providing a slightly weaker grip on its handle, thus installing a feeling of weakness and instability while drilling. If Harry notices that its user becomes really reckless in its drilling, it gets angry and will shut down by retracting its drill-point and control handles, to make its owner aware of his inappropriate behavior and give him some time to reflect. After a calm-down period, Harry then comes back to life again and is ready to be used once more.

### 3.5.1.3 Waggle the Shoe

Waggle is a pair of smart running shoes designed for people who run but sometimes cannot find the motivation to hit the track. Waggle acts like an enthusiastic puppy that loves to go outside, which it makes clear by flapping with the flexible sides of the shoes. Once Waggle senses that his owner has noticed its need, it will express its enthusiasm for running by increasing the frequency and amplitude of this flapping motion. Waggle will then make it very easy for its owner to put his running shoes on, by opening and fastening itself. Any actions that might stall the actual running are being counteracted with even more enthusiasm as Waggle starts to get jumpy; physically pushing its owner to move his or her legs. Once back after running, Waggle joins the runner in feeling both proud and tired, by loosening itself up so that the running shoes can be easily kicked off its owner's feet.

## 3.5.2 Methodology: Interactive Formgiving

### 3.5.2.1 Conceptual Phase

Starting from the initial design brief, students first conducted a brainstorm to come up with daily activities that people might need help with. This resulted in three types of users who experienced some kind of difficulties in concrete situations: (1) *teenagers* that struggle to connect with peers at school due to feelings of insecurity, (2) *construction workers* that needed to be urged to work safely, and (3) *runners* who want to go for a run regularly but lack the motivation to do so.

These situations were then further analyzed with the aim to identify an everyday object connected to the problematic situation, that – being made interactive and intelligent – would empower a person to tackle this specific problem. Thus, one

group decided to choose a *backpack* to help teenagers feel more confident. The second group focused on a *power drill* that promotes safe use during the physical act of drilling, rather than providing verbal safety instructions. The third group settled on *running shoes* to motivate runners to go for a run.

Then, students were challenged to frame their everyday object *as a character*. This involved a complex and creative leap, as they needed to characterize the object as a purposeful product and as an intelligent agent. Thus, Pat was conceived to be a *popular friend* who helps the teenager carrying stuff around but also wants to socially connect with other backpacks, while drill Harry was envisioned to be a *professional perfectionist*, who is keen to help you to drill safely and efficiently, but gets annoyed when it is not being used in the proper way. Finally, Waggle, the pair of running shoes, was framed as an *enthusiastic puppy* who motivates people to go for a run by eagerly expressing its own desire to go outside.

After having framed the character of the object, its behavior could now be defined, as being meaningful to its pragmatic function and its motive as an agent. In the course, the number of behaviors was limited to three. For example, Harry could act *dedicated* or *reluctant*, and could even *refuse* to act. Harry's *dedicated* behavior provides the construction worker support while drilling, by securing firm and stable grip controls and fast drill-rotations. However, when the construction worker applies too much physical pressure to the drill, it starts to behave *reluctantly* by loosening its grip controls and hereby becoming less stern. As a consequence, the net-amount of physical force that can be applied to the drill is stabilized, preventing the power drill from overheating and being damaged. Then, when the construction worker applies pressure at unacceptable levels that might damage the power drill, the power drill-controls are drawn-in abruptly. The drill stops and *refuses* to continue.

Pat the backpack was designed with *teasing* behavior to nudge the teenager to approach other backpacks, engaged in *comforting* behavior to motivate the teenager to continue with making the approach, and to show *disappointment* when the teenager turns away. Based on a combination of sensory inputs (heartbeat frequency as an indication of anxiety combined with the proximity of Pat to other backpacks), Pat could deduce whether its owner is approaching or retreating from other backpacks, and if the teenager feels anxious about this or not. According to this data, Pat could then decide to *tease* the teenager, if the anxiety-level is not that high, *comfort* the teenager when he or she feels anxious when making an approach, or might show *disappointment* if the teenager decides to abort. Making a state-transition diagram helped the students to formalize the information these smart objects required to act and provided students with a deeper technical understanding of their intelligence (Fig. 3.3).

### 3.5.2.2 Embodiment Phase

In the embodiment phase, students created physical mock-ups to explore the form and composition of their objects as informed by their product category (i.e., *backpack*, *power drill*, *running shoes*) and by their character (*popular friend*,




Fig. 3.3 State-transition diagram made for Pat the backpack

*professional perfectionist, enthusiastic puppy*). This involved using existing everyday objects and modifying their materials, structures and dynamic properties to enable their expressive capabilities. For example, rigid types of plastics, flexible soft materials such as textiles, and industrial rubber, etc. were applied. Students were discouraged to add elements to the object that felt alien to the product category or that felt ‘out of character’.

Within the boundaries established by these mock-ups, students were then further urged to *animate* these inherent structures and materials through human motion (e.g., *puppeteering*). For this purpose, transparent wires or other mechanisms were used, sometimes supplemented by using the hands directly, enabling students to get a grip on the object’s expressive behavior. For instance, how the backpack could slide down the back of a person in such a way that it showed *disappointment*, or how *reluctance* could be expressed through the power-drill by cleverly manipulating the positioning and experienced firmness of the grip-handles while drilling. For the shoes, this meant discovering how the distinct dynamic product attributes of shoes, such as their flexible leather sides and soles, could express *enthusiasm*.

The next step involved recording these manipulations on video and playing them back to the students in order to better understand and pick-up on the nuances of the expressions and to help students focus on their temporal qualities. For example, how fast should the backpack drop of the shoulders and which trajectory should it follow? Are the movements of the running shoes enthusiastic enough, and can subtle differences between expressions be distinguished from one another?



**Fig. 3.4** Image showing students in enacting the interaction with Pat the social backpack

### 3.5.2.3 Enactment Phase

The goal of this last phase was to understand and develop the agency of the object, as it is situated and experienced in context: what influence does the object have on the situation, how do human and object engage in collaboration, and can we assess if the object is actually empowering its users? In these enactments, the activity and context were staged as informed by the storyboard that was developed previously by the students. One student puppeteered the object, while another student played the role of the end-user, reflecting on the experience of using the object and discussing possibilities for improving the design. Video was used as an observation and communication tool to capture and reflect on the enactment (Fig. 3.4).

During the enactments, the focus was on understanding the agency of the objects. For example, some enactments showed how the expression of *disappointment* of the backpack had an emotional impact on its wearer: when students noticed Pat sliding down their backs, they felt abandoned by it as well as made them look silly. The power drill's *refusal* behavior, as expressed by its sudden retraction of its drill-head and control-grip, felt so sudden that it stirred discussion about how this might cause a dangerous situation in itself, and raised the issue of how the transition from *reluctance* to *refusal* should be made more gradual. The enactment and improvisations with the running shoes also revealed the need for a stronger type of stimulation to motivate the person to go for a run than as previously anticipated.

## 3.6 Discussion

We will discuss how our formgiving approach to smart objects contributes to the work being done on formgiving practices in third wave HCI, and also on how intelligence can be considered as a design material in these practices. Based on our

experiences in the design course, we will then reflect on (1) how students conceptualized smart objects as characters that reconciled their nature of being familiar everyday products and intelligent agents, (2) how students embodied the intelligence of smart objects in authentic ways by using a product's inherent structures and materials as a means to express their intelligence, and (3) how students approached the interaction with smart objects pragmatically (linked to their function as an everyday product) and as a form of social interaction.

### 3.6.1 *Intelligence as Material in Formgiving Practice*

Inspired by formgiving practices in Product Design and HCI, we have explored intelligence, not as a technology on the level of programming or algorithms, but as something that can be shaped and communicated through form. Wiberg's methodology (2014) deals with similar issues. Approaching intelligence *as a material*, allows one to explore intelligence as texture, which is the "formal relation that appeals to the feel, appearance, or consistency of surface or substance" Wiberg and Robles (2010: 68). The concept is useful to describe computational compositions that unite physical, digital and spatial components into an overall experience. This methodology is inspired by design processes that emphasize iteration – working back and forth – between details and wholeness, materials and textures of a computing composition. We have learned from our design course how moving through multiple iterations allowed the object's character to take form, on increasing levels of detail and integration.

Further, a design process using puppeteering and enactments was used to work with intelligence as a material in a so-called *third space*. A third space is created when reality is suspended and explored to provide a productive and creative space at the boundary between the actual and the possible, the *real* and the *fictional* (Halse et al. 2010). These enactments made students aware of the agency of design materials (Schön 1984; Tholander et al. 2012) and provided a reflective conversation with intelligence as a material that is expressive and social. The use of video in particular, allowed students to adopt a more objective view on their work and helped them to investigate behavior in its temporal details (Pasman and Rozendaal 2016).

### 3.6.2 *Intelligence as Character*

In designing smart objects as characters that reconcile their nature of being familiar everyday products and intelligent agents, we argue that aspects like *familiarity*, *authenticity* and *sociability* should be taken into account.

### 3.6.3 *Familiarity*

Designing the intelligence of smart objects as part of their character required students to understand and define intelligence *through* familiar everyday products. In this way, these products *themselves* established the grounding metaphor for their intelligence, and thus, sets the intelligence of smart objects apart from the intelligence of social robots and conversational agents, which take the *human* as their grounding metaphor. Taylor takes a similar approach as he talks about *machine intelligence*, as intelligence that is performed in everyday settings and does not require a human equivalent (Taylor 2009). We argue that framing the intelligence of smart objects within familiar product categories will help people to accurately predict and explain their behavior. Thus preventing them to develop overblown expectations about their intelligence (Norman 1994).

Practice-oriented theories (Kuutti and Bannon 2014) such as activity theory (Kaptelinin and Nardi 2006) consider everyday artifacts to mediate our activities in the world. Knowledge about this relationship can help designers to identify which objects would become empowering given a specific target group and use context (Ghajargar 2017). Further, adopting a thing-centered perspective (Cila et al. 2017) can be used to help reframe these familiar everyday artifacts as agents, for example by developing object personas as a design technique (Cila et al. 2015; Giaccardi et al. 2016).

### 3.6.4 *Authenticity*

We have learned how the intelligence of smart objects can be expressed in an *authentic* way by using a product's inherent structures and materials. For example, students were urged to think about how the flexible sides of shoes, laces and soles, straps, zippers and textiles of backpacks, and grip-controls of power-drills, etc. could be utilized as expressive means to communicate intent and affect. During the course, students were therefore discouraged to use materials that felt alien to their chosen product category.

To design the intelligence of smart objects in authentic ways, designers need to be aware of how movement can express inner life. i.e., motions and behaviors that communicate intentions, emotions and attitudes (Siegman and Feldstein 2014). This allows designers to take advantage of people's ability to perceive inanimate objects as being expressive, comparable to, yet different from, humans and animals (Heider and Simmel 1944; Hoffman et al. 2008; Hoffman and Ju 2014; Marenko and van Allen 2016). New design sensitivities and vocabularies are required to design for such perceptions.



### 3.6.5 Sociability

During the course, we have learned how students approached the interaction with smart objects pragmatically (linked to their function as an everyday product) and as a form of social interaction. For example, seeing Harry the power drill act like a *professional perfectionist* made it apparent that alongside its main drilling function, Harry could *invite, guide* and *stop* the construction worker from drilling based on its judgements about how to drill in the ‘right’ way. Interacting with smart objects then changes from *use-oriented* to *collaboration-oriented* relations, where smart objects can influence, take control, or even overrule the actions of their users, as governed by their intent.

Designing smart objects requires a critical understanding of how they collaborate with us. When objects complement humans on the level of ability, competence or awareness, designers needs to be aware of how people would experience this dependency. For example, people could feel patronized by the object or feel dominated by it. In case of Pat the social backpack, these issues were brought to the foreground quite early in the design process. Pat’s behavior of nudging the teenager forward or sliding of the shoulders, was experienced as dominating. This stirred discussion whether such (social) behavior is acceptable and useful to empower teenagers in this situation. Thus, to evaluate its appropriateness, traditional design criteria would need to be supplemented by social-oriented criteria, such as an object’s ability to communicate and act politely, firmly or cautiously (Niess and Diefenbach 2016).

## 3.7 Conclusion

In this chapter, we have reviewed related work that centers around understanding Artificial Intelligence (AI) in the context of HCI, spanning from the 60s to contemporary views on smart objects in the Internet of Things. We then have introduced an approach that brings traditional values of formgiving practices into interaction design, in particular to the design of smart objects. This approach considers intelligence as being part of, and expressed through, an object’s character, hereby integrating it into a product’s overall physical form. In the development of this approach, we have been informed by material notions on interaction in HCI that deal with the immaterial (data, media, algorithms) and we have added intelligence to the discourse.

We suggest that in order to be successful with this proposed approach, designers need to develop smart objects through iterative cycles, in which their intent, materiality and agency are explored and defined on an increasing level of detail and integration. Furthermore, we have elaborated about how aspects like familiarity, authenticity and sociability should be taken into consideration when designing the object’s character. After having discussed the implications of our approach and findings, within the frame of third wave HCI, we express our hope that this work will

contribute to the development and use of intelligence as a design material, as well as spark new theoretical understandings and approaches in interaction design. We would like to end with a word of gratitude by thanking the students who have participated in our design course and by thanking our colleagues for their support.

## References

- Allmendinger G, Lombreglia R (2005) Four strategies for the age of smart services. *Harv Bus Rev* 83(10):131–145
- Amazon Echo (2017) <https://www.amazon.com/dp/product/B00X4WHP5E>. Last checked 14 Sept 2017
- Apple HomePod (2017) <https://www.apple.com/homepod/>. Last checked 14 Sept 2017
- Bergström J, Brendon C, Frigo A, Mazé R, Vallgård A (2010) Becoming materials: material forms and forms of practice. *Digital Creativity* 21(3):155–172
- Bødker S (2006) When second wave HCI meets third wave challenges. In: Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles (NordiCHI'06). ACM, New York, pp 1–8
- Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. *Interactions* 22(5):24–31. ACM, New York
- Bødker S, Klokosek CN (2012) Dynamics in artifact ecologies. In: Proceedings of the 7th Nordic conference on Human-Computer Interaction: making sense through design. ACM, New York, pp 448–457
- Cila N, Giaccardi E, Tynan O'Mahony F, Speed C, Caldwell M (2015) Thing-centered narratives: a study of object personas. Paper presented at Research Network for Design Anthropology Seminar 3: Collaborative Formation of Issues (January 2015), Aarhus, Denmark
- Cila N, Smit I, Giaccardi E, Kröse B (2017) Products as agents: metaphors for designing the products of the IoT age. In: Proceedings of the 2017 CHI conference on human factors in computing systems. ACM, New York, pp 448–459
- Crilly N, Moultrie J, Clarkson PJ (2009) Shaping things: intended consumer response and the other determinants of product form. *Des Stud* 30(3):224–254. Design, Helsinki
- Djajadiningrat T, Overbeeke K, Wensveen S (2002) But how, Donald, tell us how?: on the creation of meaning in interaction design through feedforward and inherent feedback. In: Proceedings of the 4th conference on designing interactive systems: processes, practices, methods, and techniques. ACM, New York, pp 285–291
- Dourish P (2004) *Where the action is*. MIT Press, Cambridge, MA
- Dourish P, Mazmanian M (2011) Media as material: information representations as material foundations for organizational practice. In: Proceedings of the third international symposium on process organization studies, pp 92–xx
- Dove G, Halskov K, Forlizzi J, Zimmerman J (2017) UX design innovation: challenges for working with machine learning as a design material. In: Proceedings of the 2017 CHI conference on human factors in computing systems. ACM, New York, pp 278–288
- Evans M, Sommerville S (2007) Seeing is believing: the challenge of product semantics in the curriculum. In: Shaping the future, proceedings of the international conference on engineering and product design education, 13–14 Sept 2007, Northumbria University, Newcastle Upon Tyne, UK
- Farooq U, Grudin JT (2016) Human-computer integration. *Interactions* 23(6):26–32
- Ghajargar M (2017) Toward intelligent environments: supporting reflection with smart objects in the home. *Interactions* 24(4):60–62. ACM, New York
- Giaccardi E (2015) Designing the connected everyday. *Interactions* 22(1):26–31

- Giaccardi E, Cila N, Speed C, Caldwell M (2016) Thing ethnography: doing design research with non-humans. In: Proceedings of the 2016 ACM conference on designing interactive systems. ACM, New York, pp 377–387
- Gibson JJ (1979) The ecological approach to visual perception. Houghton Mifflin, Boston
- Google Home (2017) <https://madeby.google.com/home/>. Last checked 14 Sept 2017
- Halse J, Brandt E, Clark B, Binder T (2010) Rehearsing the future. The Danish Design School Press, København
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: Proceeding of CHI 2007. ACM, New York
- Heider F, Simmel M (1944) An experimental study of apparent behavior. *Am J Psychol* 57(2):243–259
- Hoffman G, Ju W (2014) Designing robots with movement in mind. *J Human-Robot Interact* 3(1):89–122
- Hoffman G, Kubat R, Breazeal C (2008) A hybrid control system for puppeteering a live robotic stage actor. In: The 17th IEEE international symposium on Robot and Human Interactive Communication, 2008. RO-MAN 2008. IEEE, pp 354–359
- Holmquist LE (2017) Intelligence on tap: artificial intelligence as a new design material. *Interactions* 24(4):28–33
- WikiPedia Pixar's Luxo 'Junior'. [https://en.wikipedia.org/wiki/Luxo\\_Jr](https://en.wikipedia.org/wiki/Luxo_Jr). Last checked 14 Sept 2017
- Jacucci G, Spagnolli A, Freeman J, Gamberini L (2014) Symbiotic interaction: a critical definition and comparison to other human-computer paradigms. In: International workshop on symbiotic interaction. Springer International Publishing, Cham, pp 3–20
- Janlert LE, Stolterman E (1997) The character of things. *Des Stud* 18(3):297–314
- Janlert L-E, Stolterman E (2017) The meaning of interactivity—some proposals for definitions and measures. *Human Comput Interact* 32(3):103–138
- Kaptelinin V, Nardi BA (2006) Acting with technology: activity theory and interaction design. MIT Press, Cambridge, MA
- Kortuem G, Kowsar F, Fitton D, Vasughi S (2010) Smart artifacts as building blocks for the internet of things. *J Internet Comput IEEE* 14(1):44–51
- Krippendorff K (1989) Product semantics: a triangulation and four design theories. In: Väkevä S (ed) Product semantic '89. University of Industrial Arts, Helsinki
- Krippendorff K (2006) The semantic turn: a new foundation for design. Taylor & Francis CRC, New York
- Kuutti K, Bannon LJ (2014) The turn to practice in HCI: towards a research agenda. In: Proceedings of the 32nd annual ACM conference on Human factors in computing systems. ACM, New York, pp 3543–3552
- Laurel B (1997) Interface agents: metaphors with character. In: Friedman B (ed) Human values and the design of computer technology. Cambridge University Press, Cambridge, MA, pp 207–219
- Licklider JC (1960) Man-computer symbiosis. *IRE trans Hum Factors Electron* 1:4–11
- Lim YK, Stolterman E, Jung H, Donaldson J (2007) Interaction gestalt and the design of aesthetic interactions. In: Proceedings of the 2007 conference on Designing pleasurable products and interfaces. ACM, New York, pp 239–254
- Marenko B, van Allen P (2016) Animistic design: how to reimagine digital interaction between the human and the nonhuman. *Digital Creativity* 27(1):52–70
- McKim RH (1980) Thinking visually: a strategy manual for problem solving. Wadsworth, Belmont
- Muller W (2001) Order and meaning in design. Lemma, Utrecht
- Niess J, Diefenbach S (2016) Communication styles of interactive tools for self-improvement. *Psychol Well Being* 6(1):3
- Norman DA (1994) How might people interact with agents. *Commun ACM* 37(7):68–71
- Papanek V (1971) Design for the Real World: human ecology and social change. Pantheon Books, New York
- Pasman G, Rozendaal M (2016) Exploring interaction styles through video. In: Proceedings of the 18th international conference on engineering and product design education (E&PDE16), Aalborg, Denmark, 8th–9th September 2016

- Petersen MG, Iversen OS, Krogh PG, Ludvigsen M (2004) Aesthetic Interaction: a pragmatist's aesthetics of interactive systems. In: Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques. ACM, New York, pp 269–276
- Pevsner N (1991) *Pioneers of modern design: from William Morris to Walter Gropius*. Penguin Books, London
- Raizman D (2004) *A history of modern design: graphics and products since the industrial revolution*. Laurence King Publishing, London
- Ramdueny-Ellis D, Dix A, Evans M, Hare J, Gill S (2010) Physicality in design: an exploration. *Des J* 13(1):48–76
- Reeves B, Nass C (1996) *How people treat computers, television, and new media like real people and places*. CSLI Publications and Cambridge university press, Stanford/Cambridge, MA, pp 3–18
- Robles E, Wiberg M (2010) Texturing the material turn in interaction design. In: Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction. ACM, New York, pp 137–144
- Rosner DK (2012) The material practices of collaboration. In: Proceedings of the ACM 2012 conference on computer supported cooperative work. ACM, New York, pp 1155–1164
- Ross PR, Wensveen SA (2010) Designing behavior in interaction: using aesthetic experience as a mechanism for design. *Int J Des* 4(2):3–13
- Rozendaal MC (2016) Objects with intent: a new paradigm for interaction design. *Interactions* 23(3):62–65
- Sabou M (2010) Smart artifacts: challenges for semantic web research, semantic web – interoperability. *Usability Applicability* 1:1–5
- Schön DA (1984) *The reflective practitioner: how professionals think in action*, vol 5126. Basic Books, New York
- Siegmán AW, Feldstein S (eds) (2014) *Nonverbal behavior and communication*. In: Psychology press, Oxford
- Smets G, Overbeeke K, Gaver W (1994, April) Form-giving: expressing the nonobvious. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, New York, pp 79–84
- Taylor AS (2009) Machine intelligence. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, New York, pp 2109–2118
- Tholander J, Normark M, Rossitto C (2012) Understanding agency in interaction design materials. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, New York, pp 2499–2508
- Vallgård A (2014) Giving form to computational things: developing a practice of interaction design. *Pers Ubiquit Comput* 18(3):577–592
- Vallgård A, Redström J (2007) Computational composites. In: Proceedings of the SIGCHI conference on Human factors in computing systems. ACM, New York, pp 513–522
- Verplank B (2009) *Interaction design sketchbook: frameworks for designing interactive products and systems*. <https://crma.stanford.edu/courses/250a-fall-2004/IDSketchbok.pdf>. Last access 14 Sept 2017
- Vihma S (1995) *Products as representations: a semiotic and aesthetic study of design products*. University of Arts & Design, Helsinki
- Wiberg M (2014) Methodology for materiality: interaction design research through a material lens. *Pers Ubiquit Comput* 18(3):625–636
- Wiberg M (2018) *The materiality of interaction: notes on the materials of interaction design*. MIT Press, Cambridge, MA
- Wiberg M, Robles E (2010) Computational compositions: aesthetics, materials, and interaction design. *Int J Des* 4(2):65–76
- Wiberg M, Ishii H, Dourish P, Vallgård A, Kerridge T, Sundström P, Rosner D, Rolston M (2013) Materiality matters – experience materials. *Interactions* 20(2):54–57
- Wooldridge M, Jennings NR (1995) *Intelligent agents: theory and practice*. *Knowl Eng Rev* 10(2):115–152

# Chapter 4

## De-instrumentalizing HCI: Social Psychology, Rapport Formation, and Interactions with Artificial Social Agents



Ritwik Banerji

**Abstract** Decisions in designing artificial social interactants to reproduce culturally-specific forms of human sociality evince a range of conceptions of the norms and cognitive processes involved in the human social interactions themselves. Regarding the use of machine learning (ML) in such systems, decisions whether or not to use this approach implicitly presents questions on the nature of the interpersonal adaptation that takes place and indicate a range of conceptions of the values which structure these interactions. In the design of virtual performers of musical free improvisation, several designers assume that the experience of equal partnership between improvisers can only be afforded through deployment of ML in such systems. By contrast, tests of agents not based in ML reveal that human beings experience illusions of “adaptation” in interactions with systems which lack any adaptive capacity. Such results suggest that HCI research with artificial social interactants may be used to raise new questions about the nature of human interaction and interpersonal adaptation in the formation of relationships over time.

### 4.1 Technological Re-embodiments of Human Practice

At its core, artificial intelligence (AI) research aims to create technologies which either match or surpass the natural or acquired cognitive capacities for creativity and productivity which human beings readily exhibit in a variety of activities. Inevitably, developing machines to engage in human practices involves creating computational representations of how fluent and skillful practitioners sense, feel, think, and act in response to the world and other human beings. For most AI researchers, designing

---

R. Banerji (✉)  
Department of Music, University of California, Berkeley, CA, USA  
e-mail: [ritwikb@berkeley.edu](mailto:ritwikb@berkeley.edu)

a system that does accounting, drives a car, or engages in a customer service interaction necessitates identifying objective standards by which those practices are normally executed and developing a computational description of them.

Nevertheless, as numerous critics of science and technology research have suggested (Friedman and Nissenbaum 1996; Nissenbaum 2001; Edwards 1997; Lewis 2000; Helmreich 2001; Forsythe 2002; Suchman 2006; Seaver 2012; Wilf 2013a, b; Burrell 2016; Seyfert and Roberge 2016), AI systems often reproduce at least a portion of the designer's subjective interpretation of both the tasks themselves and the culturally-specific values and principles that shape them. If one designs a machine to do what a human does, then one is inescapably also invested in a project of producing a representation of the cultures that provide meaningful contexts for human action. Moreover, like any other representation of culture (e.g., an ethnographic text, film, or novel), there exists at least a small gulf between how an AI system characterizes human action and the actual actions of human practitioners per se.

As a subfield of AI, much the same is true for machine learning (ML). Like AI, ML research includes projects dealing with reproductions of actual human behavior as well as those that aim to transcend it. Similarly, in the effort to simulate or outperform human capacities for learning, adaptation, or even improvisation, ML is a technical project that often necessitates an investigation of how human beings process their experiences of interacting with the world and successively shift their behavior through further experience. In the context of human computer interaction (HCI) with an artificial social agent, ML functions not only as a component in simulating adaptive human sociality, but as a means of posing questions and hypotheses about how human beings form rapport and adapt to one another over time. Even if HCI researchers do not frame their work as an investigation of social psychology, their work may nevertheless have significant implications for the study of human consciousness in social interaction.

With these perspectives in mind, this chapter analyzes the relationship between ML research in the context of HCI along with the culture, phenomenologies, and processes of social behavior implicitly investigated and characterized in the creation of such technologies. In light of recent efforts to move the field of HCI beyond its strict focus on functionality and ease of use (Bødker 2006; Harrison et al. 2007), this chapter focuses on three questions regarding the relationship between ML and human culture:

1. How does a designer's attempt to encode the human values that shape particular activities influence (1) their choice of whether to use ML, and (2) the decision to use certain ML techniques rather than others?
2. To what degree is ML effective in performing the kind of interpersonal adaptation processes a designer seeks to simulate? In which culturally-specific domains of human action is ML necessary? How do these differences reflect cultural variations and individual discrepancies in how subjects understand the norms of a given cultural milieu?
3. How can human experiences with ML technologies be used to better understand processes of interpersonal adaptation and rapport formation? Similarly, how can

human experiences with such systems be used to investigate the ways cultural values affect the form and degree of interpersonal adaptation expected over subsequent interactions?

To address these questions, this chapter surveys numerous stances on the necessity of ML techniques for designers of artificial social agents that engage in a form of social interaction unique to a particular musical practice. Specifically, this chapter looks at the design and design rationale of a range of virtual performers of free improvisation, a form of experimental music in which the performance is not the realization of a musical score, but results from the impromptu social and musical interactions between players. Far beyond simply reproducing the sonic surface of this artistic practice, designers of these systems seek to ensure that these artificial social interactants embody the values which shape how artists engaged in this musical form coexist with one another. In the process of pursuing this goal, these designers offer a range of perspectives on the efficacy of ML methods as a means of simulating the kinds of interactions these performers engage in and adhering to the principles that define this form of music as a culturally-specific mode of face-to-face interaction.

For a variety of reasons, several designers working in this area assume that ML constitutes an effective means of encoding the socio-political ideals at the center of this musical practice. However, tests of such systems with actively performing free improvisers suggest that ML constitutes, at best, just one of many possible ways of programming an artificial social interactant to function as an ideal human interlocutor in this domain. Surprisingly, tests of systems not designed using ML techniques reveal that performers, and perhaps human beings more generally, are often prone to experiencing these non-adaptive systems as “adaptive” even though the system has no such capacity encoded.

This illusion has two likely explanations. On the one hand, the sensation of adaptation may result from the well-documented psychological phenomenon known as the effect of “mere exposure” (Zajonc 1968), or the principle that increased contact with a given stimulus corresponds with more positive evaluations of it over time. On the other, the impression of adaptation may arise from a dynamic feedback loop resulting as the human participant’s adjustment to the non-adaptive system elicits a different side of the system’s behavior (Hsu and Sosnick 2009). As a result of their own contributions to the interaction and the system’s reactions to this novel input, the human player may be led to believe that the system has “adapted” despite the absence of any real computational means for the system to do so (i.e., ML).

Overall, I argue that even where its efficacy is doubtful or contested, the use of ML in the design of artificial social interactants offers a means of examining the social psychology of rapport formation and how human beings develop intimate knowledge of their interlocutors over a series of interactions. As will be discussed in greater detail, recent empirical work (Hsu and Sosnick 2009; Banerji 2012; Linson et al. 2015) suggests that ML may not be necessary for achieving the kinds of interactions that free improvisers prefer to engage in. In turn, this data raises questions about the nature of the interpersonal adaptation between players which

theorists of free improvisation claim is central to musical interaction in this practice (Waterman 2008; Young 2010; Beins 2011).

More importantly, such experiences also suggest a degree of ambiguity in the notion of “rapport formation” and may have implications for the theorization of this aspect of human interaction in the field of social psychology (Cappella 1990; Tickle-Degnen and Rosenthal 1990; F. Bernieri and Gillis 1995; F. J. Bernieri et al. 1996; Grahe and Bernieri 1999; Lakin and Chartrand 2003). At face value, the idea of rapport formation suggests that a dynamic process of interpersonal adaptation and learning takes place between human beings. By contrast, the fact that improvisers report that they experience “adaptation” in interactions with systems which lack any meaningful capacity for adaptation over time suggests that “rapport formation” may actually be at least partially the result of relatively static processes.

Therefore, beyond “implications for design” (Dourish 2006), the shortcomings of ML in this or other domains of simulated human interaction may offer a source of data about the nature of rapport formation and interpersonal adaptation which may not be available through other methodologies. Additionally, the limitations of ML offer a counterpoint to debates about how an algorithm would depict culturally-specific principles and cognitive processes of interpersonal adaptation. Though numerous designers assume that ML is necessary for living up to specific human values, empirical studies show that this issue is far more complicated and raise several questions about the concept of “rapport.”

Finally, though many perspectives in HCI argue that a focus on the study of user culture poses a risk of derailing HCI research from the goal of refining the fit between human tendency and system design (Crabtree et al. 2009), this interpretation underestimates what the study of culture and sociality still offers to design. Whether one is interested in researching culture or social behavior for their own sake or interested in refining a system to the needs of a human population, the study of how people conduct themselves and pursue various goals remains essential. If HCI is committed to the creation of technologies that intuitively complement the ways and means of real human users, then any data which uncovers aspects of the sociality or culture of a technology’s human interactants is likely to prove useful for this original, pragmatic goal of HCI research.

## 4.2 Delimitations and Scope

While third-wave HCI research raises numerous new questions and retheorizes several old issues in HCI such as the fit between user and system or the proper role of technology in everyday life (Bødker 2006; Harrison et al. 2007), this chapter deals with only a selection of the issues foregrounded by the third-wave perspective. Though the third-wave has foregrounded issues such as embodiment (Bardzell and Bardzell 2011), emotion (Boehner et al. 2007), and the proliferation of interactive technologies beyond the workplace, this chapter focuses on the inscription and encoding of cultural values in interactive systems (Sengers et al. 2004; Fuchsberger



et al. 2012) as well as a phenomenological perspective on how users experience technologies (Harrison et al. 2007).

Regarding phenomenology, the present discussion focuses on an account of human experience of technology without respect to whether what one is experiencing is “real” or not. In order to do so, this chapter takes after Edmund Husserl’s articulation of the concept of “bracketing” or “epoché” in his classic work *Ideas* (Husserl 1913/2012: 56–60). For Husserl, bracketing describes a basic conceptual commitment of phenomenology in which one attunes to human experience without necessarily being concerned with whether those experiences are grounded in an empirically verifiable reality. For example, this would be an account of one’s experience of hearing a “dog barking” irrespective of the fact it may either be an actual canine vocalization or another auditory stimulus which nevertheless evokes the illusion of a bark.

For this discussion (and perhaps for third-wave HCI more generally), the principal relevance of bracketing is that it is of utmost importance to examine how reality (or technology) is experienced by an individual without necessarily being burdened by the issue of whether that experience is evidence of a scientific fact or an utter hallucination. Regardless of their “reality,” the perceptions arising from encounters with humanlike technologies are often felt to be “real” even when one is aware of the illusion of these sensations. This is not to assert that such hallucinations possess veracity, nor is it to assert that it does not matter whether a system is designed using ML or not. Rather, it is to recognize that human beings live by many illusions and that it is quite likely that such illusions play a role in human encounters with humanlike technologies like an artificial social interactant. This is especially significant given that ML-driven HCI applications aim to effectuate an illusion that the human interactant is “in fact” engaging with another member of their species.

While this chapter addresses work at the intersection of HCI and ML, the principle focus will be agents that are built with the explicit goal of simulating forms of social interaction which human beings regularly engage in. This area of HCI and ML work includes research into the development of interactive conversational agents which produce a sense of realness and humanness in how they interact with a human being in terms of linguistic competence, sound production, and perception (Eklund 2002; Maatman et al. 2005; Morrissey and Kirakowski 2013). It also includes forms of “artificial life” (Langton 1997) in which machines perform as if they were human beings (e.g., artificial general intelligence, social robotics, video-game characters, etc.). With this focus in mind, other areas of research taking place at the intersection of ML and HCI are less relevant. This includes the domain of “interactive machine learning” (Fails and Olsen Jr 2003; Fiebrink et al. 2011; Amershi et al. 2014), in which users are enabled to correct the process by which ML parses and analyzes information.

Although interactive ML (or IML) research overlaps with the present discussion’s concern with the simulation of face-to-face human interaction, current work in IML does not seek to reproduce human interactions through speech, gesture, gaze, and other embodied communication in real time. Instead, current IML research largely focuses on reproducing an interaction between a supervisor and a subordi-

nate. Just as a higher-level worker might train an assistant to analyze or parse certain types of incoming information, IML partially resembles this professional mode of human interaction. Even so, the modality of this form of HCI is relatively artificial when compared with conversation or the kind of social interaction through music discussed here. Eye contact, gesture, or intonation (to say nothing of speech itself) have yet to be used as a mode of interaction for IML research, where visualization of the ML protocol has been the primary sensory output relayed to the user for their critique.

Still, an interaction between a human user and an interactive ML system, insofar as it simulates a supervisor-worker dyad, remains relevant because it resembles a real form of social interaction in business or other organizational contexts. More importantly, like any type of human sociality, values such as autonomy, responsibility, transparency, or personal integrity form much of the psychological infrastructure of an interaction between two workers. Similarly, it is quite likely that IML applications, particularly where they simulate work-related human interactions, are inflected by the same values and expectations that shape interpersonal relations in organizations even though the “worker” in that case is a nonhuman user technology. Or, considering the more experientially-oriented angle of third-wave HCI, it is possible that an interaction with an IML system may remind users of similar interactions they have had with co-workers as they mentor them in how to deal with certain kinds of information.

### 4.3 Encoding Egalitarianism

Since George Lewis’ *Voyager* (1993, 1999, 2000), researchers in computer music have developed a variety of interactive virtual musicians built for performers of musical free improvisation to play with as if these systems were just another human improviser (Blackwell and Bentley 2002; Assayag and Dubnov 2004; Hsu 2005; Casal and Morelli 2007; Yee-King 2007; Young 2008; Bown 2011; Carey 2012; Linson et al. 2015). Based on each designer’s experience with and conception of this musical practice, these systems are constructed with the goal of reproducing the same kind of inspiring, challenging, and spontaneous social interactions through sound that human performers of free improvisation hope to engage in with one another. In order to create the experience of playing with another free improviser, these virtual performers are designed to process live acoustic input from the human player and respond in real time. Thus, in the ideal, these artificial improvisers recreate the sensory and interactive experience of making music with a real improvising partner.

Before turning to the issue of ML, there are a number of ways that these systems are built to encode and mechanically re-embodiment the social values that shape how improvisers engage with one another. Specifically, as is discussed across a wide range of literature on this practice (Spellman 1966; Kofsky 1970; Bailey 1980/1993; Stanyek 1999; Steinbeck 2010; Lange 2011; Carles and Comolli 1971/2015; Corbett 2016; Rodriguez 2016), free improvisation purports to liberate musicians from the

typical interpersonal hierarchies which define many forms of music-making (e.g., composer and performer, conductor or bandleader and ensemble, soloist and accompanist, critic and performer, etc.). Instead, performers of free improvisation assemble themselves in a flat, nonhierarchical, egalitarian arrangement in which no player serves as leader and no composition is used as a guide for how sonic events will take place in the performance. Additionally, free improvisers seek to avoid relying upon any traditional structures for musical performance (i.e., pulse, harmony, form, genre, style, convention, etc.) in order to prevent the implicit hierarchies that often result when some individuals are more proficient in certain culturally-specific musical structures (e.g., jazz, Indian classical music, baroque music, etc.). By eliminating such structures, improvisers aim to create a situation in which the only determinant of what takes place in performance is the dynamic interaction of the personalities of the improvisers themselves (Blackwell and Young 2004). Nevertheless, for all that they may strive to avoid such constraints, recent critiques have duly noted the limitations of these emancipatory pursuits (Backstrom 2013; Canonne and Garnier 2015).

Although individual programming approaches of these designers vary greatly, each of these researchers agrees about some basic principles for how the lofty goal of egalitarianism should be translated into the creation of such a system. Across the board, designers concur that systems cannot reproduce an egalitarian interaction if they are built for the performer to use as an instrument. Instead they must be designed in order to allow for musicians to engage with these systems as if they were a fellow player (Rowe 1992; see also Lewis 2000). In practical terms, this means that the mode of interaction with the system should not allow the human performer to “veto” (Lewis 1999: 104) or otherwise directly control the system’s sonic output in real time as one might with an instrument. Unlike numerous other interactive performance systems (see Chadabe 1997), these systems do not incorporate any kind of haptic or tactile interface which would enable the human subject to retain their position of mastery and control<sup>1</sup> over the mechanical musical object.

In turn, this is analogous to the way that improvisers themselves interact musically and socially. Before, during, and after they play, free improvisers habitually refrain from directing, criticizing, or instructing other musicians, irrespective of how irritated they may have been with their playing or whether they had specific desires for how the performance should have taken place (Banerji 2016). Even when disgust may seem obvious from a player’s choice to stop playing (see Fischlin et al. 2013: 203–219) or play more loudly in order to drown others out, it is often unclear what these kinds of interactive actions mean or if they represent a critical judgment of one player against others. Moreover, direct expressions of criticism after performance between players is implicitly regarded as a threat to the nonhierarchical ideal (Borgo 2002; Pras et al. 2017). This reticence<sup>2</sup> is both a means of

<sup>1</sup>Again, in the case of a virtual free improviser, the designer does retain control over the system’s behavior. But unlike the control exerted over a system built to function as an “instrument” (Rowe 1992), a “player” system cannot be directly controlled on a moment to moment basis.

<sup>2</sup>While peer criticism is not seen in this way by participants of other egalitarian social projects (Chaudron 1984; Snyder and Fessler 2014), improvisers view criticism as a kind of speech act that instantly nullifies equality by placing the speaker in a position of authority with regard to the actions of the addressee.

respecting the equal status of fellow players as much as it is a performance of an ethos of openness to a “diversity” of practices within free improvisation (Bailey 1980/1993: 83).

Similarly, improvisers do not interact through any sort of visual cues between performers (Lewis 2007), as has been occasionally suggested (Andersen and Brooks 2003). Though moments of gaze between performers do occur, it is not entirely clear what meaning they carry nor what impact they have on how the interaction progresses. In sum, researchers in computer music concur that if these systems are to interact with human performers of free improvisation as their equals, then (1) there should be no mechanism that allows direct control of their behavior, and (2) it is not necessary for them to be able to respond to gestural cues.

### **4.3.1 *Is ML Necessary for the Experience of Equal Partnership?***

Beyond these points of agreement, however, designers differ significantly on the question of whether ML is essential for allowing systems to produce interactions that yield an experience of equal partnership and agency in the collective outcome of the performance. At the computational and interactive levels, these designers take divergent approaches to encoding egalitarianism in the behavioral capacities of these systems. By no means do they all assume that ML or any other adaptive systems technique (i.e., genetic co-evolution) is the best means of encoding this social ideal. For example, about free improvisation generally, George Lewis writes that “the possibility of internalizing alternative value systems is implicit from the start” (Lewis 1999: 102). For most proponents of ML methods, Lewis’ description of free improvisation would immediately suggest that it is imperative that a designer incorporate ML if they aim to encode the values that these musicians strive to live by. After all, what else would be implied by the term “internalizing”? Curiously, however, Lewis’ system *Voyager* (2000) does not incorporate ML or any other adaptive systems technique.

#### **4.3.1.1 Yes**

While Lewis and other designers do not use ML techniques in the design of their systems (Hsu 2010; Linson et al. 2015), several other researchers in this domain assume that a capacity for iterative adaptation to the human player in real time is a necessary part of the successful design of such a system. Implicitly, this line of thinking limns several crucial assertions these researchers make about the social psychology of free improvisation and how it is shaped by values like egalitarianism and multiculturalism. Specifically, the use of ML or other adaptive techniques assumes that in order to properly engage in an egalitarian style of social interaction,

two players participating in a collective free improvisation must both make an effort to adapt to one another as they make music together. Broadly, then, the use of ML in such systems assumes (1) that the human player wants to feel that the other improviser is learning about their habits and tendencies in interaction and (2) that this ongoing interpersonal learning process is essential for yielding the experience of authentically egalitarian social interaction.

However, the use of ML implicitly poses a hypothesis about what happens in the psyche of the human improviser and how they experience musical interaction with another. Because improvisers generally seek to preserve an equity of status between performers, designers of ML-based virtual free improvisers assume that the human improviser is forming a memory of how the other plays and tends to react in order to draw on these memories later on. Overall, the application of ML in this domain also assumes that if one player (whether human or machine) fails to engage in this process of adaptation while the other makes efforts in this direction, the player that maintains their original manner of interaction has acted autocratically or egotistically, thereby disrupting the equity of status between participants.

For designers using ML or other adaptive techniques like genetic co-evolution (Eiben and Smith 2003; Casal and Morelli 2007) or particle swarm optimization (Kennedy and Eberhart 1995; Blackwell and Bentley 2002) in the creation of these artificial performers, the principal objective behind deploying these methods is to allow the system to have a better knowledge of the habits and tendencies of its human interlocutor. While adaptive techniques can be used to allow the system to simulate the sympathetic and cooperative interactive style of a human player, they can also enable the system to exhibit defiant or oppositional playing. Though many players prefer more supportive styles of interaction, other players find that this interactive attitude lacks the kind of drama and tension they prefer to produce in performance (Banerji 2016). For example, Oliver Bown's *Zamyatin* system (Bown 2011, 2015) uses information learned about the human performer to create both oppositional and sympathetic responses to their playing. Likewise, Michael Young's *NN Music* (Young 2008) acquires information about the player's current performance "state" (i.e., the average and standard deviation of various timbral characteristics) in order to develop a catalog of their playing tendencies. Nevertheless, this index is not necessarily used to create sympathetic behaviors and can also be used to simulate the intentional, audible divergence many players value in musical interaction (Banerji 2016).

Although it uses neither ML nor any other adaptive systems technique, Ben Carey's *\_derivations* system (Carey 2012) builds a "memory" of its interactions with a specific performer by making a running catalog of phrases produced by the human player for later use as the basis of the system's improvisatory output. If an incoming phrase from the human player is similar<sup>3</sup> to a phrase currently stored in the system's database, the system calls up this phrase and uses it as base material for

---

<sup>3</sup>Similarity is judged by an analytical comparison of incoming and stores phrases on the basis of their loudness, pitch content, spectral centroid (or "brightness"), noisiness (or ratio of tone to noise) as well as the mean and range for these values in a given phrase.

improvisatory exploration (e.g., by manipulating the pitch or rhythmic content of the phrase). This evokes a sensation that the system has the kind of knowledge of the human player that might develop from repeatedly working with the same duo partner over time. At the same time, it avoids reproducing a type of personality which is too prone to cooperation and fails to generate the sense of interpersonal contrast that many improvisers value.

In the same vein, Gerard Assayag's *OMAX* project (Assayag and Dubnov 2004; Cont et al. 2006; Assayag et al. 2010) at IRCAM (Institut de Recherche et Coordination Acoustique/Musique), exclusively frames the use of ML around the goal of simulating a sense of "style learning" (see Pachet 2003; Assayag et al. 2006) between human and machine. That is, the system learns the player's "style"<sup>4</sup> in terms of their tendencies in the use of pitch and rhythm over time. Again, this assumes that because the human being tries to learn the style of the computer, the computer must also make an effort to learn the style of the human. However, while "style learning" may imply that the system will tend to exhibit a spirit of collaboration, *OMAX* is designed to exploit what it learns to react both sympathetically and antagonistically.

In the end, the purpose of deploying ML in this domain is not necessarily to bring the system to adapt to every move the player makes, but rather to build its capacity for acquiring knowledge about the player in order to try and predict their actions (Assayag et al. 2010). This predictive capacity is used not only to create responses which closely resemble the human performer's actions but responses which deviate and contrast with their suggestions as well. Similar in approach to Assayag's *OMAX* project, Nicholas Collins' *Improvagent* system (Collins 2008) tries to both predict the responses of the human player and also to understand the consequences of its own actions in terms of the player's response. The system assesses its own ability to predict the player's actions by comparing its predictions with the actual outcomes of the player's behavior. If there is a discrepancy between the player's actions and the system's prediction, the method of prediction is adjusted accordingly.

However, for all that there may be an intuitive association between ML and the kind of adaptation implied by certain notions of egalitarianism as a logical choice for realizing these ideals in musical interaction, it is unclear as to whether ML succeeds convincingly. This is mainly due to a consistent trend in which designers of these systems do not test them with active performers of free improvisation. As a result, the efficacy of ML as a means of embodying the sensation of nonhierarchical interaction is not known. Still, for David Plans Casal (2008), the use of an adaptive system proved to be one that he himself found frustrating as a performer collaborating with the system. In his experience, the system's constant tendency to try and improve its responses to his playing resulted in a vexing situation in which the system failed to continue in one direction for long enough in order to create a sustained musical idea.

---

<sup>4</sup>As Eitan Wilf notes (2013b), this is a very specific notion of the term "style" which essentially predetermines what can and cannot even count as style.

#### 4.3.1.2 Maybe Not

While several designers have chosen to use ML and related techniques in order to build systems which uphold values central to free improvisation as a socio-cultural practice, Lewis and others have created systems which do not assume that ML is a necessity in working towards the enactment of the principles of egalitarianism or multiculturalism. In essence, this means that across this range of designers, two fundamentally distinct notions of egalitarianism are at work. The first, which assumes ML is essential, interprets egalitarianism as an ethico-political ideal which requires that improvisers engage in a process of mutual adaptation to one another's playing and improvisatory tendencies. According to this view, the production of a flat, nonhierarchical social structure, or "leveling" (Woodburn 1982; Boehm 1993), is best achieved through the efforts of each member to assimilate to the group.

Contrasting with this assimilationist or integrationist interpretation, however, those who design interactive virtual free improvisers which are not based in ML or related adaptive techniques make a design decision that implicitly suggests their belief that such interpretations of egalitarianism are questionable. Systems like those of George Lewis (2000), Adam Linson et al. (2015), or William Hsu (2010) which do not use adaptive techniques suggest a second, alternate interpretation of egalitarianism that is rooted far more in a sense of autonomy than it is in cooperation or coalescence. For such designers (as well as the actual improvisers whose views of the situation they may reflect), adaptation to the other is incongruent with the nonhierarchical ideal because it introduces a sense of hierarchy as one player makes an effort to yield to the other. In this perspective, it is far better to maintain a sense of independence in terms of one's own playing protocol in relating to the group than it is to acquiesce to or accommodate the stylistic tendencies of others.

The meaning of these varying attitudes towards the use of ML in the creation of synthetic musicianship extends far beyond the use of these techniques to simulate human behavior or create a convincing artistic result. Rather, these differing stances on the need to integrate ML bespeak the diversity of interpretations of what "egalitarianism" is as a moral or political value in a context far removed from the immediate context of HCI. Design decisions in the creation of these virtual musicians indicate the presence of two distinct concepts of what it means to occupy a status equal to one's interlocutor and what might be required to preserve the flatness of the distribution of power between two individuals in a real time social encounter. For those that adopt ML methods in pursuit of creating a free improviser from computing machinery that performs in an egalitarian manner, their choices suggest an assumption that the experience of equality is only possible through the gradual adaptation of one social interactant to the other and vice versa. Failure to do so is regarded as a setback to realizing the nonhierarchical ideal. Conversely, for those who do not use ML in these systems, such design choices suggest an interpretation of egalitarianism in which adaptation is regarded as capitulation. In this view, disuse of ML represents and performs a rendering of egalitarianism built upon the sense that all are equal if no one feels compelled or expected to lose their autonomous sense of self and conform to the rest of the group.

Crucially, however, none of these claims are made explicit in the written documentation of these systems in published literature. Rather, it is only implicitly that such design decisions refer to egalitarianism at all from the fact that these systems aim to emulate how improvisers socially interact through sound in performance and that the discourse around free improvisation pervasively emphasizes the notion that the practice is predicated upon egalitarian principles. Returning to the third-wave HCI concern for analyzing how systems encode values, it is imperative that designers articulate how they interpret particular cultural values when they attempt to encode them in interactive systems. Doing so would not only reduce any ambiguities in this issue but also allow their systems to be understood as they should be: hypotheses on the nature of these values themselves expressed in the language of computation and the experience of embodied interaction with virtual social agents.

#### **4.4 The ML You Never Asked for: Illusions of Adaptation in the Absence of ML**

But regardless of the diversity of interpretations of egalitarianism implied in the range of attitudes about the utility of ML in this domain of arts technology, it is debatable as to whether ML is necessary for achieving the sense of adaptation and rapport that an improviser may desire. This is not a programmatic statement suggesting that the use of ML should not be investigated in this research domain nor is it to suggest that “the implication is not to design” (Baumer and Silberman 2011). Whether or not ML is effective in satisfying an improviser’s expectation or desire for a sense of adaptation over time, creating — and more importantly, rigorous testing of such systems with human interactants — is likely to provide important data which enables a more detailed understanding of how improvisers form rapport. More generally, research in this area also contributes to the ongoing social psychological investigation of what “rapport” really consists of in human relationships and how it develops over time (Cappella 1990; Tickle-Degnen and Rosenthal 1990; F. Bernieri and Gillis 1995; Grahe and Bernieri 1999; Cassell et al. 2007; Gratch et al. 2007; Huang et al. 2011).

In the context of virtual performers of free improvisation, recent empirical work suggests that ML may not be necessary for human beings to experience a particular system as adaptive. Such results are rather paradoxical since it is reasonable to assume that a system using ML would be more likely to produce the sensation of adaptation to the human interactant than a system using no such technique. Yet in several cases, researchers have noted that human interactants report that they have nevertheless experienced a sense (or really, an illusion) of adaptation and evolution in the behavior of systems which are in no way based in ML or any other kind of adaptive systems technique. For example, after designing his *Odessa* system, Adam Linson tested this virtual free improviser with a group of eight performers (Linson et al. 2015). In each test, a single human improviser and the *Odessa* system performed a series of three “duets.” Curiously, Linson notes that despite the fact that



his system uses neither ML nor any other adaptive systems technique, improvisers felt that the system's behavior seemed to have adapted to theirs over the course of the three takes.

Similarly, as part of my own research in this domain of computer music, I have tested a virtual free improviser of my own design, *Maxine* (Banerji 2010, 2012, 2016), with over 90 improvisers in Berlin, San Francisco, and Chicago. Like *Odessa*, *Maxine* is not based in ML nor does it use any other type of adaptive systems technique. Within this larger project to test this system with improvisers, I conducted a small experiment in which I asked eight improvisers in the San Francisco Bay Area to play a series of 10 short takes with *Maxine* in a studio setting at the Center for New Music and Audio Technologies (CNMAT) in the fall of 2010. After each take, the improviser was asked to complete a simple numerical evaluation of the system according to four criteria<sup>5</sup> and also provide an open-ended written commentary on their experience of that piece. Though the purpose<sup>6</sup> of these tests was not directly related to the present discussion, the qualitative data collected in this experiment resonate with Linson's observations of a perceived sense of adaptation despite the absence of any computational architecture which would directly contribute to such an experience. Specifically, several participants of this experiment noted that the system was "better" over the course of the experiment, using this term several times to describe the improvement of one take over the next.<sup>7</sup> For example, for the first take, all but one player had negative comments about their duo with the system. For the next take, however, five of eight participants reported a strongly positive difference.

Outside the context of this experiment, I have also noticed that players have frequently commented on the "improvement" of the system after their first experiences of it. I find such comments perplexing given the fact that the changes I have made to the system's original design in 2009 are so minimal as to be inconsequential<sup>8</sup> in a

---

<sup>5</sup>These criteria were (1) the degree to which the system inspired you to respond to its playing, (2) satisfaction with unexpected or surprising responses from the system, (3) the overall sense that the interaction was meaningful, and (4) whether the system's responses seemed relevant or random.

<sup>6</sup>The main question for this experiment focused on the issue of whether or not the active listening of another improviser increases or decreases an improviser's level of aesthetic or social-interactive satisfaction of the experience of playing music. In order to investigate this question, in a random selection of the 10 takes, the system was set to listen to a prerecorded track (and therefore, not listen to the sonic events of the current take) whereas in the remainder of takes the system listened to the combination of itself and the human performer, this being the way the system was originally designed to receive input in a performance setting. (For further discussion, see Banerji 2012.)

<sup>7</sup>To be clear, the quantitative data from the experiment does not necessarily suggest a clear sense of evolution in the player's experience across the set of 10 takes. However, the quantitatively-graded criteria do not directly correspond with positive or negative sentiments about the system's interactivity as an experience. With the exception of one criterion ("meaningfulness"), the criteria evaluated refer to the player's observations about the interaction overall and do not inherently convey judgments about the aesthetic value of the experience.

<sup>8</sup>These were mainly minor tweaks in order to enable the system to start and stop at the push of a single button. Such changes had no effect on how the system would begin to play, behave during the improvisation itself, or how it would "end" pieces.

player's overall experience. This was particularly striking at a concert to celebrate the release of a duo recording of myself and Maxine (Banerji et al. 2014) which took place at CNMAT in 2014. At that event, I invited several improvisers to play with Maxine in a series of duos. Despite the fact that I had done almost nothing to improve the system since their last experiences with it, many of them congratulated me to tell me that the system was finally sounding "better." One player reported that the system had "come a long way" and "really grown up" since his initial encounter with it in 2010. Beyond this incident, it is a routine occurrence that a concertgoer who has previously seen a performance with Maxine or has previously played with it will suggest that the system now exhibits a greater sense of "maturity," that it has "improved," or that it "sounds better these days." Even though nothing about the system has changed in terms of how it processes information or responds to live input, numerous interlocutors have expressed a more positive judgment of its behavior over repeated interactions.

## **4.5 Sources of the Illusion of Adaptation and a Benchmark for ML in HCI**

The sensation that a system with no computationally explicit capacity for adaptation has nevertheless somehow improved or adjusted to its interactants underscores the need for a phenomenological perspective in HCI. Just as Husserl insists that the study of experience must involve a suspension, or "bracketing" (Husserl 1913/2012), of analytical concern for the factual basis of such sensations, it is imperative that HCI researchers attune to what human interlocutors experience without being immediately concerned with whether those experiences have a factual basis. Regardless of what a designer (or any creative artist) intends, one must recognize that users or an audience are likely to receive or experience something that diverges from those intentions. Still, while Husserlian bracketing implies the need to ignore the factual, physical basis of experience in order to examine experience itself, the task of understanding what physical or technical facts produced those experiences (however illusory they may be) remains essential.

### ***4.5.1 Dynamic Feedback Loops***

Even when virtual free improvisers are designed with no capacity for adapting to other players over time, improvisers repeatedly experience them as adaptive for some reason or another. As odd as this collective hallucination may be, the consistency of such misperceptions about the supposedly "adaptive" behavior of a system which possesses absolutely no computational capacity for such a thing reveals quite a bit about the nature of such social interactions and the human social psychology into which ML-driven HCI fits. On one level, this illusion of adaptation is partially a result of the fact that human interactants naturally tend to gradually adapt to such

systems over time. Due to the human effort to adapt to and accommodate another social interactant, the collective behavior resulting from interactions with the system (or human agent) may shift. Even though the system's protocol of translating input to output remains static, the fact that the human player has adapted to the system means that the system is now receiving input of a different kind. As a result of changes in the nature of this input, it is reasonable to expect that the system's behavior will also be different from the initial encounter between the human player and the system. In sum, because of the changes that take place through this (rather one-sided) process of adaptation, the human player is likely to experience an illusion that the system has somehow "adapted" to them.

Regardless of whether the system is actually capable of it or not, the human player's shifting approach to playing with the system conjures a different element of the system's behavior, thereby suggesting that the system might have now adapted. Reporting on a comparative test of two interactive virtual free improvisers (Hsu and Sosnick 2009), computer scientists William Hsu and Marc Sosnick suggest that one system's interactive tendencies caused two human test subjects to drastically alter their typical behavior in improvisational interaction. Because one system, *ARHS*, is built to be more sensitive to short-term changes within an improvised piece, this system seems to encourage both musicians to play with rapid transitions and "choppy" material. This change in the musicians' performance in turn causes the *ARHS* system to make frequent adjustments, resulting in a dynamic feedback loop (ibid, 28).

Even though neither system tested was based in any ML technique, the results of Hsu' and Sosnick's tests have implications for an investigation into how a human being might experience the evolution of their rapport with an ML-based interactive system. A human being's natural tendency to adapt to the system's behavior yields a noticeable change in the system's behavior, resulting in an illusion or at least a superficial sensory and experiential trace of adaptation. Consequently, such results suggest that it is necessary to demonstrate that the adaptation that occurs through the coupling of the human being and an ML-based system is superior to or distinct from the illusory adaptation that occurs with a non-ML system.

Notably, the human tendency to adapt in this manner is hardly unique to an interaction with a non-human social agent built to perform as if it were a human social interactant. In numerous studies of human social interaction, human beings have been shown to exhibit a tendency to adjust to the behavioral tendencies of their interlocutors, even if these adjustments are only for the purposes of facilitating communication in a specific context (Kendon 1990; Giles et al. 1991; see also Lakin and Chartrand 2003). The human adaptive behavior that Linson and Hsu describe are likely examples of sociologist Erving Goffman's classic theory of "face-work" (1955; see also 1967), or the hypothesis that human beings tend to avoid exposing themselves or others to embarrassment. In encounters with interactive virtual performers, human improvisers simply adapt to the non-human interactant in order to try to make the most of the musical occasion. The impulse to make do and cope with the shortcomings of one's musical collaborator is repeatedly referenced both by Linson's test subjects (2014) as well as in numerous ethnomusicological studies

of musical interaction (Brinner 1995; Monson 1996; see also Sunardi 2011). When faced with the task of working with a musical partner whose skills are weaker than one's own, most musicians do what they can to prevent the partner's failings from being exposed to the audience by engaging in a variety of improvisatory means of dealing with their inadequacies.

### ***4.5.2 The Exposure Effect***

Aside from the general human tendency to accommodate one's interlocutors, the so-called effect of "mere exposure" (Zajonc 1968) or "familiarity principle" (see Moreland and Zajonc 1982) is another factor contributing to the almost hallucinatory judgment that a system with no capacity to adapt has nonetheless "adapted" to its partner. As proposed by Robert Zajonc, the exposure effect is a basic psychological inclination for human beings to have more positive opinions of stimuli they have previously encountered. As troubling as it may be to accept that human beings have a high proclivity to have more favorable evaluations of that with which they are already acquainted, an analysis of over 200 published experimental investigations of this phenomenon consistently indicates that both brief and repeated exposure yields more positive appraisals (Bornstein 1989).

To be certain of how or if familiarity and mere exposure effects play a role in this context, such effects would need to be studied more precisely in the context of the human experience of artificial social interactants like the virtual free improvisers at the center of the present discussion. Nevertheless, the exposure effect has been documented in numerous domains of activity including interpersonal attraction (Reis et al. 2011), musical taste (Huron 2006), and scholarly reputation (Serenko and Bontis 2011). Therefore, it is likely that the same effect is at work in the perception of adaptation in the human experience of interaction with these nonadaptive artificial socio-musical interactants. In other words, it is just as likely that one experiences a positive sensation of rapport formation over iterative interactions with a nonadaptive social agent as one would with an agent that is designed to exhibit adaptation through ML or related methods. Just as the human tendency to adapt to an interlocutor whose overall interactive processes are static still produces a sensation of "adaptation," the exposure effect presents a challenge to any researcher working with ML in the context of designing an artificial social interactant.

### ***4.5.3 A New Benchmark for ML in Artificial Social Interactants?***

Overall, the experience of adaptation with systems which do not fundamentally change their process of interacting with human interlocutors suggests the need for new benchmarks for the evaluation of such systems. That is to say, researchers

designing virtual social interactants using ML must convincingly demonstrate that the adaptation that the system's human interlocutors experience exceeds the rate of improvement which would result from mere exposure. Similarly, designers might also want to critically examine whether an ML-based system results in a sense of adaptation which is superior to what would result from the interactions of a human interactant who generously accommodates the inadequacies of its nonhuman partner. Ultimately, if one wants to claim that an ML-based artificial social interactant exhibits greater levels of adaptation than one which uses no adaptive systems technique, then it must be shown that the experience of adaptation is somehow superior to that which results from the exposure effect or a dynamic feedback loop.

## 4.6 Inverting HCI for Social Psychology

The regularity with which improvisers feel that they are forming a rapport with a machine that has no memory of the history of their interactions presents a fascinating set of problems for the intersection of ML and HCI. Such illusions in human-machine interaction have been noted since Joseph Weizenbaum wrote of his consternation at the way human beings were gleefully bamboozled into thinking they were making real psychological progress on personal issues by talking to "Eliza," an early chatbot designed to simulate the neutral therapy style of a Rogerian psychologist (Weizenbaum 1976; see also Hofstadter 1995). These and other perplexingly gullible experiences in the human encounter with technology further corroborate Spiro Kiouisis' provocative yet pragmatic suggestion that "interactivity" is less a feature inherent to the design of a system so much as it is a phenomenological dimension of how a human user frames and understands their experience with a given technology (Kiouisis 2002). Systems with strong interactive capacities may nevertheless lead human interactants to regard them as static or intransigent. Conversely, systems with relatively fixed interactive approaches may be regarded by their human interactants to have exhibited a strong ability to change their behavior over time.

But while the illusions of adaptation discussed above challenge the efficacy of ML as a way of recreating the kind of rapport that emerges in a real interpersonal relationship, what implications, if any, do they have for an understanding of the mutual attunement that takes place between human beings? Regardless of the reality or fantasy of adaptation between human and machine improvisers, such illusions raise questions about the precise nature of the interpersonal process of rapport formation that takes place as social beings get to know one another over time through a series of interactions. When we feel that others have adapted to us, what has actually taken place? Has a true change in their interactive behavior occurred? Or have we simply found a more favorable impression as a result of an exposure effect or a dynamic feedback loop?

For improvising musicians, the formation of a group dynamic is both a widely discussed experience but also a normative value which defines what improvisers

expect in the conduct of their peers. For example, improvising saxophonist Evan Parker, in a discussion with guitarist Derek Bailey, describes the experience of regularly playing in a group over a long period of time as one in which each participant had “accepted long ago those aspects of each other’s playing that we are never going to be able to change” (Parker, qtd. in Bailey 1980/1993: 141). Likewise, percussionist Burkhard Beins characterizes adaptation and acceptance of the tendencies of one’s peers as an inevitable fact of playing over and over again. Strikingly, Beins writes that “collective spaces of possibility already begin to establish themselves when the same group constellation meets for a second time after having formed some initial common experiences.” Furthermore, he notes that “this phenomenon appears to take place regardless of whether those who are involved are actually are of it” (Beins 2011: 171).

Likewise, testing the virtual free improviser Maxine has prompted numerous improvisers to assert that the system’s lack of memory of previous interactions is a significant barrier to their ability to feel that they are interacting with another human player. According to these players, unless I were to redesign the system to be able to acquire this kind of memory, this project is doomed to fail in the goal of reproducing the intuitive mutual understanding and rapport of human relationships. Such comments mainly arise when musicians press me to explain more about how the system works and clarify whether it can recall past events. While my goal has been to solicit their opinions of the system solely on the basis of their experience of playing with it (and not based on their evaluation of my verbal account of how it works), I have readily admitted that the system lacks any capacity to recall a previous interaction with a given performer. For several performers, this lack of memory is a major source of disappointment and immediately gives them a strong bias against the possibility that their experience playing with the system will be satisfying. For one performer, Maxine’s lack of memory of previous interactions was the aspect of the system’s makeup that most strongly made her feel that it would never be “life-like.” Without this sense of memory, she felt that the system lacks the basic human ability to be self-critical and improve one’s abilities to work with other musicians over time in order to create more compelling performances, howsoever this may be defined.

Clearly, these musicians share a strong conviction that a process of adaptation is (1) taking place regardless of one’s intentions and (2) a fundamental aspect of the successful development of a musical group engaged in this particular performance practice. But no matter how strongly these musicians may feel that such adaptation is happening, it remains a mystery as to what actually occurs that contributes to their experience of adaptation. Moreover, the fact that musicians also experience the formation of a rapport in interactions with a system whose interactive processes stay static raises the question of just how or to what degree improvisers actually adjust their approach to interaction over successive meetings to play.

If one experiences “adaptation” with an artificial social agent whose ways of interacting with input are essentially intransigent, then what does this experience suggest about the interpersonal attunement that musicians assert is both a fact and necessity in their improvisatory transactions with others? Is it possible that when

musicians claim to experience adaptation with their peers that what actually takes place is similar to the illusory one-sided “adaptation” that transpires between an adaptive musician and a nonadaptive system? What, if anything, does one learn about another player’s playing from improvising with them regularly both in private “sessions”<sup>9</sup> and concerts? What features of another’s playing does one adapt to? Minute features such as rhythm or timing habits or the spectral features of their playing (e.g., tone-to-noise ratio, spectral centroid, etc.)? Or more general features such as their tendency to create form or their disinterest in doing so? And besides the possibility of taking note of how others produce such features, what does one actually do to react to them? Mimic them? Deviate from them? A combination of the two?

As it currently stands, the answers to these questions are not yet known, though recent experimental work in the behavioral science of music offers a few beginnings of an idea of what actually happens in this process of interpersonal adaptation (Canonne 2013; Wilson and MacDonald 2015). In order to answer such questions, one obvious approach would be to analyze how a group of individuals develops a particular dynamic as a result of each other’s tendencies and their interplay over time. This could be done through the standard musicological means of transcribing the musical events captured in a recording and then using this data to conceptualize how the group’s habits of interaction change over time. But while this approach may reveal a handful of general tendencies in how human beings habituate themselves to the quirks of their interlocutors, there is no way of knowing if a truly mutual adaptation has taken place or if this is a repetition of the unilateral self-reconfiguration in which one musician copes with the inflexibility of another player, whether human or machine. Though the resultant dynamic feedback loop may resemble a process of transformation in group dynamic due to mutual, bilateral adjustment, it is just as likely that what has actually occurred is the result of the one party’s adjustment to the other. Similarly, as social psychologists have also observed, human subjects have a difficult time accurately assessing the nature of the rapport between two individuals when asked to do so from a third-person perspective in which one is not a participant of the interaction (Bernieri and Gillis 1995; Bernieri et al. 1996; Grahe and Bernieri 1999).

Therefore, in addition to the standard musicological and social-scientific methods for addressing the mystery of group dynamic formation, a comparison of players’ experiences improvising with (1) systems that do not adapt versus (2) systems that do may offer a clearer picture of what actually happens to produce the sensation of rapport. Though the field of HCI traditionally instrumentalizes the testing of such

---

<sup>9</sup>When improvisers meet in private, it is rare for them to “rehearse” materials. Instead, it is far more common to play for a duration similar to that of an actual concert (ranging from 20 minutes to an hour) without break. Afterwards, some discussion may take place about the music. However, given the fact that recalling specific details of such a long duration of improvisation, in which temporal coordination (i.e., pulse) is often absent and each player is engaged in significantly independent lines of action, it is doubtful that one will have a clear recollection of specific events. Therefore, it is unlikely as well that one will have the epistemological certainty required to make a comment about what has happened and how it should have been done differently (Corbett 1994).

systems as a path towards better coupling with the human interactant, I suggest that testing offers results which transcend the canonical expectation for “implications for design” (Dourish 2006). Instead, this approach to testing could provide a better picture of how interpersonal attunement occurs and the kind of adjustments a given kind of interlocutor expects or wants from the other. For the social psychology of rapport, comparative testing of artificial social interactants which use ML versus those which do not could offer an alternative means of studying the nature of rapport itself. The illusory sensation of adaptation discussed previously raises the question of whether what we take as a shared history with an interactant is an experience containing elements which are less the result of their adaptation and perhaps partially the result of a more static interactive process. Thus research into the design and refinement of virtual social interactants (Gratch et al. 2006; Cassell et al. 2007; Huang et al. 2011) not only allows programmers to create more believable interactants, but to also address questions in the social psychology of rapport and interpersonal adaptation.

“Repurposing” (Banerji 2012) testing in this manner leads HCI away from design and towards sociocultural study. At the same time, if HCI, regardless of the paradigm one is most aligned with, remains concerned with a fit between machines and human beings, then a study of the culture and experience continues to be a resource for this goal. Likewise, this approach to testing and reconceptualization of its true value begins to suggest a whole range of applications of HCI research that reach over and above this practical agenda to examine questions of significance for the behavioral sciences or humanities (Bardzell and Bardzell 2016). While HCI can, and should, continue to focus on increasing functionality and usability, these other goals beyond design itself are what make the field’s third-wave thinking have purchase for a broader realm of intellectual inquiries.

**Acknowledgements** Earlier versions of this chapter benefitted from the insightful critiques of Nick Seaver, Zachary Chase Lipton, as well as editors and reviewers of this volume. Financial support for this project came from the Fulbright U.S. Young Journalist’s Fellowship in Germany, the Berlin Program for Advanced German and European Studies, and the Berkeley-Mellon Fellowship.

## References

- Amershi S, Cakmak M, Knox WB, Kulesza T (2014) Power to the people: the role of humans in interactive machine learning. *AI Mag* 35(4):105–120
- Andersen K, Brooks R (2003) Special guest: Rodney Brooks. Studio 360: WNYC. National Public Radio, Washington, DC
- Assayag G, Dubnov S (2004) Using factor oracles for machine improvisation. *Soft Comput* 8(9):604–610
- Assayag G, Bloch G, Chemillier M, Cont A, Dubnov S (2006) Omax brothers: a dynamic topology of agents for improvisation learning. In: *Proceedings of the 1st ACM workshop on audio and music computing multimedia*, pp 125–132



- Assayag G, Bloch G, Cont A, Dubnov S (2010) Interaction with machine improvisation. In: Argamon S, Burns K, Dubnov S (eds) *The structure of style*. Springer, Berlin and Heidelberg, pp 219–245
- Backstrom MJ (2013). The field of cultural production and the limits of freedom in improvisation. *Crit Stud Improv/Études critiques en improvisation* 9(1)
- Bailey D (1980/1993) *Improvisation: its nature and practice in music*. Da Capo Press, New York
- Banerji R (2010) Maxine Banerji: the mutually beneficial practices of youth development and interactive systems development. *eContact! J Can Electroacoust Commun* 12(3). [https://econtact.ca/12\\_3/banerji\\_maxine.html](https://econtact.ca/12_3/banerji_maxine.html)
- Banerji R (2012) Maxine's Turing test: a player-program as co-ethnographer of socio-aesthetic interaction in improvised music. In: *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE'12)*
- Banerji R (2016) Balancing defiance and cooperation: the design and human critique of a virtual free improviser. In: *Proceedings of the international computer music conference*, pp 49–54
- Banerji R, Maxine, the Astromusicologist (2014) Palmer square. pfMentum, San Diego
- Bardzell J, Bardzell S (2011) Pleasure is your birthright: digitally enabled designer sex toys as a case of third-wave HCI. In: *Proceedings of the SIGCHI conference on human factors in computing systems*, pp 257–266
- Bardzell J, Bardzell S (2016) Humanistic HCI. *Interactions* 23(2):20–29
- Baumer EP, Silberman M (2011) When the implication is not to design (technology). In: *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, New York, pp 2271–2274
- Beins B (2011) Entwurf und Ereignis (Scheme and Event). In: Beins B, Kesten C, Nauck G, Neumann A (eds) *echtzeitmusik berlin: selbstbestimmung einer scene (self-defining a scene)*. Wolke Verlag, Hofheim
- Bernieri F, Gillis JS (1995) The judgment of rapport: a cross-cultural comparison between Americans and Greeks. *J Nonverbal Behav* 19(2):115–130
- Bernieri FJ, Gillis JS, Davis JM, Grahe JE (1996) Dyad rapport and the accuracy of its judgment across situations: a lens model analysis. *J Pers Soc Psychol* 71(1):110
- Blackwell T, Bentley P (2002) Improvised music with swarms. In: *Proceedings of the World conference on computational intelligence*, pp 1462–1467
- Blackwell T, Young M (2004) Self-organised music. *Organised Sound* 9(2):123–136
- Bødker S (2006) When second wave HCI meets third wave challenges. In: *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*. ACM, New York, pp 1–8
- Boehm C (1993) Egalitarian behavior and reverse dominance hierarchy. *Curr Anthropol* 34(3):227–254
- Boehner K, DePaula R, Dourish P, Sengers P (2007) How emotion is made and measured. *Int J Hum Comput Stud* 65(4):275–291
- Borgo D (2002) Synergy and Surrealstate: the orderly disorder of free improvisation. *Pac Rev Ethnomusical* 10:1–24
- Bornstein RF (1989) Exposure and affect: overview and meta-analysis of research, 1968–1987. *Psychol Bull* 106(2):265–289
- Bown O (2011) Experiments in modular design for the creative composition of live algorithms. *Comput Music J* 35(3):73–85
- Bown O (2015) Player responses to a live algorithm: conceptualising computational creativity without recourse to human comparisons? In: *Proceedings of the sixth international conference on computational creativity*
- Brinner B (1995) *Knowing music, making music: Javanese gamelan and the theory of musical competence and interaction*. University of Chicago Press, Chicago
- Burrell J (2016) How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data Soc* 3(1):1–12
- Canonne C (2013) Focal points in collective free improvisation. *Perspect New Music* 51(1):40–55
- Canonne C, Garnier N (2015) Individual decisions and perceived form in collective free improvisation. *J New Music Res* 44(2):145–167

- Cappella JN (1990) On defining conversational coordination and rapport. *Psychol Inq* 1(4):303–305
- Carey B (2012) Designing for cumulative interactivity: the \_derivations system. In: Proceedings of the international conference on new interfaces for musical expression
- Carles P, Comolli J-L (2015 [1971]) *Free jazz/Black power*. University Press of Mississippi, Mississippi
- Casal DP (2008) Time after time: short-circuiting the emotional distance between algorithm and human improvisors. In: Proceedings of the international computer music conference
- Casal DP, Morelli D (2007). Remembering the future: applications of genetic co-evolution in music improvisation. In: Proceedings of the European conference on artificial life
- Cassell J, Gill AJ, Tepper PA (2007) Coordination in conversation and rapport. In: Proceedings of the workshop on embodied language processing. Association for Computational Linguistics, pp 41–50
- Chadabe J (1997) *Interaction*. In: *Electric sound: the past and promise of electronic music*. Prentice-Hall, Upper Saddle River, pp 286–323
- Chaudron C (1984) The effects of feedback on students' composition revisions. *RELC J* 15(2):1–14
- Collins N (2008) Reinforcement learning for live musical agents. In: Proceedings of the International Computer Music Conference (ICMC), Belfast
- Cont A, Dubnov S, Assayag G (2006) A framework for anticipatory machine improvisation and style imitation. In: Proceedings of the conference on Anticipatory Behavior in Adaptive Learning Systems (ABiALS)
- Corbett J (1994) Evan Parker: saxophone botany. In: *Extended play: sounding off from John Cage to Dr. Funkenstein*. Duke University Press, Duke, pp 201–208
- Corbett J (2016) *A Listener's guide to free improvisation*. University of Chicago Press, Chicago
- Crabtree A, Rodden T, Tolmie P, Button G (2009) Ethnography considered harmful. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp 879–888
- Dourish P (2006) Implications for design. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp 541–550
- Edwards PN (1997) *The closed world: computers and the politics of discourse in Cold War America*. MIT Press, Cambridge, MA
- Eiben AE, Smith JE (2003) *Introduction to evolutionary computing*. Springer, Berlin/Heidelberg
- Eklund R (2002) Ingressive speech as an indication that humans are talking to humans (and not to machines). In: Proceedings of the 7th International Conference on Spoken Language Processing (ICSLP)
- Fails JA, Olsen DR Jr (2003) Interactive machine learning. In: Proceedings of the 8th international conference on Intelligent user interfaces. ACM, New York, pp 39–45
- Fiebrink R, Cook PR, Trueman D (2011) Human model evaluation in interactive supervised learning. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp 147–156
- Fischlin D, Heble A, Lipsitz G (2013) Take two: sounds of surprise; rights, risks and responsibilities in improvised music. In: *The fierce urgency of now: improvisation, rights, and the ethics of co-creation*. Duke University Press, Duke, pp 203–219
- Forsythe D (2002) *Studying those who study us: an anthropologist in the world of artificial intelligence*. Stanford University Press, Stanford
- Friedman B, Nissenbaum H (1996) Bias in computer systems. *ACM Trans Inf Syst (TOIS)* 14(3):330–347
- Fuchsberger V, Moser C, Tscheligi M (2012) Values in action (ViA): combining usability, user experience and user acceptance. In: *CHI'12 extended abstracts on human factors in computing systems*. ACM, New York, pp 1793–1798
- Giles H, Coupland N, Coupland I (eds) (1991) *Contexts of accommodation: developments in applied sociolinguistics*. Cambridge University Press, Cambridge
- Goffman E (1955) On face-work: an analysis of ritual elements in social interaction. *Psychiatry Interpersonal Biol Process* 18(3):213–231
- Goffman E (1967) *Interaction ritual: essays in face-to-face behavior*. Anchor Books, New York

- Grahe JE, Bernieri FJ (1999) The importance of nonverbal cues in judging rapport. *J Nonverbal Behav* 23(4):253–269
- Gratch J et al (2006) Virtual Rapport. In: Gratch J, Young M, Aylett R, Ballin D, Olivier P (eds) *Intelligent Virtual Agents. IVA*. Springer, Berlin/Heidelberg, pp 14–27
- Gratch J, Wang N, Gerten, J., Fast, E., & Duffy, R. (2007). Creating rapport with virtual agents. In: Pelachaud C, Martin J-C, André E, Chollet G, Karpouzis K, Pelé D (eds) *International Workshop on Intelligent Virtual Agents*. Springer, Berlin, pp 125–138
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. *alt. chi*
- Helmreich S (2001) After culture: reflections on the apparition of anthropology in artificial life, a science of simulation. *Cult Anthropol* 16(4):612–627
- Hofstadter DR (1995) The ineradicable Eliza effect and its dangers. In: Hofstadter DR (ed) *Fluid concepts and creative analogies: computer models of the fundamental mechanisms of thought*. Basic books, New York, pp 155–168
- Hsu W (2005) Using timbre in a computer-based improvisation system. In: *Proceedings of the international computer music conference*
- Hsu W (2010) Strategies for managing timbre and interaction in automatic improvisation systems. *Leonardo Music J* 20:33–39
- Hsu W, Sosnick M (2009) Evaluating interactive music systems: an HCI approach. In: *Proceedings of the international conference on New Interfaces in Musical Expression (NIME)*
- Huang L, Morency L-P, Gratch J (2011) Virtual rapport 2.0. In: Vilhjálmsson HH, Kopp S, Marsella S, Thórisson KR (eds) *Proceedings of the 10th International Conference on Intelligent Virtual Agents*. Springer, Berlin/Heidelberg, pp 68–79
- Huron DB (2006) *Sweet anticipation: music and the psychology of expectation*. MIT Press, Cambridge, MA
- Husserl E (1913/2012) *Ideas: general introduction to pure phenomenology* (WRB Gibson, Trans.). Routledge
- Kendon A (1990) *Conducting interaction: patterns of behavior in focused encounters*. Cambridge University Press, Cambridge
- Kennedy J, Eberhart R (1995) Particle swarm optimization. In: *Proceedings of the IEEE conference on neural networks*, pp 1942–1948
- Kiousis S (2002) Interactivity: a concept explication. *New Media Soc* 4(3):355–383. <https://doi.org/10.1177/146144480200400303>
- Kofsky F (1970) *Black nationalism and the revolution in music*. Pathfinder Press, New York
- Lakin JL, Chartrand TL (2003) Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychol Sci* 14(4):334–339
- Lange BR (2011) Teaching the ethics of free improvisation. *Crit Stud Improvis/Études critiques en improvisation* 7(2)
- Langton CG (1997) *Artificial life: an overview*. MIT Press, Cambridge, MA
- Lewis GE (1993). *Voyager* (G Lewis, R Mitchell, V. I. V. I. System, Trans.). Avant, Japan
- Lewis GE (1999) Interacting with latter-day musical automata. *Contemp Music Rev* 18(3):99–112
- Lewis GE (2000) Too many notes: computers, complexity and culture in *Voyager*. *Leonardo Music J* 10:33–39
- Lewis GE (2007) *Mobilitas animi: improvising technologies, intending chance*. Parallax 13(4):108–122
- Linson A (2014) *Investigating the cognitive foundations of collaborative musical free improvisation: experimental case studies using a novel application of the subsumption architecture*. Doctoral thesis, The Open University
- Linson A, Dobbyn C, Lewis GE, Laney R (2015) A Subsumption agent for collaborative free improvisation. *Comput Music J* 39(4):96–115
- Maatman R, Gratch J, Marsella S (2005) Natural behavior of a listening agent. In: *Intelligent virtual agents*. Springer, Heidelberg, pp 25–36
- Monson I (1996) *Saying something: jazz improvisation and interaction*. University of Chicago Press, Chicago
- Moreland RL, Zajonc RB (1982) Exposure effects in person perception: familiarity, similarity, and attraction. *J Exp Soc Psychol* 18(5):395–415

- Morrissey K, Kirakowski J (2013) 'Realness' in Chatbots: establishing quantifiable criteria. In: Kurosu M (ed) Proceedings of the 15th international conference on human-computer interaction. Springer, Berlin/Heidelberg, pp 87–96
- Nissenbaum H (2001) How computer systems embody values. *Computer* 34(3):120–119
- Pachet F (2003) The continuator: musical interaction with style. *J New Music Res* 32(3):333–341. <https://doi.org/10.1076/jnmr.32.3.333.16861>
- Pras A, Schober MF, Spiro N (2017) What about their performance do free jazz improvisers agree upon? A case study [Original Research]. *Front Psychol* 8(966). <https://doi.org/10.3389/fpsyg.2017.00966>
- Reis HT, Maniaci MR, Caprariello PA, Eastwick PW, Finkel EJ (2011) Familiarity does indeed promote attraction in live interaction. *J Pers Soc Psychol* 101(3):557
- Rodriguez AW (2016) Harmolodic pedagogy and the challenge of omni-musicality. *Jazz Perspect* 9(2):173–192
- Rowe R (1992) *Interactive music systems: machine listening and composing*. MIT Press, Cambridge, MA
- Seaver N (2012) Algorithmic recommendations and synaptic functions. *Limn* 1(2)
- Sengers P, Kaye J, Boehner K, Fairbank J, Gay G, Medynskiy Y et al (2004) Culturally embedded computing. *IEEE Pervasive Comput* 3(1):14–21
- Serenko A, Bontis N (2011) What's familiar is excellent: the impact of exposure effect on perceived journal quality. *J Informet* 5(1):219–223
- Seyfert R, Roberge J (2016) *Algorithmic cultures: essays on meaning, performance and new technologies*. Routledge, New York
- Snyder JK, Fessler DM (2014) Narcotics anonymous: anonymity, admiration, and prestige in an egalitarian community. *Ethos* 42(4):440–459
- Spellman AB (1966) *Four lives in the bebop business*. Pantheon Books, New York
- Stanyek J (1999) Articulating intercultural free improvisation: Evan Parker's synergetics project. *Resonance* 7(2):44–47
- Steinbeck P (2010) "Patience, sincerity, and consistency": Fred Anderson's musical and social practices. *Crit Stud Improvis/Études critiques en improvisation* 6(2)
- Suchman L (2006) *Human-machine reconfigurations: plans and situated actions*. Cambridge University Press, Cambridge
- Sunardi C (2011) Negotiating authority and articulating gender: performer interaction in Malang, East Java. *Ethnomusicology* 55(1):32–54
- Tickle-Degnen L, Rosenthal R (1990) The nature of rapport and its nonverbal correlates. *Psychol Inq* 1(4):285–293
- Waterman E (2008) Naked intimacy: eroticism, improvisation, and gender. *Crit Stud Improvis/Études critiques en improvisation* 4(2)
- Weizenbaum J (1976) *Computer power and human reason*. W. H. Freeman and Company, San Francisco
- Wilf E (2013a) Sociable robots, jazz music, and divination: contingency as a cultural resource for negotiating problems of intentionality. *Am Ethnol* 40(4):605–618
- Wilf E (2013b) Toward an anthropology of computer-mediated, algorithmic forms of sociality. *Curr Anthropol* 54(6):716–739
- Wilson GB, MacDonald RA (2015) Musical choices during group free improvisation: a qualitative psychological investigation. *Psychol Music*:1–15
- Woodburn J (1982) Egalitarian societies. *Man* 17(3):431–451
- Yee-King MJ (2007) An automated music improviser using a genetic algorithm driven synthesis engine. In: Giacobini M (ed) *Applications of evolutionary computing (EvoWorkshops 2007) proceedings*. Springer, Berlin/Heidelberg, pp 567–576
- Young M (2008) NN music: improvising with a 'living' computer. In: Kronland-Martinet R, Ystad S, Jensen K (eds) *Computer music modeling and retrieval: sense of sounds*. Springer, Berlin, pp 337–350
- Young M (2010) Identity and intimacy in human-computer improvisation. *Leonardo Music J* 20(online supplement)
- Zajonc RB (1968) Attitudinal effects of mere exposure. *J Pers Soc Psychol* 9(2, part 2):1–27

# Chapter 5

## Interaction Design for Metacreative Systems



Oliver Bown and Andrew R. Brown

**Abstract** In this paper, we examine digital creativity as a collective activity performed through socio-technological networks of agency. We introduce metacreation—the automation of creative tasks with machines—as a domain that is usefully examined from a 3rd wave HCI approach. We discuss four general human-computer interaction activities that commonly appear in metacreation: (1) metagenerating form; (2) searching/finding; (3) helping machines learn; and (4) evaluation/iteration. These are not necessarily specific to metacreation, but nevertheless point to particular design considerations in a metacreative context. Four creative interaction design themes are considered in their relation to metacreation: direct manipulation and real-time control; supporting playful interaction and divergent goals; the programmatic design of behaviours, and; managing distributed creativity. We then identify three paradigms of interaction design for metacreation: operation-based interaction, involving the direct manipulation of generative algorithms; request-based interaction, involving the submission of requests to a system that returns results; and ambient interaction, that involves the operation of autonomous metacreative processes in the background. Our discussion of these suggests possible trends for design: an increasingly complex and modular future for networked human-machine digital creativity; an increasing role for request-based metacreative systems where users specify, rather than construct, outcomes; the increasing role of metacreation in ‘prosumer’ content creation; and, consequently, the reduction of labour involved in creating media. The chapter makes clear, we hope, that metacreative practices present unique challenges and opportunities for interaction design.

---

O. Bown (✉)  
University of New South Wales, Sydney, Australia  
e-mail: [o.bown@unsw.edu.au](mailto:o.bown@unsw.edu.au)

A. R. Brown  
Griffith University, Nathan, Australia

## 5.1 Introduction

Over the past half century, the design of creative software systems that semi-automate creative tasks—referred to here as metacreative systems—has moved steadily from concept to reality. Such systems might generate melodic content for use by a film composer, automatically mix or master music recordings, create humorous tweets relevant to a particular topic, design buildings that satisfy structural and aesthetic constraints, or perform stylistic renderings of images. Although it remains a largely experimental field, in recent years real applications of metacreation<sup>1</sup> have necessitated the consideration of how we interact with these systems, raising issues of human computer interaction (HCI) and the adaptation of existing HCI practices to the field (Bown 2014; Kantosalo et al. 2015). Indeed it has become apparent that some of the substantial challenges in metacreation involve issues of the interaction between people and smart machines. Metacreation requires new interaction paradigms that involve semi-autonomous computational processes, and that may also be applicable to other areas of interaction design.

*Metacreation*<sup>2</sup> involves the use of computational methods to shift the balance of creative agency from people toward machines with the assistance of artificial intelligence (AI) or related methods. Metacreation has been conceptualized in various ways. The terms *algorithmic* or *procedural content creation* (e.g., Langston 1989; Togelius et al. 2011) tend to refer to the use of computers to procedurally generate patterns and variations - for example, using parametric and combinatoric methods. *Generative art* often refers to the creative practice of software-based artists in which the software has some degree of autonomy (Galanter 2003). *Computational creativity* (CC) is another catch-all term that some within the community have used to emphasise non-trivial system autonomy, which is distinguished from ‘mere generation’ because CC systems perform evaluation of the artefacts they produce (Ventura 2016). In other words, systems fulfilling this ‘lofty’ idea of CC have an ‘aesthetic sense’—or, one could say—they exhibit something approximating ‘taste’. Another term in use is *creative AI*, which is less strictly concerned with the creative agency of computational systems, and includes any application of AI in creative domains.

Interaction design for metacreative systems requires us to reframe HCI thinking to take into account the greater contextual awareness and decision making processes that devices can take on. 3rd wave HCI thinking is particularly well-suited to designing software that takes into account a modern understanding of the performance of co-creative tasks. These tasks include a focus on open-endedness, the collective construction of practices amongst a community of users, the creative re-appropriation

---

<sup>1</sup>From machine learning-based music composition (Ghedini et al. 2016) to distributed evolutionary computing websites (Secretan et al. 2008).

<sup>2</sup>The term *metacreation* itself can be attributed to Whitelaw (2004), and although his interest in using it was with respect to artificial life-based art, we interpret it as having a wider remit. Quoting Schöffer: “we are no longer creating a work, we are creating creation... we are able to bring forth... results... which go beyond the intentions of their originators” (Whitelaw 2004: 17).

of technologies, personalisation, and the configuration of complex hybrid networks of technologies (Bødker 2006, 2015).

Creativity theorists are largely agreed that creative problems are, by their nature, unmapped; we cannot follow a known path to a solution, but must conduct searches for solutions and employ heuristic methods of navigation (Wallas 1926; Perkins 1994; Simonton 1999; Amabile 1996). Exploratory creative strategies exist in creative behaviours at a multiplicity of levels. Classical thinking about individual creative cognition recognises ‘incubation’ as a key process, in which multiple candidate solutions are evaluated subconsciously before being presented to conscious awareness in a stage called ‘illumination’ (Wallas 1926). Over longer timescales, individual creative practitioners exhibit a tendency to speculatively spread their attention in order to conduct broad searches (Perkins 1994; Simonton 1999). At the super-individual level, we see undirected search operating as a social process, with social institutions performing heuristic functions: the open market, prizes, funding bodies, and the social prestige of creative achievement. A consequence of this is that creative success is unevenly distributed, although it is related to the intensity and directedness of activity and influenced by factors like education, intrinsic motivation and talent (Simonton 1999).

Because creativity occurs as a set of exploratory, cumulative, and stochastic processes, then creative agency—the attribution of influences on creative outcomes—is necessarily found to be distributed across groups of people and things, that together form complex and fluid *networks of agency* (Brown 2016). This distributed agency is illustrated in Malafouris’ (2008) description of a potter shaping clay on a wheel, that involves a process of dynamic interaction where the potter’s hands shape the clay but also respond to it. Achieving a creative goal requires the potter to be aware of the physical affordances of the clay and, in some sense, to ‘negotiate’ with it rather than merely dictating its form. Contrary to our instinctive view of humans as the sole agents in such processes, this example highlights the value of conceptualising the various actors involved as each exerting some agency, with *interactions* within the human+clay+wheel+culture hybrid being the source of creativity.

Likewise, in the authors’ own practices, making electronic music involves creative input inherited from the makers of the electronic instruments, but also the instruments’ own musical affordances which may transcend the intentions of their inventors, for example in the nonlinearities of electronic circuits.

These preliminary references set up the context in which we see the automation of creativity as meaningful. Making machines that are active contributors to the production of aesthetic artefacts seems at first glance—as was the case for many of the field’s early experimenters—to mean making *virtual artists*. Such was the dominant nature of the discussion, surrounding the work of well-known generative practitioners such as Harold Cohen (Cohen 1995) in visual art and David Cope (Cope and Mayer 1996) in music. In contrast, we bypass the question of whether computers can be ‘truly creative’, can ‘originate anything’, or fulfil the role of an artist, and focus instead on how the types of action and interaction that occur in hybrid socio-technological networks lead to creativity.

In this chapter, we provide an overview of some of the interaction design issues in metacreation, then proceed to consider the impact for the design and evaluation of metacreative systems within the extended contextual frame of 3rd wave HCI. We base our analysis on examples from various areas of musical metacreation. We propose three metacreation design paradigms that can introduce some structure to the field. Finally, we reflect on how metacreation is a constituent part of an advanced AI future and consider its potential impact in the spheres of creative work and daily life.

## 5.2 Making Metacreative Systems Usable: Some Areas for Consideration

Creative practices involving semi-autonomous dynamic systems present new interaction design challenges—from traditional HCI issues surrounding perceived affordances and usability, to emergent issues of how users react to software systems that convey a sense of autonomy and authorship, or that perform unorthodox actions such as challenging or directing a co-creative<sup>3</sup> activity. The latter imply profound transformations to the practice of working with computers, and are exemplary of the need for 3rd wave HCI frameworks and methods.

Algorithms that support metacreation include complex hand-coded rule-based systems, machine learning algorithms, evolutionary and other types of search or optimisation algorithms, along with reasoning and concept-manipulation methods.<sup>4</sup> There are a range of different users, from expert programmers to end-user creative practitioners, for whom different interaction design considerations will apply. We are particularly interested in the challenges of designing metacreative tools for non-programmer end users as a way of amplifying professional practices.

## 5.3 Human-Computer Interaction Activities Common to Metacreation

In this section we discuss four general human-computer interaction activities that commonly appear in metacreation: (1) generating forms, (2) searching and finding, (3) helping machines learn, and (4) evaluation and iteration. These are not

---

<sup>3</sup>Co-creativity refers to a situation in which human user and computer system are interacting, active participants in a creative process.

<sup>4</sup>An indication of the range of algorithms used in computational creativity research can be found by looking through the free online proceedings of the International Conference on Computational Creativity (<http://www.computationalcreativity.net>) or of edited volumes such as McCormack and D’Inverno (2012) and Besold et al. (2015).



necessarily specific to metacreation, but nevertheless point to specific design considerations in a metacreative context.

Following this, four creative interaction design themes are considered in their relation to metacreation: (1) direct manipulation and real-time control, (2) supporting playful interaction and divergent goals, (3) the programmatic design of behaviours, and (4) managing distributed creativity.

### 5.3.1 *Generating Forms*

The simplest application of metacreation is the programmatic generation of variation, whereby a creator defines a method of *representation* and a *space* of transformational possibilities. This space can then be explored by the creator using the metacreative system, or by third party creative producers, or an audience. For example, Jon McCormack's (McCormack 2004) series of 3D virtual plant forms used a generative structure called an L-system to establish an infinite search space of possibilities which he then searched using an interactive genetic algorithm. Similarly, in architecture, it is common to create parametric forms that vary by simple adjustment of a set of predefined parameters (Hudson 2008). Variations of these forms can then be explored to optimise building performance using readily measurable factors.

As a general observation, rich metacreative systems are likely to be complex and convoluted, and hence tend to appear as opaque black boxes to users. This sets up a tension with a core principle in interaction design—that of the system's increased visibility to the user through perceived affordances and constraints, mappings and system feedback (Norman 1988). This can be resolved through designs that present relevant information to a user that helps them maintain a model of what the system is doing, as best it can. But with complex metacreative systems that cannot easily be modelled there is a tension between the power of the system and the capacity that a user has to control it. Further, many machine learning systems used in such processes do not clearly reveal, even to their expert operators, the logic of their pattern matching. For the creation of metacreative systems we prefer, instead, to understand how the system operates to empower a user to pursue a plan of action or, at least, to understand the limits of actions (Winograd and Flores 1986). Equally, we might consider how playful and engaging the system is, so as to improve engagement, even if the user cannot maintain a clear model of the system behaviour (Pask 1976; Gaver et al. 2004).

### 5.3.2 *Searching and Finding*

Perhaps the most obvious way in which metacreative tools can support creativity is through search. Consider, for example, an advanced search feature on a database of artworks. Such a feature might combine filters or use natural language processing to

enable a query like: *show me artworks produced in nineteenth century France that are colourful, and involve people and city scenes*. Interfaces that support such queries might be further be used to direct a search, not through existing content, but through spaces of generative possibilities looking for instances that meet the descriptive criteria. With the current search infrastructure in place it is potentially a small step for companies like Google or Spotify to insert generated content into their search results.

*Target-based* automated searches involve a user specifying a set of target properties (in evolutionary computing this is known as a fitness function, and targets can also be presented in terms of sets of constraints) and having the system find outcomes that satisfy these targets (Davis 1991). Alternatively, forms of *novelty search* can be employed to offer diversity instead of specificity (Lehman and Stanley 2008). For example, given a parametric system such as a parametrically controlled creature morphology for a game, a novelty search process would show a range of possibilities offered by a system: for example, a set of prototypical creatures each of which is different in its own way.

Because in creative search a person may not necessarily have a clear idea of what they are searching for, outcomes may be more or less precisely specified at the beginning of the process. A co-creative search process may begin with a more speculative set of suggestions led by the system, which are then reflected upon and refined by the user. The user may be led to attend to areas they were not originally interested in, i.e., to explore rather than to optimise to a target.

Many hybrids of these basic search paradigms can be constructed. An automated system might be used to generate a new set of parameters by which a user can conduct a manual search (a nested search). For example, in multiobjective search, the system is asked to satisfy multiple conflicting objectives. The result is not a single best individual, but a set of potentially optimal individuals, defined by a *Pareto front*: a surface of possible compromise solutions that the user might interrogate manually (Konak et al. 2006). Alternatively, a system might do its own filtering to generate a suitably small set of prototypes that a user can feasibly skim through (as in novelty search) (Woolley and Stanley 2014).

### 5.3.3 *Helping Machines Learn*

Machine learning systems are increasingly used for content generation. This can involve learning specific relations (supervised learning), reinforcement learning (in which only good/bad feedback is given), or data-driven (unsupervised) learning. The most common machine learning applications are in perception or data mining, but certain types of machine learning systems are inherently capable of generation. Sequence learning systems such as Markov models and LSTM networks, for example, learn to predict the next event in a sequence given previous events. This means that when seeded with a sequence they can iteratively generate new content. It is not obvious that such systems would do anything besides regurgitate existing content

with little creative value, but it is interesting to begin to see the extent to which such systems' blind generation can be the basis for producing creatively insightful and novel outputs. Some kinds of image-based neural networks can adapt or generate images, utilising techniques such as style transfer—applying the low level features of one set of images to the high-level features of an input, e.g. painterly rendering (Bruckner and Gröller 2007). Machine learning systems that are incapable of generation can still be used as evaluation nodes in hybrid systems—for example, playing the role of a fitness function in an evolutionary process.

The selection of training data is key to training a machine learning system. A music system might need an interface to specify what style to emulate: feed it Bach and it plays Bach-like music, feed it Gershwin and it plays Gershwin-like music. However, for good on-the-fly learning with current state-of-the-art systems, it has proven necessary to configure the system itself and possibly select the correct way to represent the data (Martin et al. 2011). This requires knowledge of machine learning principles and/or expert knowledge of the domain, with novel interface design challenges to overcome if we want to put such tools in the hands of non-expert end users.

### 5.3.4 *Evaluation and Iteration*

Given an algorithmic system that generates candidate outputs, we would generally expect a user to engage in a feedback process, typically over several iterations. This is analogous to a client-producer relationship, such as when a company director interacts with a graphic designer regarding a logo design. Alternatively, an audience might give feedback about the output of a system over longer time scales, as in the DarwinTunes system (MacCallum et al. 2009, 2012), where an audience listens to a live radio stream of evolved music and votes on preferred segments of the output, feeding back into an evolutionary algorithm that aims to improve the music.

An important aspect of aesthetic evaluation is that individual judgement is far from fixed, and can be influenced by any number of factors external to the thing being evaluated. For example, Salganik and Watts (2008) show how music preferences are influenced by knowledge about what *other* people think of the same music; a winner-takes-all effect where music that is perceived to be popular is subsequently more likely to be rated highly (this effect could potentially be reversed amongst those with contrarian inclinations, where 'popular' co-equates with 'lowest common denominator'). Meanwhile Bloom (2010) presents examples in which the value associated with something is influenced by beliefs about the artefact such as its provenance (for example that it was owned by someone famous and therefore once in physical contact with them). Indeed, furthermore, evaluation should be understood as multi-dimensional, a rich engagement with objects and people involving disparate factors, rather than as a fixed scalar 'score' given to aesthetic artefacts.

Taking a longer-term perspective, we note that the lifelong iteration of creative techniques can be seen as part of a more substantial reflexive development of a creative practice for an individual. Currently we adapt the computational systems around us by configuring them, and we adapt our working practices to their affordances. Software systems do not automatically coevolve with us in the development of our creative practice, but they could. This necessitates looking at long term patterns of software use, and processes of mutual adaptation and habituation between people and the systems they use, an area in which HCI has started to establish good strategies as it develops into the 3rd wave.

### ***5.3.5 Creative Interaction Design Themes***

The forms of interaction discussed above occur in a range of metacreation contexts and are core foci of HCI design for metacreative systems. However, a number of other themes cross the boundary between existing digital creative practice and metacreative practice.

#### **5.3.5.1 Direct Manipulation and Real-Time Control**

Directly controlling a metacreative system is similar to the control of any other device. While advanced computing techniques such as machine learning or evolution may be used to generate artefacts—or to generate systems that generate artefacts—thought is required when interfacing them with direct manipulation by a user; such as a creative practitioner employing generative techniques. It can be productive, therefore, if generative processes can apply to standard objects, such as images or musical scores, and that multiple generative and direct manipulation processes are interoperable. In other scenarios, live *configuration* of some more or less complex network of interactions becomes a core part of the activity. For example, the ecosystemic compositions of Agostino di Scipio are predicated on using a human-managed feedback loop between speakers, an acoustic space, a microphone, and digital signal processing to simulate the emergence of complex patterns from the background room acoustics (di Scipio 2003).

#### **5.3.5.2 Supporting Playful Interaction and Divergent Goals**

Analyses of creative tasks indicate that goal-oriented behaviour—the primary focus of much software interaction design—is less applicable in creative domains where many tasks are open-ended, exploratory or playful in nature (Perkins 1994). Correspondingly, interface guidelines emphasise predictability and reliability (Nielsen 1993). This presents a potential misalignment between the static or dynamic nature of a system and its user interface. Complicating interaction design further is the potential for users to vary their objectives, especially between aims

that are utilitarian and speculative. For example, every aspect of a professional creative program such as Adobe Photoshop is designed to be stable and predicated on the idea of being ‘powerful’, offering the user a large number of goal-directed actions. Yet the creative practitioner might use it in very playful ways, engaging in odd experiments, appropriating quirks and errors to creative effect, and imposing peculiar constraints (Ferguson and Brown 2016). The intended interface goals of software are typically to provide leverage and ease of use. By comparison, computer games offer an alternative paradigm of software that does not simply aspire to make a task easy—quite the opposite—and the unexpected is often part of the pleasure of game play. This diversity of objectives exemplifies the reasons ‘user experience’ exists as a separate category from ‘usability’, and why the design of interfaces for metacreative systems often needs to privilege exploratory power so such systems can be effective in stimulating idea generation and discovery.

Even in the most professional context, creative production involves an element of playful search. Professional creative software tools can support this search by being predictable and powerful. But generative creative software might support such search by being proactive, challenging and surprising.

### 5.3.5.3 The Programmatic Design of Behaviours

Typically, metacreative systems are complex, bespoke pieces of software, and thus their behavioural algorithms are programmed as software algorithms. Given this, and despite the potential of graphical user interfaces (GUIs) for controlling specific metacreative systems, it may be the case that the expert manipulation of metacreative systems requires programming. In such systems the user is effectively operating in an open-ended context. A direct-manipulation widget-based GUI does not easily allow such open-ended configurations.

Because the current development of metacreative systems often involves software development in the traditional sense, programming is an important form of interaction to consider. The work on end-user programming, such as that of Blackwell (2002), serves this discussion. He identifies practitioners, besides professional software developers, who have to make regular choices between programming as a creative act, and what as we have seen, is termed ‘direct manipulation’ (from Shneiderman 1993). Programming is defined by a set of core cognitive tasks: requirements gathering, specification, design, coding and debugging (Blackwell 2002). With generative systems, regardless of the interface, the designer/user may face such tasks—particularly those resembling debugging or structural change—as a natural part of the exploratory nature of creative practices. The use of code as an interface enables this flexibility, but is not practicable as an interface for all, perhaps not even most, metacreative interaction contexts.

Meanwhile, software languages are becoming increasingly expressive and fluid. Some practitioners have taken programming as the key mode of interaction with metacreative systems in a performative framework. This is particularly prevalent in the practice of live coding which “involves writing and modifying computer pro-

grams that generate music in real time” (Brown and Sorensen 2009: 17). This performative approach to software development involves the on-the-fly programming of system behaviors and the management of them through code modification in real time. A core theme here is that of switching to an improvisational form of dynamic program creation using code as the interface.

### 5.3.5.4 Managing Distributed Creativity

The design of metacreative systems necessitates a distributed view of creative agency involving networks of people and things (Suchman 1987; Latour 1995; Rammert 2008).

In such networks, intelligence is extended beyond its conventionally assumed place in the human mind to include objects, machines and environments. In this view, “the human organism is linked with an external entity in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right” (Clark and Chalmers 1998: 12) and the interface between agents is critical to the effectiveness of interactions. Correspondingly from a HCI perspective, Bødker argues that frameworks that properly represent the multiplicity of interactions of real working environments, described as webs-of-technology, are increasingly needed: “we never design single, monolithic devices or systems but technology that must be seen and used in relation to many other devices, applications and systems.” (Bødker 2006: 3).

Echoing Latour’s concept of the *script* (Latour 2011) where individuals can be either independently (above the script) or passively (under the script) following the directions of their circumstances, Bødker also highlights how individual creative practitioners are themselves shaped in networked socio-technological interactions: “users’ shared capacities and experiences are not just based on individual acting and learning in the world; rather, they are bound to shared practices, joint activities, and so on, in artifact ecologies. It is against this background that the relationship between the user and the artifact exists.” (Bødker 2015: 27).

Such considerations, in our view, reinforce the sense that creativity occurs as a socio-technological process, understood by identifying and revealing the nature of various temporary people-object aggregates, and in particular the various relations between the functions of different subsystems. In our practice-based work in musical metacreative systems, we have both independently assumed a framework in which creativity occurs through these agency networks (Bown 2012, 2016).

A consequence of the evolution of digital systems is that the ecosystems of creativity are becoming more complex. A creative practitioner working with software may access multiple software tools, plugins, and assets created by others and distributed via the internet, involving increasingly digitally mediated collaborative practices such as remixes and mashups (Lessig 2008). Those assets might be made from other assets that have in turn been put together by the same process. The dynamic topologies of digital systems transform the network properties of creativity. As individual humans organise into fluid networks of distributed cognition, so

computational systems exhibit porous boundaries: machine-to-machine interactions such as web-services create new computer-agency hybrids with new affordances. Bringing these together into human-object agency networks amplifies the challenges of coordination through interaction design.

A final comment on this theme concerns attribution. In artistic domains, distributed creativity relates to issues of authorship, an ever-shifting and contended topic throughout the twentieth century and into the twenty-first. Sometimes we care very much who the author is; we care if a Renaissance painting is a forgery, even if it *looks* like the real thing. At other times authorship is less important – advertising jingles, for example. Metacreative systems may encounter different degrees of leeway in terms of how they can be used in place of human creative work, depending on the importance of authorship in any given context. For example, computer generated (authorless) townhouses and elevator music might be more permissible than computer generated war monuments or popular music hit singles.

## 5.4 Metacreative Design Paradigms

The features of metacreative systems described above are far from exhaustive, but they offer an insight into the issues surrounding metacreative interaction, pointing to design strategies for metacreation.

Based on the above, we can split interaction scenarios into three top-level categories, which cut across the distinction used by Blackwell (2002) and others between direct manipulation and programming. These top-level categories are:

1. Operation-based. The user operates the system (controlling parameters, interactively evolving, training, etc.). The system may have degrees of autonomy and computational intelligence but is presented as a production tool; the user uses the system to produce work.

Operation-based interfaces are what we commonly deal with in most HCI. The idea behind operation-based paradigm for metacreation is that we operate complex metacreative algorithms in much the same way we operate any other digital tool. A neural network, for example, could have an interface allowing us to set its parameters, train it, and operate it generatively.

2. Request-based. The user puts requests to the system, such as “generate this” or “evaluate this”. The system is a (likely remote) service that would resemble a search interface. The system would have a clear role that resembles the activities of a person, such as producing something or evaluating something, and would give a clear sense of autonomy, even if its output is predictable and directly manipulable. The interface can take the form of a text box, programming API or GUI and has a relative degree of open-endedness.

In creative domains, we can imagine smart systems that respond to such requests: “*make 100 deep house tracks with a walking jazz bassline and cut up excerpts from Martin Luther King’s speeches*”. A search engine-type interface parses the request, converts it into a machine specification, and then seeks generative services which perform various subtasks. Upon receiving the results, further correc-

tive instructions may be issued, where some dialogue may be used to better identify the criteria.

3. **Ambient.** The system operates ambiently and proactively. The user does not directly manipulate it nor submits requests. In one scenario, the system could be an evolving creative tool that is constantly updating based on improvements, possibly adapted to an individual user's needs. In another scenario, the system would take the form of a human assistant and make suggestions via language. In a third scenario, the system would perform forms of background analysis which is presented to the user as ambient information (e.g., the user begins to draw an image and the assistant tiles the background with related or complementary images, or indicates suggested completions of the image).

Content-aware fill functions in software such as Adobe Photoshop provide a glimpse of what is possible with ambient metacreation. With these tools, users can delete sections of an image and the software will fill them in by extrapolating from surrounding materials. For example, when removing an object from a photo the content-aware fill will fill in the background effectively, dealing with complex structures such as trees and clouds. In conjunction with learning and predictive algorithms, such ambient systems could have the capacity to not only complete actions for the user, but anticipate actions and take some appropriate complementary steps.

This analysis applies at a similar level to a categorisation offered by Lubart (2005), which is somewhat richer in metaphor—computer as nanny, as pen-pal, as coach and as colleague—but our list places the boundaries between categories differently. Current creative practice focuses largely on operation-based interaction, although we are familiar with request-based and ambient interaction scenarios in non-generative contexts. We also note that these paradigms should not be seen as absolute or clearly demarcated, and any allocation of a given system into any of these categories requires interpretation. For example, a user interacting with a request-based interface can still be seen as operating that system.

Existing design concerns related to operation-based interaction include the successful creation of modular systems that enable cumulative development and easy debugging, powerful visualisation of system dynamics, tools that support rapid exploration, systems that can be edited in multiple ways and interacted with from multiple viewpoints. The operation-based paradigm is more associated with older waves of HCI but has great potential to be rethought through the 3rd wave HCI focus on artefact ecologies and distributed creativity.

Existing design concerns related to request-based interaction include producing results that exhibit consistency, relevance, clarity, meaningfulness, and a form of manipulability. Generative systems should produce meaningful diversity. They should successfully handle ambiguity: rather than throw an error or return nothing, they should return something, along with ways to cater for misunderstanding. In the Google search engine this includes alternative suggestions and autocomplete to help narrow in on the 'correct' search term. Request-based systems should ideally produce outputs that can be manipulated in multiple ways, which can be supported by technology that enables manipulation of arbitrary material. For example, source-



separation algorithms in music promise to allow the manipulation of elements that have previously been mixed together. Above all, generative systems should produce ‘good’ results, the central challenge of metacreation.

Request-based systems offer great potential to scale, and can be rapidly and fluidly reconfigured (or self-reconfiguring). They can naturally occur in a distributed manner which truly breaks from a one-user-one-system perspective, to one of an ecosystem of many-to-many interactions between multiple machines and users. We see smart knowledge and computation engines/interfaces such as Wolfram Alpha<sup>5</sup> as indicative of what form such systems might take.

Existing design principles related to ambient interaction include not disrupting workflows (good timing, not creating distractions, relevance), and being sensitive to context, such as the current activity the user is engaging in. This in part involves a recognition of different phases of creative production, some of which are more open-ended and exploratory, whilst others are more objective and goal-driven. In the former case (where the overall style of a work might not yet have been determined), the generation of diverse ideas, or of thought-provoking reconfigurations of concepts may be appropriate, whereas in the latter case (where the style is locked down and the work needs to be completed), focused, goal directed support in the style of text autocomplete might be more valuable.

The request-based and ambient paradigms prompt consideration of how metacreative technologies could be interacted with effectively in these contexts. Such analysis must take into account the evolution of creative practice in distributed networks of users, including remixes and mashups, the interaction between multiple individuals and software components in the creation of works, tool configuration activities such as appropriation and personalisation, and open-ended, non-goal oriented exploration. Research must ultimately take place in the wild—looking at users in context, seeing how they use metacreative software to achieve real creative outcomes—and we point to our practice-based work in the context of digital music practice as a productive model. However, as Bødker also notes (2015), it is not always appropriate to ‘dump’ software on people and study them. Bridging studies and approaches to bootstrapping metacreative technologies needs to be considered.

## 5.5 Examples and Speculative Futures

In this section we discuss current examples of each of the three paradigms, and briefly speculate on how each paradigm might develop into the future. We draw our examples predominantly from our own area of expertise: the world of digital music.

---

<sup>5</sup><https://www.wolframalpha.com/>

## 5.5.1 *Operation-Based*

### 5.5.1.1 Examples

Game-like mobile music apps such as those made by Brian Eno and Peter Chilvers (2008–2016) provide an example of the simplest kinds of operation-based metacreative interfaces, emerging in the context of smartphones and tablets. In these works, a touch interface is used to trigger sounds or manipulate icons that represent musical fragments. The apps include underlying generative processes, typically of repetition or gradual variation, on top of which the user can collage a range of preset sound and musical fragments. Thus, the evolving soundscape results from automation and manipulation. Digital music artists now commonly release such apps as extensions of their music production. These are clear examples of the metageneration of form in an operation-based paradigm. More professionally oriented tools for metagenerating musical form using rules include Noatikl<sup>6</sup> and Patter.<sup>7</sup> In these cases the user is generally not assumed to have a technical knowledge of the algorithms involved, which are packaged to have a user-friendly interface which reveals the system behaviour.

Operation-based control of search systems include interactive genetic algorithms such as Picbreeder (Secretan et al. 2008), where the user can interactively breed new image forms. With these and with learning tools such as Jnana,<sup>8</sup> which learns musical style in an interactive context, interfaces are usually very simple and the direct control of parameters (such as the genetic structure that is evolved, or the statistical model that is learnt) is obscured (see further discussion in Bray and Bown 2017).

### 5.5.1.2 Futures

We see the complexification of digitally-mediated creative practice in phenomena such as remix and mashup culture. Copying and reinterpreting is a critical component in all creative practice, but in remix and mashup practices, it is not just the ideas that are copied, but the artefacts themselves: for example a sound file, or source code fragment (Lessig 2008). We are seeing this integration of practice through digital artefacts intensifying, with more instances of people working together on the same material: a Wiki entry or a DAW project. We believe the growth of creative coding infrastructure and communities will play an increasingly important role in this area, as code continues to be one of the best ways to interact with advanced metacreative algorithms, and the sharing of code allows for unlimited combinatoric creativity in hybrid, distributed, human-computer agency networks.

---

<sup>6</sup><https://intermorphic.com/noatikl/>

<sup>7</sup><http://playpatter.com/>

<sup>8</sup><https://ccrma.stanford.edu/~colinsul/projects/jnana/>

## 5.5.2 Request-Based

### 5.5.2.1 Examples

Request-based interactions with machines are increasingly prevalent. Common examples include internet search and recommendation systems. Common examples include internet search and recommendation systems and conversational user interfaces. This interaction design paradigm requires a medium of enquiry - typically text - familiar to the user, and a computational system with domain knowledge and an ability to service the request. Through developments in data mining and machine learning, we see these request-based systems being increasingly utilised in the creation of artistic works.

Style Machine<sup>9</sup> is an example of a music generation system that produces complete musical works in a given style or combination of styles from a range of electronic dance music options. Like Jnana and Patter, the system is built directly into a DAW, Ableton Live, so that users can have immediate low-level access to the generated content. Startups that are using machine learning engines to drive request-based interaction in music generation are emerging. Aiva and Jukedeck are two recent arrivals applying this type of interaction design to two different use-cases, one that is client focused and mediated by technicians and the other that is end-user focused, with a web interface.

Both of these examples show the hazy boundary between operation-based and request-based paradigms. Both have parametric controls that affect their behaviour (operation-based) but are also presented as black boxes that respond to requests (request-based).

### 5.5.2.2 Futures

Given recent developments, we foresee scenarios in which people will interact with request-based outputs from generative systems for their own pleasure: for example by using a service similar to Spotify to listen to original generative music. This might mean in turn that artists output generative systems rather than fixed media works. People may still be listening to *U2*, but via a *U2* generator that might be capable of generating an infinite number of *U2*-style songs.<sup>10</sup> Another alternative is that the generated artefacts may be authorless without need for the kudos of renowned creators.

---

<sup>9</sup><https://www.youtube.com/watch?v=Xw2I9B4yt7I>

<sup>10</sup>One way to continue the fruitful collaboration between *U2* and *iTunes*, started with “*Songs of Innocence*”?

### 5.5.3 *Ambient-Based*

#### 5.5.3.1 Examples

Like DarwinTunes discussed above, Draves' Electric Sheep (1999) is a well-known ambient generative system in the visual arts domain. It operates as a computer screen saver using a set of evolving algorithms to generate complex visual animations. The system runs 'ambiently' and autonomously during times of inactive computer use. Computers with Electric Sheep installed communicate with each other via the internet to share the load of creating new morphing abstract animations.

Another form of ambient interaction occurs where the system is simply conceived of as a contributor to a collaborative work. This is found to the greatest extent in the scenario of live improvisation. A common strategy here depends on corpus-based matching algorithms. A (typically short) musical query is submitted, and the system attempts to find a match, which could mean something similar to the query, or complementary to it either sequentially (call and response) or in parallel (harmonisation or counterpoint). For example, the interactive music system Frank (Plans and Morelli 2007) captures audio and uses gestural matching techniques (k-means clustering) in combination with genetic co-evolution. The system interacts with a human player by responding to audio queries (performed audio input) by evolving solutions and outputting the resulting audio in real time.

Interactive computer improvisation agents such as Frank, and our own work with systems such as CIM (Brown et al. 2013) and Zamyatin (Bown 2011) fit an ambient interaction paradigm in the sense that the primary form of interaction is to 'jam' with them, as you would with a fellow human musician. Musicians interacting with such systems do not need to operate or explicitly submit queries to the system. They just play.

#### 5.5.3.2 Futures

Along with user-led control, metacreation seems likely to involve a move away from the staple interaction paradigm of direct manipulation toward a more passive use of computers where the computer is actively suggesting, with users benefitting from these suggestions being ready to hand when required. Such change can be framed in the context of Harrison et al.'s (2007) survey of 3rd wave HCI trends, including shifts from generalised to situated knowledge, from information to interpretation, from clean to messy formalisms, and from scientific to unscientific strategies. To this list we would add a shift in the role of the computer from passive tool to active collaborator.

### ***5.5.4 Professional and Social Implications of Metacreative Systems***

Beyond generation, metacreative systems often employ processes of selection and judgement over generated (or input) materials, leading to situations where the computer takes over increasingly evaluative roles in a client-producer relationship (for example, suggesting why an aesthetic decision is valid). Automated evaluation can lead to increased efficiencies in traditionally labour intensive creative domains. But what other effects will such changes in technology have on creative professions and on society?

#### **5.5.4.1 In Professions**

Creative teams and workflows will inevitably be transformed. We foresee a continuation of the trend towards smaller groups involved in the development of works, which will take place more rapidly through the speed-up enabled by automation. We speculate that this may lead to greater domain generalism, with algorithms managing to capture domain-specific knowledge well, and successful operators of these algorithms consequently becoming empowered to work across domain boundaries. For example, in advertising, a single person may be involved in the production of an entire animated infomercial with characters, music, storyboard and focus-group tested messaging. Similarly, a single individual may create the sound-design for a movie in a day, mostly playing the role of a client to the software-as-producer. An architect may present a myriad of design options to a client, all tested and simulated with building analytics.

Metacreation offers new ways to engage in micro-branding, communicate via rich media, develop and modify memes on demand, and generate digital content on-the-fly for rich media environments such as AR, VR and game-based worlds. How and when to intervene in these processes and what modes of interaction to use when doing so remain design challenges best resolved with knowledge of particular situations and domains.

#### **5.5.4.2 In Society**

Metacreative systems that automate content generation will impact the contexts in which creative content production is conducted, not just in professional contexts, but in the everyday ‘prosumer’ creation of content. With such generative power, the casual creation of artefacts is set to become easier. Use cases might resemble Bødker’s observations of individual iPhone use innovation, where she points to “both the shared and explorative phases that are part of this development, and how the iPhone moved from a fancy telephone to, for example, a highly individual ‘poetry machine’, in the hands of one of the interviewees.” (Bødker 2015: 26).

The ability for people to rapidly create content without specialised training using metacreative systems will have an effect on many forms of human-computer interaction. A vast amount of computer use in casual social interaction takes the form of micro-content creation, primarily text, photos, video, audio, and types of structured data such as playlists. In the smartphone era, creation tools that enable rapid on-the-fly creation via minimal and intuitive interfaces have proven to hold special value. Inviting someone out, or wishing them happy birthday, could be augmented with a contextually relevant generative artwork whipped up on a smartphone. We can imagine how our list of interaction activities, themes, and overarching paradigms (operation-based, request-based, and ambient) might fit with the design of rapid metacreation tools for such scenarios. For example, stark emotional markers may be key in constraining a search of generated possibilities (e.g. ‘give me something happy’). Meanwhile, the novelty requirements of artefacts might be minimal in these scenarios, but a personal style signature developed by system creators or users as part of their digital identity may play an important role in making them feel that they are creating something that is not just generic.

Casual and social computer use also includes various forms of evaluation such as ratings (including of artists, aesthetic artefacts, or creative tools), along with a search for valued artefacts, which provides the big data required to build massive models of cultural phenomena that can be used to feed generative processes. We tend to think of metacreative technologies as fun and on the whole positive, or at least harmless. But recent global political events suggest that the combination of big data analytics of cultural trends, with the ability to generate rich emotionally salient media optimised for cultural impact could be a potent mix with dangerous uses. Metacreative researchers should be mindful of the deeper political implications when creative practices become a site for automation.

## 5.6 Conclusion

Metacreation technologies form part of the drive toward 3rd wave HCI not least because they privilege the agency of computing systems in interactive contexts. The consequences that come with this shift filter through many of the issues that this era of HCI is concerned with.

As in other areas of automation, metacreation will have an impact on how creative work is done. We are witnessing an ongoing complexification of socio-technological creative systems, and the evidence from our work suggests that metacreative systems will play a part in (and benefit from) the natural continuation of this complexification. Systems such as neural networks, generative processes, and web-services using big-data generation, depend increasingly on being situated in interactive contexts. Metacreative systems are often more complex than existing software in terms of their opacity and their potential for autonomy. We suggest that these new systems should coexist alongside existing tools to enrich the ecology of participants in human-object agency networks. Metacreative tools can be seen as

adding new dimensions to existing technological infrastructures and thus to human creativity.

Although the focus of our research is on making metacreative systems that work well, as this technology evolves we are also becoming increasingly aware of wider, potentially negative social implications, as alluded to in our comments above. Beyond the erosion of creative work, which we do not think is a significant threat, the affordances such tools might have in social control, through the creation of emotionally salient and attention grabbing content is now worthy of serious consideration. In light of the rapid changes the world is facing, such as recent trends in political manipulation, a field such as metacreation could lose its claims of innocence rapidly. We hope that research will be conducted top-to-bottom, and from a range of technical and analytical perspectives, to connect algorithmic and HCI innovation with these possible wider social implications.

## References

- Amabile T (1996) *Creativity in context*. Westview Press, Boulder
- Besold TR, Schorlemmer M, Smaill A (eds) (2015) *Computational creativity research: towards creative machines*. Atlantis Press, Berlin
- Blackwell AF (2002) First steps in programming: a rationale for attention investment models. In: *Proceedings of the IEEE 2002 Symposia on human centric computing languages and environments, 2002*. IEEE, Piscataway, pp 2–10
- Bloom P (2010) *How pleasure works: the new science of why we like what we like*. Random House, New York
- Bødker S (2006) When second wave HCI meets third wave challenges. In: *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*. ACM, New York, pp 1–8
- Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. *Interactions* 22(5):24–31
- Bown O (2011) Experiments in modular design for the creative composition of live algorithms. *Comput Music J* 35(3):73–85
- Bown O (2012) Generative and adaptive creativity: a unified approach to creativity in nature, humans and machines. In: *Computers and creativity*. Springer, Berlin/Heidelberg, pp 361–381
- Bown O (2014) Empirically grounding the evaluation of creative systems: incorporating interaction design. In: *Proceedings of the fifth international conference on computational creativity*, pp 112–119
- Bray L, Bown O (2017) How can we deal with the design principle of visibility in highly encapsulated computationally creative systems? In: *Proceedings of the international conference on computational creativity, June 19–23, Atlanta*
- Brown AR (2016) Understanding musical practices as agency networks. In: *Proceedings of the international conference on computational creativity*. [Online]. 2016 Paris, Association of Computational Creativity. Available from: [http://www.computationalcreativity.net/iccc2016/wp-content/uploads/2016/01/paper\\_21.pdf](http://www.computationalcreativity.net/iccc2016/wp-content/uploads/2016/01/paper_21.pdf)
- Brown AR, Sorensen A (2009) Interacting with generative music through live coding. *Contemp Music Rev* 28(1):17–29
- Brown A, Gifford T, Voltz B (2013) Controlling interactive music performance (CIM). In: *Proceedings of the fourth international conference on computational creativity, vol 221*
- Bruckner S, Gröller ME (2007) Style transfer functions for illustrative volume rendering. In: *Computer graphics forum (Vol. 26, No. 3, pp. 715–724)*. Blackwell Publishing Ltd, Cambridge

- Clark A, Chalmers DJ (1998) The extended mind. *Analysis* 58:10–23
- Cohen H (1995) The further exploits of AARON, painter. *Stanf Humanit Rev* 4(2):141–158
- Cope D, Mayer MJ (1996) Experiments in musical intelligence, vol 12. AR Editions, Madison
- Davis L (1991) Handbook of genetic algorithms. Van Nostrand Reinhold, New York
- Di Scipio A (2003) Sound is the interface: from interactive to ecosystemic signal processing. *Organised Sound* 8(03):269–277
- Draves S (1999) Electric sheep [Online]. Spotworks LLC. Available from: <https://electricsheep.org/>
- Eno B, Chilvers P (2008–2016) Bloom, trope, air, quarta & reflection. iPhone App Store, Opal Limited
- Ferguson J, Brown AR (2016) Fostering a post-digital avant-garde: research-led teaching of music technology. *Organised Sound* 21(2):127–137
- Galanter P (2003) What is generative art? Complexity theory as a context for art theory. In: GA2003–6th generative art conference. 10–12 December, 2003, Milan, Italy
- Gaver WW, Bowers J, Boucher A, Gellerson H et al (2004) The drift table: designing for ludic engagement. In: CHI'04 extended abstracts on Human factors in computing systems. ACM, New York, pp 885–900
- Ghedini F, Pachet F, Roy P (2016) Creating music and texts with flow machines. In: *Multidisciplinary contributions to the science of creative thinking*. Springer, Singapore, pp 325–343
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In *Alt. Chi*. Session at the SIGCHI conference on human factors in computing systems. San Jose, California, USA, pp 1–18
- Hudson R (2008) Frameworks for practical parametric design in architecture. In: *Proceedings of the 26th eCAADe conference*. Antwerp, Belgium, pp. 847–854
- Kantosalo A, Toivanen JM, Toivonen HT (2015) Interaction evaluation for human-computer co-creativity. In: *Proceedings of the sixth international conference on computational creativity*. June 29–July 2, Park City, Utah, USA
- Konak A, Coit DW, Smith AE (2006) Multi-objective optimization using genetic algorithms: a tutorial. *Reliab Eng Syst Saf* 91(9):992–1007
- Langston P (1989) Six techniques for algorithmic music composition. In: *15th international computer music conference*. ACMA, Columbus, pp 2–5
- Latour B (1995) Mixing humans and nonhuman together: the sociology of a door-closer. In: *Ecologies of Knowledge*. New York, State University of New York Press, pp 257–277
- Latour B (2011) What's the story? Organizing as a mode of existence. In: Passoth JH, Peuker B, Schillmeier M (eds) *Agency without actors*. Routledge, London
- Lehman J, Stanley KO (2008) Exploiting open-endedness to solve problems through the search for novelty. In: *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*. MIT Press, Cambridge, MA, pp 329–336
- Lessig L (2008) *Remix: making art and commerce thrive in the hybrid economy*. Penguin, New York
- Lubart T (2005) How can computers be partners in the creative process: classification and commentary on the special issue. *Int J of Hum Comput Stud* 63(4):365–369
- MacCallum B, Leroi A, Mauch M, Welburn S et al (2009) DarwinTunes [Online]. London, Imperial College and Queen Mary University. Available from: <http://darwintunes.org/>
- MacCallum RM, Mauch M, Burt A, Leroi AM (2012) Evolution of music by public choice. *Proc Natl Acad Sci* 109(30):12081–12086
- Malafouris L (2008) At the potter's wheel: an argument for material agency. In: Knappett C, Malafouris L (eds) *Material agency: towards a non-anthropocentric approach*. Springer, New York, pp 19–36
- Martin A, Jin CT, Bown O (2011) A toolkit for designing interactive musical agents. In: *Proceedings of the 23rd Australian computer-human interaction conference*. November 28–December 02, 2011, Canberra, Australia, pp 194–197



- McCormack J (2004) Aesthetic evolution of L-systems revisited. In: Workshops on applications of evolutionary computation. Springer, Berlin/Heidelberg, pp 477–488
- McCormack J, D’Inverno M (eds) (2012) Computers and creativity. Springer, Berlin
- Nielsen, Jakob. 1993. Usability Engineering. New York: Morgan Kaufmann, pp 115–164
- Norman D (1988) The psychology of everyday things. Basic books, New York
- Pask G (1976) Conversation theory. Applications in education and epistemology. Elsevier, Amsterdam/New York
- Perkins DN (1994) Creativity: beyond the darwinian paradigm. In Boden M (ed), Dimensions of creativity, Bradford Books, Cambridge, MA. pp 119–142
- Plans Casal D, Morelli D (2007) Remembering the future: an overview of co-evolution in musical improvisation. In: Proceedings of the international computer music conference. Copenhagen, International Computer Music Association, pp 200–206
- Rammert W (2008) Where the action is: distributed agency between humans, machines, and programs
- Salganik MJ, Watts DJ (2008) Leading the herd astray: an experimental study of self-fulfilling prophecies in an artificial cultural market. *Soc Psychol Q* 71(4):338–355
- Secretan J, Beato N, D Ambrosio DB, Rodriguez A, Campbell A, Stanley KO (2008) Picbreeder: evolving pictures collaboratively online. In: Proceedings of the SIGCHI conference on human factors in computing systems. April 27–May 02, 2013, Paris, France, pp 1759–1768
- Shneiderman B (1993) 1.1 Direct manipulation: a step beyond programming languages. *Sparks of innovation in human-computer interaction*, 17
- Simonton DK (1999) Origins of genius: Darwinian perspectives on creativity. Oxford University Press, New York
- Suchman LA (1987) Plans and situated actions: the problem of human-machine communication. Cambridge university press, New York
- Togelius J, Yannakakis GN, Stanley KO, Browne C (2011) Search-based procedural content generation: a taxonomy and survey. *IEEE Trans Comput Intell AI Games* 3(3):172–186
- Ventura D (2016) Mere generation: essential barometer or dated concept. In: Proceedings of the seventh international conference on computational creativity, ICCO. June 27–July 1, Paris, France
- Wallas G (1926) The art of thought. Harcourt Brace, New York
- Whitelaw M (2004) Metacreation: art and artificial life. The MIT Press, Cambridge, MA
- Winograd T, Flores F (1986) Understanding computers and cognition: a new foundation for design. Addison Wesley, Reading
- Woolley BG, Stanley KO (2014) A novel human-computer collaboration: combining novelty search with interactive evolution. In: Proceedings of the 2014 annual conference on genetic and evolutionary computation, 12–16 July, 2014, Vancouver, BC, Canada, pp 233–240

# Chapter 6

## 3D Printing Technologies: A Third Wave Perspective



Deborah Lupton

**Abstract** Three-dimensional (3D) printing is a novel digital technology that has gathered momentum and public recognition over the past few years. In this chapter, I examine the sociocultural and political dimensions of 3D printing technologies. I begin with an overview of the third wave human-computer interaction (HCI) approach to digital technologies and contributions made by social and cultural theory that are relevant to understanding the broader contexts of 3D printing technologies and how they are represented, discussed and experienced. This is followed by a discussion of the sociotechnical imaginaries that animate speculations about their possibilities and the agential capacities identified by research investigating the lived experiences of those who have tried using these technologies. The chapter ends with some brief reflections on future research directions.

### 6.1 Introduction

Three-dimensional (3D) printing is a novel digital technology that has gathered momentum and public recognition over the past few years. 3D printing is a process of fabricating objects using computer-aided design (CAD) software and hardware that responds to instructions from the software. It uses materials such as plastics, powders, ceramics, metals and organic matter such as tissue cells or foodstuffs propelled through jets that extrude substances in layers to form objects. Additive manufacturing is the technical term for this process. However, the term 3D printing has entered into common parlance, as part of attempts to capture public interest and domesticate these technologies as appropriate for mass use (Fordyce 2015). Additive manufacturing has been used in industry since the 1980s. It was initially employed to quickly make prototypes and fabricate customised parts on-demand (Birtchnell and Urry 2013a). It is only in the last few years, with the development of smaller and

---

D. Lupton (✉)  
University of Canberra, Bruce, Australia  
e-mail: [Deborah.Lupton@canberra.edu.au](mailto:Deborah.Lupton@canberra.edu.au)

less expensive hardware targeted at mainstream consumer use, that the potential for these technologies to be taken up in other social domains has been explored.

Various popular books, news stories, industry blogs and research publications have presented sometimes fervid visions of the possibilities of 3D printing technologies. Contentions that ‘3D printing could change the world’ (Campbell et al. 2011), lead to a ‘new industrial future’ (Birtchnell and Urry 2016), ‘disrupt’ industries and manufacturing (Bilpin 2014) or even ‘destroy the world’ (Armstrong 2014) have represented these technologies as revolutionary and ground-breaking. The World Economic Forum’s list of the top emerging technologies of 2016 included additive manufacturing (Meyerson 2015). A range of commentators, including entrepreneurs, researchers and journalists, have imagined a future in which the 3D printer will be commonly available in homes, offices, libraries and schools, similar to personal computers and inkjet printers (Birtchnell and Urry 2016; Lipson and Kurman 2013). These claims and promises for 3D printing technologies provide part of the social contexts for how users respond to these technologies. There are many other elements that also contribute to shaping responses to, understandings, and uses of novel technologies such as 3D printers. These include legal and policy environments, geographical location and infrastructure, embedded norms, moral meanings and assumptions, social identities and institutions, and membership of social groups.

It would appear that many of the possibilities and promises that have been expressed for 3D printing technologies have yet to eventuate, remaining the stuff of boosterish speculation or science fiction (Bosqué 2015; Hielscher and Smith 2014). Critics have begun to assert that some 3D printing technologies, such as home printers, have failed to capture consumers’ interest (Cranz 2016). A recent assessment of the state of 3D printing, seeking to look ‘beyond the hype’, concluded that its future lie in industrial applications, and that there is ‘still a long way to go’ to realise this potential (Holt 2017).

In this chapter, I examine the sociocultural and political dimensions of 3D printing technologies. I begin with an overview of the third wave human-computer interaction (HCI) approach to digital technologies and contributions made by social and cultural theory that are relevant to understanding the broader contexts of 3D printing technologies and how they are represented, discussed and experienced. This is followed by a discussion of the sociotechnical imaginaries that animate speculations about their possibilities and the agential capacities identified by research investigating the lived experiences of those who have tried using these technologies. The chapter ends with some brief reflections on future research directions.

## 6.2 Sociocultural Perspectives on Digital Technologies

Researchers in HCI have identified three waves or paradigms influencing their area (Bardzell 2010; Harrison et al. 2011, 2007; Sellen et al. 2009). The first wave focused on engineering factors, while the second wave took up social psychological perspectives, seeking to identify the cognitive processes underpinning the ways in

which people interact with digital technologies. The third, and most recent paradigm moves on from and broadens these approaches by emphasising sociocultural perspectives. In particular, the phenomenological dimensions of HCI have been emphasised: or the situated cultures and contexts in which people give meaning to and adopt practices in relation to digital technologies as part of their everyday lives, and in which developers, designers, researchers and coders generate ideas for objects and services. The third wave perspective attempts to site users beyond their computers to the world in which both reside (Sellen et al. 2009), including the embodied interaction of users with technologies as they move through time and space (Harrison et al. 2011).

A more critical, ethical and political approach to HCI research has also been advocated as part of the third paradigm. From this perspective, design is always a political process, shaping artefacts and services from particular viewpoints and assumptions while ignoring or closing off others. The underlying values, norms and assumptions underpinning user experience are identified, and the ways in which design operates to support the dominance of specific social groups and marginalise others are brought into sharp focus (Bardzell and Bardzell 2013; Bardzell et al. 2012; Khovanskaya et al. 2013; Khovanskaya et al. 2015). As feminist scholars have long argued, the design, development and use of digital technologies is profoundly gendered (Wajcman 2013). Feminist HCI researchers (Bardzell 2010; Fox and Rosner 2016; Rode 2011) and those interested in adversarial design (Bardzell et al. 2012; Björgvinsson et al. 2012; DiSalvo et al. 2012) direct their attention to the political dimensions of digital technologies. This may include ‘designing for provocation’ (Bardzell et al. 2012), or to inspire critical reflection and social change, rather than simply attempting to design a utilitarian solution to a problem. It also includes involving the participation of marginalised groups in design, prototyping and evaluation (Fox and Rosner 2016; Toupin 2014). HCI researchers have also recently called for a more ethical and empathic approach – sometimes referred to as ‘sensitive HCI’ (Waycott et al. 2015) – in both interacting with users and designing technologies that better suit their needs, particularly in relation to vulnerable, stigmatised or disadvantaged social groups such as children, the elderly, or people living with disabilities or mental and chronic health conditions (Thieme et al. 2014; Vines et al. 2013; Waycott et al. 2015).

Outside the realm of HCI studies, social and cultural theory and empirical research have progressively moved towards acknowledging the embodied, sensory and spatial dimensions of digital technology use. Scholars in science and technology studies employ the term ‘sociotechnical imaginaries’ to describe the defining discourses and practices that give meaning to novel technologies. Sociotechnical imaginaries are broad-scale visions that animate discourses and practices. Jasanoff (2015: 4) defines these imaginaries as ‘collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of, advances in science and technology’. Sociotechnical imaginaries work to shape

responses to novel technologies. These imaginaries relate to aspirations and normative aspects of social order. They build on broader sociocultural meanings that can relate to dominant ideas about how social and technological progress should be defined and develop, government and industry should operate, policy be developed and citizen/consumers respond. Multiple imaginaries can co-exist around a specific technology, either in tension with or productively responding to each other. Powerful voices and institutions, such as the media, government and industry, work to privilege some imaginaries over others (Jasanoff 2015).

While it is important to identify and understand these imaginaries when researching new technologies, a focus on the details of the lived experiences of users can provide key insights into the phenomenological dimensions of the enactments of technology use. In new materialism theory, humans and nonhuman actors such as digital technologies are viewed as mutually entangled, co-habiting and co-evolving (Barad 2003; Haraway 1995; Hayles 2012; Lupton 2015a). This approach understands human-computer interaction as involving the configuration of complex and heterogeneous assemblages of human and nonhuman actors that are subject to constant flux. From this perspective, the concept of human-computer interaction is expanded to that of the human-computer-data-space-place-time assemblage, a configuration that is constantly subject to change (Lupton 2016b). It also involves recognising that users' experiences of artefacts are always contingent and improvised rather than fixed, and that the artefacts themselves are always unstable and emergent in terms of their meanings, logics and affordances.

As humans and nonhumans come together in different dynamic configurations, agential capacities – the ability to affect and be affected – are generated (Barad 2003). This is a more-than-human approach to understanding human-computer interaction. It involves an awareness that designed objects themselves have agential capacities that contribute to the assemblages of humans and nonhumans of which they are part. Through the process of design and the incorporation of objects or services into everyday lives and mundane routines, they develop the capacity to affect and be affected. Critical social research can consequently be directed at 'how matter comes to matter', as Barad (2003) puts it. This may include identifying the agential capacities and affective power of these assemblages.

### 6.3 The Sociotechnical Imaginaries of 3D Printing Technologies

A diverse range of benefits and potential harms of 3D printing technologies have been envisaged. The current or suggested uses of 3D printing mainly fall into six main categories: (1) industrial; (2) creative; (3) marketing and leisure; (4) health and medical; (5) cultural heritage; and (6) educational.

### 6.3.1 *Industrial*

In the industry and manufacturing context, it has been asserted that 3D printing technologies offer a way of contributing to the reduction of environmental pollution and improving sustainability. Manufacturing can begin to take place in local areas or even in people's own homes, with less energy required for transporting materials, reduced vehicle emissions and less waste, as the fabricated objects can be more readily customised for consumers' individual needs (Birtchnell and Hoyle 2014; Birtchnell and Urry 2013b; Birtchnell et al. 2017). Some commentators have claimed that if these uses become widespread, 3D printing may lead to a third industrial revolution (Petrick and Simpson 2013), particularly in developing countries (Birtchnell and Hoyle 2014). Advocates of improving environmental sustainability and alleviating world hunger have also suggested that the technologies can be employed to create food products using alternative food sources such as insects, algae or food waste (Lupton 2017b). Objects already fabricated by 3D printing technologies include car and aircraft parts, customised clothing and shoes, jewellery and furniture (Birtchnell and Urry 2013a).

### 6.3.2 *Creative*

The creative and maker possibilities of consumer and domestic 3D printer technology use have also received a high degree of attention over the past decade or so. Employing domestic 3D printers to manufacture objects as part of craft, do-it-yourself, fan, hacker or hobbyist cultures is promoted as providing people with the opportunity to generate artefacts outside traditional institutional and commercial structures (Birtchnell et al. 2017; Fordyce et al. 2015; Ratto and Ree 2012; Walter-Herrmann and Buching 2014). Websites such as Thingiverse encourage the sharing of computer files for 3D printing enthusiasts. A range of personal 3D printers can be purchased for domestic use. Printers are also made available for public use in 'FabLabs' or alternatively 'Hackerspaces' or 'Makerspaces'. These spaces are heralded as promoting sociality, open knowledge sharing and creative coding and as potentially contributing to participatory design opportunities and the democratisation of invention (Carstensen 2014; Eisenberg 2013; Fleischmann et al. 2016; Hielscher and Smith 2014; Tanenbaum et al. 2013). Some FabLabs have been set up explicitly to challenge social norms and inequalities and provide better access to disadvantaged or marginalised groups to new technologies (Hielscher and Smith 2014). This approach has been referred to as a 'critical making practice' (Ratto et al. 2014).

### **6.3.3 *Marketing and Leisure***

3D printing technologies are also used for marketing and leisure purposes: for example, providing replica figurines of people that are used as mementos of events like graduations, sporting events and weddings. Fans can order figurines of their favourite characters with their own faces superimposed. One company offers opportunities for consumers to use home 3D printing to fabricate their own sex toys, while another can generate a printed foetus from digital 3D ultrasound files as a keepsake for expectant parents (Lupton 2016a). Chefs are experimenting with offering 3D printed food as part of novel gourmet cuisine. Various food companies have produced and offered for sale chocolate and other confectionary goods using the printers, including customised items that use digital files to produce edible portraits or wedding-cake figurines of the bride and groom (Lupton 2017b).

### **6.3.4 *Health and Medical***

In the realm of health and medicine, 3D printers are currently employed in medicine and dentistry to produce customised prosthetics, crowns and implants, and to generate bio-matter for purposes such as skin grafts. 3D printed anatomical replicas of whole or parts of people's bodies are used in medical and patient education and for planning surgery (Lupton 2015b). 3D food printers also are currently used in European nursing homes to make more appealing soft food for people with chewing and swallowing difficulties (Lupton 2017b).

### **6.3.5 *Cultural Heritage***

3D printing is employed by archaeologists, cultural heritage practitioners and museums to fabricate replacements for ancient and other cultural artefacts and create virtual experiences or mementos for visitors (Antle et al. 2017; Kidd 2015; Neumüller et al. 2014).

### **6.3.6 *Educational***

Teaching children and young people how to use 3D printing has been advocated as a way of encouraging their interest in science and technology subjects and to prepare them for future employment (Carstensen 2014). Some schools, libraries and universities have installed 3D printers as a new way of learning about and

experimenting with digital technologies and creating objects (Eisenberg 2013; Fordyce et al. 2015; Nemorin and Selwyn 2017).

## 6.4 Socio-Legal and Political Aspects

While most of the discourses framing the meaning of 3D printing have portrayed them in positive ways, researchers have also begun to discuss socio-legal aspects, including possible threats and harms. They have pointed out the ways in which 3D printing could be used illegally: in the manufacture of weapons such as guns, for example, or in the flouting of intellectual property law (Birtchnell and Urry 2013b; Daly 2016; Heemsbergen et al. 2016). The income streams and business models of many industries have been transformed by the entry of digital technologies, from journalism to the music industry. 3D printing technologies offer similar challenges in terms of how products are monetised in the face of the impetus for consumers to take greater control in creating and sharing goods, including digital content. Part of the discourses of creativity surrounding 3D printing discussions is the notion that the technologies allow people to engage in pursuits that are outside traditional forms of regulation and state control (Daly 2016; Heemsbergen et al. 2016). In some cases, there are also concerns about the safety of products manufactured using 3D printing and how this should be monitored and regulated, including goods such as medical devices, human or animal tissue and foodstuffs (Daly 2016; Tran 2015; Tran 2016).

The darker sociotechnical imaginaries for uses of 3D printing technologies were identified in Fordyce's (2015) analysis of the discussion of these technologies in online platforms for far-right, white supremacist, anti-feminist and men's rights politics. He noted how users portrayed their expectations of how these technologies could assist their repression of women and minority groups. These expectations include being able to print out weapons, customised sex dolls (with the physical characteristics of specific women that would not otherwise accept users as sexual partners) and for new forms of industry manufacturing that would benefit white men and marginalise others. There were also fears expressed on these forums that 3D printing could lead to white men losing jobs and thus their power base, while women may benefit from the changes to manufacturing to which these technologies were expected to lead.

## 6.5 The Lived Experience of Using 3D Printing Technologies

The research reviewed in the previous section demonstrates the ways in which analysing public discussion and portrayals of novel technologies like 3D printing can provide insights into their sociocultural and political dimensions. Thus far, only a small number of studies have investigated how people are taking up 3D printing technologies or their products as part of their everyday lives. As I noted earlier, there



is some speculation already that these technologies' potential for domestic use has been overstated. Thus far, the majority of consumers appear reluctant to embrace home 3D printing, as representatives of one of the major developers of these machines, MakerBot, conceded in late 2016 (Cranz 2016).

Studies that have investigated the contexts in which people are invited to use 3D printers outside the home have identified a range of issues that have hampered uptake and challenged assumptions about the creative potential of these technologies and how easy they are to use. The sociotechnical imaginary of FabLabs and maker spaces is that they allow democratic participation in innovative and creative use of digital technologies and the empowerment of users. The types of people who are attracted to these spaces, however, tend to be those who have a professional interest in 3D printing, or else are technologically adept and enjoy tinkering, and are therefore not alienated by high-tech environments (Carstensen 2014). Digital technologies tend to be associated with men – and particularly, middle-class, university-educated white men – who fit the stereotype of the 'computer nerd' or 'hacker'. A survey conducted in 2012 of people who use 3D printing technologies found that the majority were highly-educated men, with an average age was 35 years, and who identified with the maker movement (Moilanen and Vadén 2013).

Feminist researchers have pointed out that the kinds of spaces and promotional material used to attract people to experiment with the technology – and indeed, the very definition of what 'hacking' or 'tinkering' involve, and who should be doing it – needs to be gender-inclusive so that girls and women do not feel marginalised and alienated. They also contend that these spaces should be inclusive of non-white people and those who feel confronted by high-tech environments, including people from socioeconomically-disadvantaged backgrounds (Carstensen 2014; Fox and Rosner 2016; Nascimento 2014; Rosner and Fox 2016; Toupin 2014).

Users of 3D printing technologies are also grappling with the legal implications of open source sharing and tinkering. For example, a study of Australians using 3D printing at workshops and trade shows (Heemsbergen et al. 2016) found that they tended to hold a naïve understanding of their rights and responsibilities concerning managing the sharing and printing of designs when using the technologies. While the participants subscribed to the ideal of sharing design files with other 3D printer users and the right for users to modify others' designs, they were also reluctant to share designs because of the lack of regulation of online material and what was viewed as the dominance of major corporations in terms of possessing legal rights over the ordinary consumer. They were fearful that their own designs might be appropriated by large digital corporations.

Observations of people learning to use 3D printers in makerspaces or workshops have shown that the machines are not simply 'plug and play'. Various forms of specialised knowledge and expertise must be developed, and printers must be closely observed by the user to ensure successful fabrication (Bosqué 2015; Ratto and Ree 2012). Ratto and Ree (2012) observed how participants engaged in hands-on experimentation with these devices in a Canadian workshop and interviewed them later. Their participants, many of whom were industrial designers, highlighted the material properties of the process of printing and the fabricated objects and drew attention

to the possible environmental implications of technology should ‘desktop fabrication’ become popular. They wondered about how these technologies might affect their own profession, in terms of challenging their expertise, trivialising design, and transforming the manufacturing process. Ratto and Ree (2012: no page) conclude that 3D printing technologies are not just another tool or device. Instead, they constitute ‘a new form of material engagement that both productively and problematically recombines knowledge work, craft, and design in novel ways’.

In another project investigating people’s engagement with and learning about 3D printing, Hudson et al. (2016) researched how casual 3D printer users fared when trying to fabricate objects at public print centres. Their interviews found that these casual makers often struggled to successfully use these technologies. They were deeply dependent on the print centre operators to help them in the process, but these operators themselves were not always well equipped to do so. A project by Ludwig et al. (2015) involved observations and interviews with two German 3D printing user communities based in an HCI research lab and an artists’ workshop. They found that a playful fascination with the technologies initially motivated members of both communities, and being able to fabricate tangible objects that could be held and touched also contributed to their enjoyment. The complexity of the software and hardware was a factor that required high motivation on the part of the users to overcome so that they could continue to pursue their interest. The technologies were viewed as black boxes; the ways in which they worked and problems could be fixed were largely mysterious to the users, involving a high degree of tacit knowledge that was not readily shared with other users. The researchers conclude that these aspects of 3D printing present barriers to greater public use.

Little research has been conducted on the use of 3D printing in educational institutions. One significant exception is the study by Nemorin and Selwyn (Nemorin 2016; Nemorin and Selwyn 2017) in an Australian secondary school. Their ethnographic research identified the affective dimensions of this experience in the context of the school setting and existing pedagogical structures, including students’ lack of engagement and the sheer hard work, tedium and frustrations they encountered when attempting to make the technology work as desired. Nemorin (2016) refers to this as ‘the affective labour of failing’: a type of work in which students and teachers are required to engage in the context of the structural constraints of the classroom and the requirement to meet the demands of the curriculum.

Based on their study’s findings, Nemorin and Selwyn (2017) question whether introducing these technologies in the school setting can readily achieve the socio-technical imaginaries attributed to them, such as promoting science and technology education and entrepreneurial and maker cultures among young people. They highlight the importance of identifying and acknowledging the complexities of the social, spatial and temporal constraints of teachers’ attempts to involve students in learning how to use these technologies. As was found in the studies of adults using 3D printers in maker spaces referred to earlier, the students encountered many technical challenges in using the software for design and then ensuring that their objects were printed out to an acceptable standard. Existing pedagogical assumptions structured the ways in which the technology was introduced into the curriculum and

space of the classroom which constrained efforts to encourage creativity and initiative from the students along the lines suggested by maker culture discourses.

Another qualitative field study reported on the introduction of 3D design and printing with children in a very different geographical and sociocultural context: a Palestinian refugee camp (Stickel et al. 2015). The researchers found that the children appreciated the playful, collaborative, self-expressive and creative aspects of experimenting with these technologies, as well as the opportunity to customise and turn their designs into physical objects. They were able to design and make objects for themselves that they could incorporate into other play. The difficulties they experienced related to socio-technical limitations such as poor usability of the interface and the printers. The children in this study were able to engage in digital technologies and creative collaboration in a context in which such opportunities were limited. This kind of research positions 3D printing technologies as contributing to humanitarian efforts to educate and engage children living in disadvantaged and deprived conditions. These social contexts work both to promote the children's interest and excitement in trying these new technologies, but also impose significant limitations, such as lack of resources to allow the children to continue or to build on their knowledge.

The reception of consumers to 3D printed objects has also been little explored thus far. This was an issue my colleague Bethaney Turner and I set out to research. We conducted an online discussion group project involving a group of Australians responding to questions about 3D printed food products (Lupton and Turner 2017, 2018). We found that consumers experienced difficulties in understanding how these food products were fabricated. They were not very familiar with how 3D printing in general worked, and none had had personal experience of using a 3D printer. 3D printing was associated with plastic, inedible objects. The study participants struggled to comprehend how such an unfamiliar process as digital fabrication could generate edible, appealing and nutritious food. When we provided them examples of 3D printed food products and asked the participants for their responses to them, they were more accepting of food products that they knew to already be highly processed, such as chocolate, confectionery, pizza and pasta. The idea of consuming a printed food made from insects, food waste or algae, however, was greeted largely with disgust. The apparent texture and appearance of 3D printed food was also off-putting for many people. As this study showed, the reception of 3D printed products such as food by consumers involves a range of acculturated norms and affective responses that go beyond ideas and understandings concerning the technologies themselves, but also include such factors as the appearance and constituents of the products.

Another focus of social research into 3D printing directs attention to the implications for people's engagements with their personal data. These technologies offer a way to materialise data into tangible forms. Many digital technologies collect information about the people who use them. In most cases, these personal data can be viewed in the form of two-dimensional materialisations such as graphs or metrics displayed on a digital device screen. If 3D printing technologies are used, however, personal data can be materialised into physical objects that can be touched, held,

displayed, and even smelt, listened to or eaten (Barrass 2016; Khot et al. 2014, 2015; Lupton 2017a). Designers and HCI researchers have experimented with using 3D printing to transform personal data into physical objects – or what are often termed ‘data physicalisations’ (Jansen et al. 2015). They have observed the extent to which the opportunity to handle, touch and display these objects help people to make sense of data in more multifaceted ways and to engage more emotionally with their data.

One research team adopted this approach in a project investigating whether printing out objects from physical exercise data tracked by digital devices would encourage people to be more active (Khot et al. 2014). These researchers have also experimented with printing chocolate artefacts as rewards for physical exercise, so that people can smell, taste and consume as well as view and touch these materialisations of their personal data and feel rewarded for their efforts (Khot et al. 2015).

Adopting a more political perspective on 3D data materialisations, Nissen and Bowers (2015) have experimented with encouraging people to engage in ‘data making’ by producing ‘data-things’ from their own data using 3D printing technologies. Nissen and Bowers refer to such activities as ‘participatory data translation’. This approach to personal data materialisation is intended to promote communal rather than individualistic modes of generating and materialising data that invite reflection on how this information is created and who has the opportunity to use it. This is a critical making approach that also resonates with contemporary critical data studies perspectives that seek to understand how people make sense of and engage with their personal data, and to what extent they are aware of how third parties might exploit their data (Lupton 2017a; Pink et al. 2017; Stark 2016).

In another critical making project, Devendorf and colleagues (Devendorf et al. 2016, 2015) devised a speculative portable digital fabrication system. This project sought to highlight the values and assumptions about making, creativity and human and nonhuman agency that are embedded in the design of 3D printing technologies. The participants were invited to become part of this new digital fabrication system, and in so doing, their role in the fabrication process was reimagined. Instead of a conventional digital 3D printer being used to make objects, the participants used their hands to lay down the materials, following the type of software instructions that typically direct 3D printing. The researchers sought to draw attention to the ways in which conventional 3D printing can close off artistic and creative making by removing hand work from the process of forming objects and deferring agency to the printing apparatus instead. In conventional 3D printing, while the finished object can be handled and touched by makers, the process of shaping its form is directed by the machine rather than the human body. By replacing the machine with human hands, the sensory dimensions of digitised fabrication are brought back in, and the interactions of humans and nonhumans rendered more open-ended and improvisatory. The agential capacities of these human-nonhuman assemblages are reconfigured, provoking reflection on the limitations of conventional 3D printing systems.

## 6.6 Future Research Directions

As I have shown in this chapter, the introduction of novel technologies such as 3D printing takes place in social, cultural, and political contexts in which publics make sense of and respond to these technologies drawing on longstanding but often tacit acculturated beliefs and norms. Some of these beliefs and norms may influence acceptance of these novel technologies and the objects they fabricate, while others work against it, even inspiring rejection or resistance. For those researchers wishing to adopt a third wave HCI perspective, it is important to recognise how novel digital technologies like 3D printers and their fabricated objects are incorporated into people's everyday lives across a range of domains, places and spaces.

A continued focus on the ways in which the popular media, industry blogs, online forums and academic researchers play roles in establishing sociotechnical imaginaries is important. Many aspects such as 3D printing technologies' social, cultural, economic and political impacts on education, the workplace, consumption and leisure cultures, including tourism, museums, exercise, gaming, fandom and food remain to be investigated. These technologies' legal and ethical implications, their contribution to social relations and identities, notions of embodiment, sensory experiences and knowledges, and to people's understandings of personal data are some intriguing possibilities for further investigation. These elements all contribute to the more-than-human worlds in which 3D printing technologies are designed, developed, marketed, sold and incorporated into everyday lives.

## References

- Antlej K, Leen R, Russo A (2017) 3D food printing in museum makerspaces: creative reinterpretation of heritage. *KnE Eng* 2:1–7
- Armstrong R (2014) 3D printing will destroy the world unless it tackles the issue of materiality. *The Architectural Review*. <http://www.architectural-review.com/home/products/3d-printing-will-destroy-the-world/8658346.article>. Accessed 11 Oct 2014
- Barad K (2003) Posthumanist performativity: toward an understanding of how matter comes to matter. *Signs* 28:801–831
- Bardzell S (2010) Feminist HCI: taking stock and outlining an agenda for design. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10). ACM, Atlanta, pp 1301–1310
- Bardzell J, Bardzell S (2013) What is critical about critical design? In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). ACM, Paris, pp 3297–3306
- Bardzell S, Bardzell J, Forlizzi J, Zimmerman J, Antanitis J (2012) Critical design and critical theory: the challenge of designing for provocation. In: Proceedings of the Designing Interactive Systems Conference (DIS '10). ACM, Newcastle, UK, pp 288–297
- Barras S (2016) Diagnosing blood pressure with acoustic Sonification singing bowls. *Int J Hum Comput Stud* 85:68–71
- Bilpin L (2014) 10 industries 3D printing will disrupt or decimate. *TechRepublic*. <http://www.techrepublic.com/article/10-industries-3d-printing-will-disrupt-or-decimate/>. Accessed 11 Oct 2014

- Birchneil T, Hoyle W (2014) 3D printing for development in the global south: the 3D/4D challenge. Palgrave Macmillan, Houndmills
- Birchneil T, Urry J (2013a) 3D, SF and the future. *Futures* 50:25–34
- Birchneil T, Urry J (2013b) Fabricating futures and the movement of objects. *Mobilities* 8:388–405
- Birchneil T, Urry J (2016) A new industrial future?: 3D printing and the reconfiguring of production, distribution, and consumption. Routledge, London
- Birchneil T, Urry J, Westgate J (2017) Design mobilities via 3D printing. In: Spinney J, Reimer S, Pinch P (eds) *Mobilising design*. Routledge, London, pp 76–86
- Björgvinsson E, Ehn P, Hillgren P-A (2012) Agonistic participatory design: working with marginalised social movements. *CoDesign* 8:127–144
- Bosqué C (2015) What are you printing? Ambivalent emancipation by 3D printing. *Rapid Prototyp J* 21:572–581
- Campbell T, Williams C, Ivanova O, Garrett B (2011) *Could 3D printing change the world? Technologies, potential, and implications of additive manufacturing*. Atlantic Council, Washington, DC
- Carstensen T (2014) Gendered fab labs. In: Walter-Herrmann J, Buching C (eds) *Fab Lab: of machines, makers and inventors*. Transcript Verlag, Bielefeld, pp 53–64
- Cranz A (2016) Home 3D printing ‘just not there yet’ admits MakerBot. *Gizmodo*. <http://www.gizmodo.com.au/2016/09/home-3d-printing-just-not-there-yet-admits-makerbot/>. Accessed 28 Oct 2016
- Daly A (2016) *Socio-legal aspects of the 3D printing revolution*. Palgrave Macmillan, Houndmills
- Devendorf L, Ryokai K (2015) Being the machine: reconfiguring agency and control in hybrid fabrication. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI ‘15)*. ACM, Seoul, pp 2477–2486
- Devendorf L, De Kosnik A, Mattingly K, Ryokai K (2016) Probing the potential of post-anthropocentric 3D printing. In: *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS ‘16)*. ACM, Brisbane, pp 170–181
- DiSalvo C, Louw M, Holstius D, Nourbakhsh I, Akin A (2012) Toward a public rhetoric through participatory design: critical engagements and creative expression in the neighborhood networks project. *Des Issues* 28:48–61
- Eisenberg M (2013) 3D printing for children: what to build next? *Int J Child Comput Interact* 1:7–13
- Fleischmann K, Hielscher S, Merritt T (2016) Making things in Fab Labs: a case study on sustainability and co-creation. *Digit Creat* 27:113–131
- Fordyce R (2015) Manufacturing imaginaries: Neo-Nazis, men’s rights activists, and 3D printing. *J Peer Prod* 6. <http://peerproduction.net/issues/issue-6-disruption-and-the-law/peer-reviewed-articles/manufacturing-imaginaries-neo-nazis-mens-rights-activists-and-3d-printing/>. Accessed 6 Aug 2017
- Fordyce R, Heemsbergen L, Mignone P, Nansen B (2015) 3D printing and university makerspaces: surveying countercultural communities in institutional settings. *Digit Cult Educ* 7:192–205
- Fox S, Rosner DK (2016) Inversions of design: examining the limits of human-centered perspectives in a feminist design workshop image. *J Peer Prod* 8. <http://peerproduction.net/issues/issue-8-feminism-and-unhacking-2/peer-reviewed-papers/inquiry-through-inversion-collisions-of-feminism-and-design-in-two-workshops/>. Accessed 6 Aug 2017
- Haraway D (1995) Foreword: cyborgs and symbionts: living together in the new world order. In: Gray CH (ed) *The cyborg handbook*. Routledge, New York, pp xi–xx
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ‘07)*. ACM, San Jose, CA, pp 1–18
- Harrison S, Sengers P, Tatar D (2011) Making epistemological trouble: third-paradigm HCI as successor science. *Interact Comput* 23:385–392
- Hayles NK (2012) *How we think: digital media and contemporary technogenesis*. University of Chicago Press, Chicago

- Heemsbergen L, Fordyce R, Nansen B, Apperley T, Arnold M, Birtchnell T (2016) Social practices of 3D printing: decentralising control and reconfiguring regulation. *Aust J Telecommun Digit Econ* 4:110–125
- Hielscher S, Smith AG (2014) Community-based digital fabrication workshops: a review of the research literature. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2742121](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2742121). Accessed 4 Aug 2017
- Holt V (2017) Beyond the hype: what's next for industrial 3D printing. [http://www.huffingtonpost.com/vicki-holt/beyond-the-hype-whats-nex\\_b\\_14751106.html](http://www.huffingtonpost.com/vicki-holt/beyond-the-hype-whats-nex_b_14751106.html). Accessed 5 Aug 2017
- Hudson N, Alcock C, Chilana PK (2016) Understanding newcomers to 3D printing: motivations, workflows, and barriers of casual makers. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '16). ACM, San Jose, CA, pp 384–396
- Jansen Y et al (2015) Opportunities and challenges for data physicalization. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15). ACM, Seoul, pp 3227–3236
- Jasanoff S (2015) Future imperfect: science, technology, and the imaginations of modernity. In: Jasanoff S, Kim S-H (eds) *Dreamscapes of modernity: sociotechnical imaginaries and the fabrication of power*. University of Chicago Press, Chicago, pp 1–33
- Khot R, Hjorth L, Mueller FF (2014) Understanding physical activity through 3D printed material artifacts. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). ACM, Toronto, pp 3835–3844
- Khot R, Pennings R, Mueller FF (2015) EdiPulse: supporting physical activity with chocolate printed messages. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15). ACM, Seoul, pp 1391–1396
- Khovanskaya V, Baumer EP, Cosley D, Volda S, Gay G (2013) Everybody knows what you're doing: a critical design approach to personal informatics. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13), Paris, ACM, pp 3403–3412
- Khovanskaya V, Baumer EPS, Sengers P (2015) Double binds and double blinds: evaluation tactics in critically oriented HCI. *Aarhus Ser Hum Cent Comput* 1:12
- Kidd J (2015) Museums are using virtual reality to preserve the past – before it's too late. *The Conversation*. <https://theconversation.com/museums-are-using-virtual-reality-to-preserve-the-past-before-its-too-late-44600>. Accessed 21 Oct 2016
- Lipson H, Kurman M (2013) *Fabricated: the new world of 3D printing*. Wiley, Indianapolis
- Ludwig T, Stickel O, Boden A, Pipek V, Wulf V (2015) Appropriating digital fabrication technologies – a comparative study of two 3D printing communities. Paper presented at the iConference 2015 Newport Beach, CA
- Lupton D (2015a) *Digital sociology*. Routledge, London
- Lupton D (2015b) Fabricated data bodies: reflections on 3D printed digital body objects in medical and health domains. *Soc Theory Health* 13:99–115
- Lupton D (2016a) 3D printed self-replicas: personal digital data made solid. In: Carnicelli S, McGillivray D, McPherson G (eds) *Digital leisure cultures: critical perspectives*. Routledge, London, pp 26–38
- Lupton D (2016b) Digital companion species and eating data: implications for theorising digital data–human assemblages. *Big Data Soc* 3. <http://journals.sagepub.com/doi/abs/10.1177/2053951715619947>. Accessed 7 Jan 2016
- Lupton D (2017a) Feeling your data: touch and making sense of personal digital data. *New Media Soc* 19:1599–1615
- Lupton D (2017b) 'Download to delicious': promissory themes and sociotechnical imaginaries in coverage of 3D printed food in online news sources. *Futures* 93:44–53
- Lupton D, Turner B (2017) 'Both fascinating and disturbing': consumer responses to 3D food printing and implications for food activism. In: Schneider T, Eli K, Dolan C, Ulijaszek S (eds) *Digital food activism*. Routledge, London, pp. 169–185.
- Lupton D, Turner B (2018) 'I can't get past the fact that it is printed': consumer attitudes to 3D printed food. *Food Cult Soc* (in press)

- Meyerson B (2015) Top 10 emerging technologies of 2015. World Econ Forum. <https://www.weforum.org/agenda/2015/03/top-10-emerging-technologies-of-2015-2/#additive-manufacturing>. Accessed 21 Oct 2016
- Moilanen J, Vadén T (2013) 3D printing community and emerging practices of peer production. First Monday 18 <http://journals.uic.edu/ojs/index.php/fm/article/view/4271/3738>. Accessed 7 Aug 2017
- Nascimento S (2014) Critical notions of technology and the promises of empowerment in shared machine shops. *J Peer Prod* 5. <http://peerproduction.net/issues/issue-5-shared-machine-shops/editorial-section/critical-notions-of-technology-and-the-promises-of-empowerment-in-shared-machine-shops/>. Accessed 31 Oct 2014
- Nemorin S (2016) The frustrations of digital fabrication: an auto/ethnographic exploration of '3D Making' in school. *Int J Technol Des Educ* (online ahead of print)
- Nemorin S, Selwyn N (2017) Making the best of it? Exploring the realities of 3D printing in school. *Res Pap Educ* 32:578–595
- Neumüller M, Reichinger A, Rist F, Kern C (2014) 3D printing for cultural heritage: preservation, accessibility, research and education. In: Ioannides M, Quak E (eds) *3D research challenges in cultural heritage: a roadmap in digital heritage preservation*. Springer, Berlin, pp 119–134
- Nissen B, Bowers J (2015) Data-Things: digital fabrication situated within participatory data translation activities. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, Seoul, ACM, pp 2467–2476
- Patrick IJ, Simpson TW (2013) 3D printing disrupts manufacturing: how economies of one create new rules of competition. *Res Technol Manag* 56:12–16
- Pink S, Sumartojo S, Lupton D, LaBond CH (2017) Mundane data: the routines, contingencies and accomplishments of digital living. *Big Data Soc* 4. <https://doi.org/10.1177/2053951717700924>. Accessed 28 Mar 2017
- Ratto M, Ree R (2012) Materializing information: 3D printing and social change. First Monday 17 <http://firstmonday.org/ojs/index.php/fm/article/viewArticle/3968>. Accessed 8 Aug 2015
- Ratto M, Wylie SA, Jalbert K (2014) Introduction to the special forum on critical making as research program. *Inf Soc* 30:85–95
- Rode JA (2011) A theoretical agenda for feminist HCI. *Interact Comput* 23:393–400
- Rosner DK, Fox SE (2016) Legacies of craft and the centrality of failure in a mother-operated hackerspace. *New Media Soc* 18:558–580
- Sellen A, Rogers Y, Harper R, Rodden T (2009) Reflecting human values in the digital age. *Commun ACM* 52:58–66
- Stark L (2016) The emotional context of information privacy. *Inf Soc* 32:14–27
- Stickel O, Hornung D, Aal K, Rohde M, Wulf V (2015) 3D printing with marginalized children – an exploration in a Palestinian refugee camp. In: Boulus-Rødje N, Ellingsen G, Bratteteig T, Aanestad M, Bjørn P (eds) *Proceedings of the 14th European conference on computer supported cooperative work (ECSCW '15)*, Oslo, Springer, pp 83–102
- Tanenbaum JG, Williams AM, Desjardins A, Tanenbaum K (2013) Democratizing technology: pleasure, utility and expressiveness in DIY and maker practice. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, Paris, ACM, pp 2603–2612
- Thieme A et al (2014) Enabling empathy in health and care: Design methods and challenges. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)* Toronto, ACM, pp 139–142
- Toupin S (2014) Feminist hackerspaces: the synthesis of feminist and hacker cultures. *J Peer Prod* 5. <http://peerproduction.net/issues/issue-5-shared-machine-shops/peer-reviewed-articles/feminist-hackerspaces-the-synthesis-of-feminist-and-hacker-cultures/>. Accessed 7 Aug 2017
- Tran JL (2015) To bioprint or not to bioprint North Carolina. *J Law Technol* 17:123–178
- Tran JL (2016) 3D-printed food. *Minn J Law Sci Technol* 17:855–880
- Vines J, Clarke R, Wright P, McCarthy J, Olivier P (2013) Configuring participation: on how we involve people in design. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, Paris, ACM, pp 429–438



- Wajcman J (2013) *TechnoFeminism*. Wiley, Hoboken
- Walter-Herrmann J, Buching C (2014) Notes on FabLabs. In: Walter-Herrmann J, Buching C (eds) *FabLab: of Machines, Makers and Inventors*. Transcript Verlag, Bielefeld, pp 9–23
- Waycott J, Wadley G, Schutt S, Stabolidis A, Lederman R (2015) The challenge of technology research in sensitive settings: case studies in 'sensitive HCI'. In: *Proceedings of the Australian Conference on Human-Computer Interaction (OZCHI '15)*, Melbourne, ACM, pp 240–249

# Chapter 7

## Denaturalizing 3D Printing's Value Claims



Gabby Resch, Daniel R. Southwick, and Matt Ratto

**Abstract** This chapter examines how 3D printing has been framed as a *liberatory technology* that confers agency to users on the one hand, and an *automated system that de-centers the user* on the other. These entangled visions, we argue, can be traced to values that are threaded into 3D printing's DNA. By historically situating the social context of 3D printing, tracing its roots to the CAD/CAM revolution of the 1950s and 1960s, we denaturalize assumptions about the technology's users, its modes of interaction, and its societal impact, offering third wave HCI new insights for broadening how it considers context and values.

### 7.1 Introduction

Third wave HCI has become a convenient handle for design of human-computer interactions that claims to be both *contextually grounded* and *values-based*. What is at stake in designing technologies informed by these two concerns if they are (a) taken for granted, and (b) unexamined with respect to where and how they arise? By developing greater sensitivity to the social context of new technologies, third wave HCI practitioners and theorists have shown how innovations might invite participation by larger and more diverse audiences; be deployed in a wider array of situations; and be used to actively intervene in problematic aspects of society. This grounding – or “situatedness” – has, if anything, strengthened the instrumental foci of HCI, despite zero-sum fears that it would undermine them. But beyond “contexts of use” that merely differentiate between collaborative work environments and the “general context of culture and human being” that constitutes non-work (Bødker 2006), what do we mean, exactly, by *context*? And *whose/which values* are new designs being based on? In this chapter, we argue that attending to social context and, in effect, surfacing the assumptions, biases, and values that are naturalized in the development of new technologies, also requires illuminating the historical

---

G. Resch (✉) · D. R. Southwick · M. Ratto  
Faculty of Information, University of Toronto, Toronto, Canada  
e-mail: [gabby.resch@utoronto.ca](mailto:gabby.resch@utoronto.ca)

context from which they emerge. This kind of sociocultural work, if effective, provides a vehicle for denaturalizing assumptions, biases, and values once they have been surfaced. In the sections that follow, we unsettle a particularly insidious notion that 3D printing, which has been underrepresented in the HCI literature to date, is an emergent technology with an entirely *de novo* set of use metaphors and interaction possibilities, divorced somehow of historical context and value-ladenness. We demonstrate how historically contextualizing 3D printing's emergence calls into relief a number of crucial assumptions about its users, its modes of interaction, and its societal impact – assumptions that are ported almost wholesale from the historical design space of computer-aided design and manufacturing (CAD/CAM).

We will use two studies from our lab to ground this work. The first is a long-term cultural-historical study of CAD/CAM and 3D printing in various contexts. The second involves the design of a software and hardware system that facilitates the production of 3D-printed prosthetic and orthotic devices in clinical settings. We claim, in this chapter, that tracing the antecedents of novel technologies is an important step toward situating context and, ultimately, fostering sustainable design practices that also improve instrumental outcomes. Furthermore, doing this work facilitates opportunities for reflection on more general assumptions about how humans and computers should and do interact – assumptions that are often baked into technical infrastructures. We demonstrate that, while denaturalization is an interesting process in its own right, it can serve to produce better instrumental and ethical outcomes when incorporated in practices of design that encourage embodied, situated, materially-aware perspectives.

## 7.2 Surfacing the Values of Desktop 3D Printing

In the sections that follow, we illuminate how 3D printing has been framed as a *liberatory technology* that confers agency to users on the one hand, and an *automated system that de-centers the user* on the other. These entangled visions, we argue, can be traced to values that are threaded into the 3D printing community's DNA. To make this argument, we leverage our lab's long-standing engagement with desktop 3D printing, which dates back to its initial appearance as a viable public technology. Desktop 3D printing emerged at the end of the first decade of the 2000s.<sup>1</sup> The MakerBot *Cupcake CNC* and *Thing-O-Matic* devices, in particular, heralded a "new dawn" of desktop manufacturing according to the techno-utopian fanfare and marketing propaganda issued at the time.<sup>2</sup> Unlike a number of open DIY

---

<sup>1</sup>Fused deposition modeling (FDM) printing (and related forms of additive manufacturing), the most common process for hobbyist-grade desktop 3D printing, traces back to the 1980s. The first commercial FDM machine was released by Stratasys in 1992. Various desktop printer kits, including do-it-yourself (DIY) projects like the open source RepRap and the MakerBot Cupcake CNC, became available between 2005 and 2009. The technology's origins, however, as we will see, are rooted in CAD/CAM technology.

<sup>2</sup>For a great example of this, see MakerBot's description of its automated build platform: <https://www.thingiverse.com/thing:4056>. This techno-utopian narrative overshadowed more cautious

technologies that came on to the market within a few years of each other (e.g. Arduino), MakerBot's devices signaled a responsibility for users to create new digital and material worlds,<sup>3</sup> a responsibility enshrined within a liberatory ideal. The narrative of 3D printing as a *liberatory technology* (Bookchin 1968) served to foster a sense of responsibility among desktop 3D printing's early adopters for developing and populating a peer-to-peer ecosystem, initiating a new kind of decentralized political formation. This peer-to-peer formation would find its home within a novel digital commons that would instantiate the emancipatory ideals espoused by a growing number of players in the industry.<sup>4</sup> An ongoing systems design project undertaken by a community of novices and experts alike, MakerBot's ecosystem fostered active participation among users responsible for building up Thingiverse (MakerBot's online repository for digital content); multivocal dialogue (on internet forums) between MakerBot engineers and a community of DIY guinea pigs; and direct engagement between the emergent Maker community and the new stars of 3D printing's startup culture, particularly Bre Pettis, MakerBot's camera-friendly co-founder. (There were precedents for all of this, as we shall see.)

MakerBot's early mandate was to bring "the future to your desktop" by producing the easiest to use and most affordable printers for a wide user base.<sup>5</sup> "Start the 3D printing revolution from your living room" MakerBot implored its customers. "If you think it, you can make it." Its kits were "designed to be hacked." Its operational devices would "make things for you" – "your own little factory" – a true blurring of the work-life technology distinction that has become a hallmark concern of third wave HCI.<sup>6</sup> Pettis carefully articulated a vision for a "personal manufacturing revolution" that the company's devices would spark.<sup>7</sup> Its users would be recognized as digital pioneers.<sup>8</sup> When MakerBot was purchased by 3D printing giant Stratasys in 2013, however, the liberatory ideals that underpinned its erstwhile open source project came grinding to a halt. MakerBot's community largely abandoned it, spurring the growth of new entries in the 3D printing market such as LulzBot and

---

arguments about piracy, remix, and the "napsterization of three-dimensional things." See, for example, <http://www.nytimes.com/2007/04/05/business/05scan.html>

<sup>3</sup> Historian Langdon Winner (1983) writes of the development of new technologies: "New worlds are being made. There is nothing 'secondary' about this phenomenon. It is, in fact, the most important accomplishment of any new technology."

<sup>4</sup> See <http://www.bloomberg.com/news/2012-06-14/3-d-copying-makes-michelangelos-of-the-masses.html>

<sup>5</sup> <https://www.makerbot.com/media-center/2012/05/08/makerbots-new-digs>

<sup>6</sup> Each of these statements is difficult to attribute, as they have variously turned up in marketing pieces, video titles, campaigns, etc. Refer to the following: <http://www.hypercastle.com/blog/diy-3d-printer/>; <https://www.shapeways.com/blog/archives/1762-3dea-a-holiday-3d-printing-pop-up-store.html>; <http://web.archive.org/web/20100206135626/http://makerbot.com/>

<sup>7</sup> <http://www.businessinsider.com/3d-printing-company-makerbot-2013-6>

<sup>8</sup> A 2007 article in the New York Times suggested that "Adopters must be ready to develop the same skill as the early photographers who juggled glass plates and egg white emulsions in total darkness." <http://www.nytimes.com/2007/04/05/business/05scan.html>

Ultimaker.<sup>9</sup> Despite this, MakerBot is still a market leader in desktop 3D printing, with devices that are more readily found on Walmart shelves than hackerspace workbenches. Today's 3D printing landscape, however, is no longer shrouded in the emancipatory claims of a decade ago. Its users have been unyoked from the responsibility of pioneering new digital ecosystems. The personal manufacturing revolution appears to have stabilized in the long trough of disillusionment.

Being attentive to the social context through which these issues have unfolded has made it possible for new entries to capture market share, push innovative design practices, and forge new communities of practice. This competitive area, ultimately, is pointing toward machines that come ever-closer to the goal of one-button printing, a trope lifted entirely from Star Trek's replicator, which Makerbot's most popular line of printers is named after (and take their design inspiration from). A longer historical perspective, though, reveals various aspects crucial to understanding the design space of 3D printing today. Are there, or have there been, common values that circulate within the 3D printing community (e.g. the immediacy of one-button printing)? How have these values been *naturalized* by early adopters, regular users, and the general public? Woven into the fabric of the desktop 3D printing community is a persistent notion of it as a liberatory technology that will enable new forms of creativity and reduce the drudgery of assembly line labour. While the focus toward decentralized power structures that supports this liberatory notion might be interpreted as an aspect of reflexive design, something else distinguishes 3D printing as an "interventionist" technology. What is it about the 3D printing community that has resulted in it being associated with anarcho-libertarian tech hackers (e.g. Defense Distributed); flipped classroom education disruption; and body modification, extension, and repair (e.g. Project e-Nable) – each community claiming some form of emancipation-via-technology-intervention narrative? What forces instill the liberatory-emancipatory-interventionist ideal as a naturalized virtue in 3D printing and desktop manufacturing projects more generally?

It is our contention that additional values, beyond emancipatory/liberatory ideals, are at play in modern conceptions of 3D printing. Specifically, these include ideas about what creative and expert labour amounts to, where it is located, and how it can be automated, stemming primarily from developments in CAD/CAM that arose in the post-WWII period at MIT.

### 7.3 Naturalization and Denaturalization

Scholars in the field of Science and Technology Studies (STS), who frequently shape theoretical developments in HCI due to the intellectual proximity of the two fields, have long been interested in the concept of *naturalization*. This term

---

<sup>9</sup>It's impossible to recount the controversy in full here, but see <http://www.tridimake.com/2014/06/do-not-buy-makerbot-3d-printers.html> and <https://www.makerbot.com/media-center/2012/09/20/fixing-misinformation-with-information>

commonly refers to processes that result in actors (scientists; computer users; the general public) assuming that prevailing conditions are natural, part of some grand teleological progress, and not caused by deliberate social, structural, or institutional design. In the context of HCI, naturalization can refer to how embedded assumptions or biases about user behaviour influence normative design approaches. Naturalization occurs as a kind of “settling in” or stabilization of a technology, media, or interaction modality as a sociocultural phenomenon. This includes both the technological artifact and the language used to describe it (Gillespie 2006). This process of stabilization generally entails ongoing processes of negotiation. Prominent STS scholar Sheila Jasanoff has written that “reason is a great naturalizer” suggesting that “once we are persuaded of the reasonableness of an argument or action, it becomes the most natural thing in the world to accept it: of course, this is how things are; of course, this is how things should be (Jasanoff 2012: 6).” Examples of this abound in HCI and design contexts, ranging from hesitation toward mixed reality and multisensory interfaces to critiques of 3D data visualization. As a counter to the naturalizing tendency, Jasanoff recommends looking at phenomena as if “through the eyes of visitors from other worlds” – a process of “making the familiar strange.” This process is what we refer to in this chapter as *denaturalization*.

Third wave HCI advocates for greater attention to social context. This remains a somewhat vague and challenging directive, decades after second wave HCI and design theorists began to discuss it, because methods for undertaking it in parallel to newly-developed engineering practices have not been fully fleshed out. Scholars in STS recognize that doing cultural-historical analysis in order to situate social context can provide a useful vehicle to denaturalize values embedded in technology. As a methodological tactic, then, denaturalization should be recognized as an important component of a more reflexive design practice. One crucial step in the process of denaturalization is to locate developments historically as a way of spotlighting the naturalized assumptions present within them. To begin this process of analysis in the domain of 3D printing, we turn to the CAD/CAM revolution of the 1960s that produced many of the digital graphical interfaces that we now take for granted. By looking culturally-historically at the development of CAD/CAM, we trace specific moves to replace material expertise with rational computational systems. Many new hardware and software developments in the domain of 3D printing appear to operate with an assumption that creativity resides in the software domain, while routine labour resides in the material/hardware domain. These assumptions, we will show, rely on an ongoing erasure of material expertise in computer-aided design. In building a system for facilitating 3D printing comprised of software and hardware elements, we initially abided by this overarching assumption that the design is complete when the digital designer finishes their work in the software domain. This notion of completeness, we would discover, was in direct contradiction to how the clinical context in which our system would be deployed actually operates. Why did we neglect to recognize this contradiction? Because the assumptions it pushed against had been naturalized in the very technology infrastructure we were trying to reimagine.

## 7.4 A Short Sociocultural History of the Values of CAD/CAM

Beginning in 1949, The Massachusetts Institute of Technology (MIT) initiated a series of research projects under the banner *Innovations in Manufacturing Technology*, which were intended, as the name of the research agenda implies, to radically alter the processes of design and manufacturing. The first project, entitled “Numerical Control” (or NC), ran from 1949 to 1954, and set in motion the development of a new type of milling machine controlled by computers. Running from 1955 to 1959, the second project saw the development of “Automatically Programmed Tool” (APT), a computer language which sought to dramatically reduce both the skill and time involved in the production of control tapes for NC mills. The final project, simply dubbed the “CAD Project,” ran from 1960 until 1969, and did not involve the production of a singular technology.<sup>10</sup> Rather, the CAD project oversaw the development of a series of tools intended to produce “human computer design teams” in order to facilitate “automatic manufacturing.”<sup>11</sup>

From a modern perspective, the light pen-driven CRT display interfaces and punched tape-fed servomechanisms developed as a result of the Innovations in Manufacturing Technology research agenda appear archaic. It is easy to cast them aside as relics of a computational past, but doing so causes us to misunderstand how interaction logics embedded within them do, in fact, have direct bearing on a wide variety of interfaces that bridge today’s software and hardware work. What follows is a brief cultural history of the design context in which NC, APT, and CAD were incubated. In calling attention to this history, we show how the HCI pioneers who worked on these connected projects sought to develop a model of design and manufacturing that would eliminate the need for tacit craft knowledge and establish a purely rational approach to production. This is a model that establishes digital representation of the design object as complete, relegating physical manifestations produced in parallel as inferior reflections of the idealized digital design. While the precise methodologies and techniques used in these technologies have adapted over time, the logics embedded within them remain under the surface of modern CAD applications, including 3D printing.

### 7.4.1 Numerical Control

In 1949, the Parsons Corporation, an engineering services firm established near the end of WWII, was awarded a contract by the United States Air Force to produce a computer controlled milling machine (Parsons Corporation 1952). Shortly after

---

<sup>10</sup> For a detailed history of these three projects, see Reintjes (1991), Noble (1984), and Ross (1978).

<sup>11</sup> A much longer cultural-historical analysis of CAD/CAM is currently being undertaken as a component of ongoing research by one of the authors. Refer to Southwick ([forthcoming](#)) or contact the authors for more details.

being awarded the contract, CEO John Parsons approached MIT hoping to sub-contract development of the control system for the mill to the university's Servomechanism Laboratory. Gordon Brown, director of the Servomechanism Laboratory at the time, saw Parsons' proposal an opportunity to engage a new area of applied research, and an agreement was established allowing MIT to develop a new system based on feedback control (Reintjes 1991). The partnership was short-lived, however, due to a fundamental disagreement between Parsons and a sub-group of MIT researchers, led by William Pease, about the project's underlying goal. From early in the project, Parsons favoured the production of a single-axis demonstration mill, specifically designed to produce wing panels, as a means of showcasing the viability of NC technology (Reintjes 1991). Pease, however, was strongly opposed to this approach, as he felt that a single-purpose mill fundamentally misunderstood what NC was capable of. In a memo circulated around the Servomechanism Lab in September 1949, Pease summed up the "general aim" of the project as being able to automatically machine "any mathematically definable surface" (Noble 1984). The breadth of this claim is important. If a part could be mathematically represented, regardless of its complexity or underlying production concerns, it could be produced on a NC mill. Pease regarded the NC mill as an information processing device that could physically manifest any possible design an engineer could dream up (Pease 1952).

Pease's vision won out as a result of the Parsons Corporation lacking the requisite technical experience to build the single-axis mill, in addition to the United States Air Force – which was funding the project – preferring the more ambitious model of NC (Reintjes 1991). In approving the NC project, the Air Force put its chips in on an approach that could increase the accuracy of machined parts used on super-sonic flights, while also reducing the time it took to make them. By framing NC as a scalable information-based technology that claimed to reduce the possibility of human error in the machining process, Pease was able to directly address each of these concerns simultaneously. The design, if successful, would dramatically reduce the need to expert machinists. This aligned favourably with political concerns in post-war America. The USAF was concerned about organized labour largely because of a series of strikes that occurred in the latter half of the 1940s. The capacity of a small group of highly-skilled labourers to grind production to a halt was something that the US government wanted to curb (Noble 1984). Pease's vision of NC seemed to nullify these political concerns by removing highly-skilled craft labour from the machining process and replacing it with "mathematically-defined" routines developed by engineers further up the chain. Within this context, accuracy was aligned with both the capacity to machine parts at higher tolerances, and with a design orientation that enforced the authority of engineers. We can consider this design orientation to be a key value of CAD/CAM.



### ***7.4.2 Automatically Programmed Tool***

The development of APT, the second project of the MIT Innovations in Manufacturing Technology research agenda, was led by Douglas Ross and initiated in the summer of 1955 (Reintjes 1991). Translating designs produced by engineers into a “readable” format for NC machines proved to be an expensive and time consuming process, one that was also frequently replete with error. It established and depended on a new class of highly specialized worker, running counter to the USAF’s goal of mitigating political concerns associated with labour strife. Despite this, the goal of the APT project ultimately became to instantiate the vision of NC that sought to replace humans with automated processes (Ross 1978). Drawing on the work of John Runyon and Arnold Siegel, two MIT Servomechanism Lab members who experimented with the use of digital computers for NC programming in the early 1950s, Ross developed APT as a series of “English-like” commands that enabled programmers to define parts using points, space curves, and regions – APT-I, APT-II, and APT-III respectively (Ross 1978). Before APT’s public unveiling in February 1959, Ross formalized the language’s underlying philosophy and logic in a succession of papers and internal documents. Outlining a “systems approach” which saw the NC mill, the general purpose computer running APT, and the parts programmer configured together into a single system, Ross argued that each “actor” within the system could be assigned to perform a task best suited to their capacities. Humans – specifically, engineers and parts programmers – would address design issues while machines handled the heavy lifting of calculation and analysis needed to manufacture or improve a design (Ross 1956, 1958). Embedded in this systems approach was a process for clearly demarcating specific actions that take place in each phase of design and manufacturing, along with a hierarchy favouring computational certainty over human/machinist creativity. This hierarchy of tasks and emphasis on the computational can be understood as a second key value of CAD/CAM.

### ***7.4.3 Computer-Aided Design***

Wanting to build off of the success of APT, members of the Servomechanism Lab met with personnel from the Design Division of MIT’s Mechanical Engineering Department in late 1949. The goal of this collaboration was to establish a project that would extend the tripartite systems approach developed by Ross into the domain of design. The two groups eventually agreed to seek funding from the USAF, and a proposal was submitted in 1960 (Reintjes 1991). The specific goal of the Computer-Aided Design (CAD) project was to facilitate the “automatic manufacturing” of parts once “the human computer design team had established the features of the design” (Coons and Mann 1960; Ross 1965). From the outset, a mandate to streamline the design and manufacturing processes was clearly articulated. Steven Coons, a member of the Design Division, summarized the scope of the CAD project in an

early analysis of human-centered design and manufacturing processes. Characterizing traditional methods as ineffective uses of resources, Coons argued that there is an inherent exponential growth of manpower and time as a project progresses (Coons and Mann 1960). His critique shed light on practices that saw a “few engineers performing highly creative tasks at the beginning, coupled with a very large number of draftsmen and technicians who perform relatively uncreative tasks over a fairly long period of time” (Mann and Coons 1965). Coons argued that, outside of the initial creative process performed by engineers, design and production processes would best be performed by computers. This would prevent “noise” from entering the design. By reducing the number of actors, there would be less opportunity for designs to be altered from the engineers’ original intent – a re-articulation of the authority of engineers manifested in NC (Coons 1966). Beyond the initial stage of engineering, additional design and production routines were fundamentally mechanical in nature according to Coons. Delegating responsibility for them to automated machine processes would prevent “man” from engaging in tasks that did not fulfil his creative impulse, while also speeding up design and manufacturing (Mann and Coons 1965). This liberatory ideal articulated by Coons, to free humans from drudgery, would strongly influence the rhetoric of future CAD-based systems (including 3D printing). This ideal can be understood as a third key value of CAD/CAM.

Unlike the NC and APT projects, the CAD project did not have a singular deliverable. Rather, it unified five broadly-stated research objectives that were common among the Servomechanism Lab and the Design Division. Included among these five objectives were the development of computer language for design, and the establishment of a system for pictorial and symbolic communication between humans and computers (Ross 1965). Although the CAD project ran until 1969, the Design Division withdrew from it in the mid-1960s as a consequence of the Mechanical Engineering Department reorganizing its research objectives to focus on applied projects with concrete deliverables (Reintjes 1991). Despite this early withdrawal, the epistemic values of the CAD project had been firmly established. Coons had been able to expand the systems approach initiated by Ross during the development of APT, ultimately proposing a radical new model for design and manufacturing that sought to upstream authoritative control to the engineer. Near the end of his official involvement, Coons co-authored an article that sought to explain the broad effects CAD might have, not only on design and manufacturing, but on society as well. “In 10 years, much of what has been outlined in this article will have come to pass. If we are able to make a human adjustment without trauma to this new era it promises to lead to perhaps the brightest period in the history of the world. Certainly in all material ways man will be the master of his universe. It remains only to be seen whether he will succeed in the deeper, more important ways of the spirit” (Mann and Coons 1965: 9). In conjuring this utopian world that CAD would help construct, Coons reiterated that what originally took an “army of men six months to perform” would be “reduced to one-man one-machine tasks taking only several seconds... A comparatively few extremely capable people, coupled with computers

and machines, should be able to provide the goods for the entire world” (Mann and Coons 1965: 2–8). This utopian ideal regarding CAD/CAM and productive capacity can be understood as a fourth key value.

The goal of this section has been to use the history of the MIT Innovations in Manufacturing Technology research agenda to surface specific values associated with CAD/CAM. To briefly restate, these values include: (1) NC’s design orientation that enforced the authority of engineers and eliminated the need for unnecessary labour; (2) APT’s systems-driven philosophy, which demarcated where specific actions would take place in each phase of design and manufacturing, and enforced a rigid hierarchy that emphasized computational certainty over human creativity; (3) CAD’s instantiation of these ideals by seeking to remove humans altogether, save for the “comparatively few extremely capable people” who would be required to seed the computational processes with suitable ideas; and (4) the emancipatory agenda promoted by Coons that promised a break from the menial toil of manufacturing. Together, these four values – engineering authority, hierarchical ordering of tasks, erasure of expertise, and liberatory technology – have continued to inform the design and production of CAD/CAM hardware and software. Equally, as the MakerBot examples described above highlight, these values remain present in current understandings of 3D printing.

## 7.5 3D PrintAbility

Nia Technologies Inc. (Nia) is a Canadian non-profit social enterprise founded by cbm Canada in 2015 in partnership with the Semaphore Research Cluster and Critical Making lab at the University of Toronto.<sup>12</sup> Supported by exploratory research on prosthetics carried out in the Critical Making lab between 2011 and 2013 (Record et al. 2013) and continuing work in the lab since, Nia’s mission is to discover, develop, and deploy innovative technologies that will improve lives in low and middle-income countries (LMICs). Since its founding, Nia’s main focus has been to address the need for mobility devices, specifically prosthetics and orthotic braces, in the developing world.

Approximately 30 million people in LMICs require prosthetic limbs, braces, or other mobility devices, but only 5–15% of the people who need these devices have access to them. Core barriers to mobility device access in LMICs include affordability, a lack of rehabilitation services and qualified orthopaedic clinicians, and the inability to finance sustainable service delivery. Hospital administrators and clini-

---

<sup>12</sup>Author Ratto is Nia’s Chief Science Officer. All three authors work in the Critical Making lab and the Semaphore Research Cluster. Support for this work included sustaining funding from cbm Canada, and grants from Grand Challenges Canada, Google Foundation, Jericho Foundation, and Autodesk Foundation. A portion of the descriptive text in this section comes from an unpublished clinical trial report. The authors would like to thank Joshua Qua Hiansen, Jennifer Marshall, Jerry Evans, and Nia Technologies for the use of this material.

cians in these countries express a need for new innovations that can help shorten production time, reduce material waste, and improve device affordability. In an attempt to meet these needs in LMICs, Nia developed a set of digital technologies under the banner *3D PrintAbility* (3DPA). Nia's primary objective with this has been to increase access to mobility devices for young people with disabilities in LMICs by supporting the development, deployment, and use of 3DPA in appropriate contexts.

### 7.5.1 *Overview of Traditional Process*

The majority of prosthetic and orthotic device production is conducted using a manual process. Transtibial (TT) sockets and ankle foot orthoses (AFO) are created by prosthetists and prosthetic technicians through a series of time-consuming routines. First, a plaster-cast wrapping is applied to the patient's limb in order to generate a negative cast. Once set, the cast is removed and plaster of Paris is poured into the new mold to create a positive of the patient's limb. Prosthetists then apply a series of build-ups or removals onto the surface of the positive cast. This process, called rectification, is done to shape the device to the patient's specific anatomy and needs. Typically, build-ups are used to reinforce weight-bearing areas on devices, and material is removed where the force of movement is not expected to be high. These processes save on weight, increase comfort, and aid in the overall function of the device. Upon completion of the rectification stage, a heated sheet of polypropylene is wrapped and shaped tightly around the rectified positive cast. Once cooled, the shaped polypropylene is removed from the built-up positive cast, and a series of cuts made by the prosthetist shape and complete the device.

Current manual methods are effective in producing sturdy, comfortable, and usable devices. However, the process is extremely time-consuming, requiring the prosthetist to be hands-on throughout the entire manufacturing cycle. In addition, the initial plaster-cast mold is lost as a consequence of the manufacturing process. Prosthetists, therefore, must repeat the entire cycle each time a patient needs a re-fit or a new device. These limitations restrict the productivity of prosthetists and exacerbate the problem of device availability. It is clear how an opportunity for innovation was brought about by the emergence of 3D printing, a technology that purports to solve the problem of ongoing reproducibility and adaptation by moving significant aspects of creativity, preparation, and manipulation into the digital realm.<sup>13</sup>

---

<sup>13</sup>To be clear, the history of CAD/CAM and prosthetics pre-dates 3D printing. Early research includes work by some of the original CAD/CAM researchers, including Robert W. Mann (<http://news.mit.edu/2006/obit-mann>). Currently, various CAD/CAM systems are used by prosthetic practitioners. This primarily involves specialized CAD software and subtractive milling; cf. Smith and Burgess (2001).

### 7.5.2 *Nia's Technical Innovation*

Nia's 3DPA platform consists of 3D scanning, 3D rectification, and 3D printing technologies woven into a single system. The combination of these processes aims to enhance the productivity of prosthetists and addresses the need for increased availability of prosthetic and orthotic devices in LMICs. Nia's 3D scanning capabilities enable prosthetists to get accurate and instantaneous true-scale 3D models of a patient's limbs, eliminating the need to apply plaster-cast bandages and the need to create a physical positive cast. OrthoGen, Nia's custom 3D rectification software, uses modern computer-aided design software to enable prosthetists to readily augment and adjust scanned digital limbs in 3D space. By working in an entirely digital environment, a digital history of patient-specific models is stored, enabling prosthetists to refer to previous steps and make adjustments as required. Nia's software also allows prosthetists to send their digitally designed devices to a fused deposition modelling (FDM) 3D printer for fabrication. These 3D printers work by converting a solid, high-strength thermoplastic filament into a semi-molten state that is deposited in a predefined and programmed shape to cool and solidify. The printing process is entirely driven by the printer, which allows the prosthetist to attend to other clinical activities, such as caring for additional patients. Nia sees its 3DPA system as a means for prosthetists to address the evolving needs of their profession, while maintaining and enhancing their skills, craftsmanship, and expertise.

### 7.5.3 *Explicit Values in Designing 3D Printing System*

As may be clear from the above narrative, Nia's work starts from a specific value position that emphasizes the importance of the clinical expert within the prosthetic domain. Rather than attempt to replace the prosthetist with an automated system, Nia's software and hardware tool chain has been explicitly designed to supplement and extend the ability of clinical personnel to provide care for prosthetic users. This is in no way accidental. This explicit value – the idea of technologies augmenting rather than replacing human capacities (Viseu 2003) – predates the commencement of Nia's relationship with the Critical Making lab. In fact, this value colours most of the work we do in the Critical Making lab, and its conceptual origins lie in a liberatory agenda drawn from the work of Critical Theory scholars (e.g. Horkheimer 1972). This value position was reinforced by the lab's initial exploratory work on prosthetics, which involved a series of workshops on 'extending the body,' wearable technology, and digitally-mediated embodied interaction (Record et al. 2013). These events brought together artists, prosthetists, prosthetic users, social scholars, and designers, and involved open-ended 'making' activities and short professional lectures. The results of this work – simple cast body parts, interactive electronic wearables, and 3D printed artifacts – acted as objects for discussion and reflexive prompts, rather than functional prototypes. An initial aspect of this work, captured

in the title of a workshop series we initiated entitled “DIY prosthetics” was the personalized creation of technologies for bodily augmentation. It was through these events that we first came into contact with prosthetist caregivers and users of prosthetics and, accordingly, how the complex clinical relationships between clinical staff and prosthetic wearers first became evident to us.

These experiences very much coloured our design ideas regarding the nexus of 3D printing and prosthetic development. From the beginning, our focus was on collaboration with clinical experts, both in developing and developed world contexts. In setting up the initial design constraints and requirements, we met multiple times with prosthetic technologists at CorSU hospital in Uganda, carrying on an extensive email conversation about needs and specifications. At the same time, we discussed our design goals with clinical experts in a local context, including prosthetists and orthotists at Holland Bloorview Kids Rehabilitation Hospital and the George Brown College Prosthetic & Orthotic program.<sup>14</sup> Once initial prototypes of the various software and hardware processes had been developed, these professionals assisted us with testing and evaluation, in many ways acting as co-designers of what ultimately became the 3D PrintAbility system. Our design activities very much followed a standard ‘values in design’ process whereby values such as social justice, equity, and autonomy were highlighted as key goals within production (cf. Nissenbaum 2001).

We can usefully compare and contrast this co-design process with the work of other contemporary organizations working in the 3D printing space. The ‘Enabling the Future’ or *e-NABLE*<sup>15</sup> project, which also began in 2013, bypasses the clinical prosthetist, relying on networks of volunteers with 3D printers to produce one or more parts of an upper extremity prosthetic (e.g. a 3D printed prosthetic hand) for those in need. The devices produced by e-NABLE use pre-designed 3D models that are then re-sized to fit a particular patient, and do not typically include one of the most important aspects of a prosthetic, the socket that interfaces the hand or ‘terminal device’ to the patient’s body. While there are many media stories about the success of e-NABLE’s strategy, clinicians have been rightfully concerned about a lack of training on the part of those producing and fitting these devices, as well as how individuals involved in e-NABLE may disseminate an incorrect understanding of the work involved in producing prosthetic devices and taking care of prosthetic users.

Much of our early work in developing Nia involved a rehabilitation of 3D printing as a technology which might serve the prosthetist (rather than a technology that can be used to challenge their expertise). This value – *support rather than replace the agency of the prosthetist* – had direct material impacts on the design choices we made in designing 3D PrintAbility. It required us to weigh the various activities our users described as creative articulations of their skill. These choices informed the workflow that underlies our software. For instance, prosthetists believe that one of their core skills lies in reshaping the internal volume of a socket to fit the specific anatomies of users. In this process, described above as rectification, prosthetists use ana-

---

<sup>14</sup><https://www.hollandbloorview.ca/> and <https://www.georgebrown.ca/prostheticsorthotics/> respectively.

<sup>15</sup><http://enablingthefuture.org/>

tomical and kinesiological knowledge to manually add or remove material from the internal dimension of a socket. Our software needed to support this activity, while granting as much agency as possible to the prosthetist-user in order for them to apply their tacit knowledge of this necessary stage. Conversely, the process whereby a socket is actually produced (the thermoplastic process described above), while understood as requiring craft skill and experience, is not where prosthetists believe their primary knowledge to reside. Whereas our software provides numerous tools for rectification and related aspects of the prosthetic design process, shape manipulations that prosthetists felt did not involve their creativity were partially automated. In this, we followed the assumption that habitual labour can and should be automated – an assumption that, as we’ve seen, can be traced to the development of CAD.

### ***7.5.4 Missed Opportunities – Naturalized Assumptions in 3D Printing***

As the above section makes clear, our work on 3D PrintAbility leveraged a co-design process intended to surface where prosthetists’ knowledge and creativity reside. We followed an explicit value objective – to reinforce the agency of the prosthetists – and designed our system around this value. Our goal here was to increase the agency of the user of our system, to support the extension of the user’s own skills, and to make the system valuable to their professional practice. Through interviews, observations, ‘talk-throughs,’ and user testing, we tried to engage our users as co-designers of the system. Initial research indicates that this process was successful.<sup>16</sup>

But in doing so, we bought into the logics of CAD – that creative work happens in the early stages of digital design; that the digital artifact becomes a kind of sufficient and complete object; that certain forms of routine work should be automated – when we made an early assumption that prosthetists would require automated support in the area of rectification. The logic that Coons helped manifest in CAD, that if creative work has to happen in the material space it should be recognized as a kind of failure, completely discounted the creativity and knowledge that prosthetists would later describe to us. This firewall between where digital models are created (in CAD software) and where they are materially instantiated (in slicing software that delivers the machine path to the printer) is indicative of a long-standing problem that takes the space of digital information processing to be where the real creative work gets done. We can recall how Pease first articulated this notion when he positioned NC as a scalable information processing technology. In fact, creativity in the hands-on rectification process was where prosthetists considered their knowl-

---

<sup>16</sup>Clinical trials carried out between 2015 and 2017 demonstrate the success of our system (Ratto et al. 2017). Equally, end of trial semi-structured interviews with users of the system indicate that the goals of supporting user agency were successful (unpublished clinical trial report; please contact authors for more details).

edge to firmly reside, and they expressed a need to have control and agency in this space. Knowing the history of CAD helped us surface and denaturalize these assumptions about creativity and expertise – to see them as historically motivated. Only then could we begin to imagine alternatives.<sup>17</sup> A truly values-based design must engage values that are surfaced not only by acknowledging their existence and ascribing them to some specific group or actor (i.e. *whose values?*). It must also strive to denaturalize their logics, even when we, as technology designers, have been captured by them. Denaturalizing our own assumptions, then, in addition to not attempting to erase the valuable expertise inherent in routine prosthetic work, became our design prompt for developing better 3D systems.

## 7.6 Conclusions and Third Wave Concerns

The historical and contemporary cases we've highlighted speak to a number of third wave HCI concerns. In particular, prescriptions that advocate context awareness, along with concomitant claims about what accounting for situatedness actually entails, are important to weigh in depth. When the seminal papers of third wave HCI (Bødker 2006; Harrison et al. 2007) extended the second wave's focus on context, they largely framed this activity as *dynamic use context*, separating work activities from non-work and leisure ones, with technical objects playing the role of mediator. Researchers began to focus not merely on how context provides meaning to design, but on how design accommodates context (Harrison et al. 2011). Local, situated practices moved to the fore, placing greater emphasis on *social context*. It is our claim, however, that taking social context into account is insufficient. Meaningful contextual awareness must consider historical antecedents and the values they instill in developer and user communities. (We must also stress that what many in the third wave HCI community refer to as *situated*<sup>18</sup> is not entirely commensurable with what many in the STS community call situated – even when they share some overlapping concerns and bodies of literature.<sup>19</sup>).

---

<sup>17</sup>There is a connection here to how Bardzell and Bardzell (2016: 26) describe as *emancipatory design* in the context of *humanistic HCI*: "... the fusing of the humanistic use of critique as speculation with design activities. Doing so enables HCI researchers and practitioners to interpretively explore alternate worlds."

<sup>18</sup>See Harrison et al. (2007: 10) for elaboration on the three ways HCI uses this term. They suggest that each notion - the interactionist, the ecological, and the cultural – is a "systems approach" that primarily seeks to understand the relationship between system elements and activities. In a 2011 follow-up paper, Harrison et al. elaborate in greater depth, suggesting that "taking situated knowledges seriously" will require researchers to explicitly articulate the intellectual and political commitments they bring to a particular project (Harrison et al. 2011: 392). It remains to be seen when or how this will commence (outside of the occasional alt.chi paper).

<sup>19</sup>See Haraway (1988) with its specific emphasis on re-drawing the boundaries of scientific objectivity, which, in many ways, calls into question the rigid empiricist methods that third wave HCI has a difficult time shaking. Suchman (2007: 13–17) and Frauenberger (2016: 344–345) offer descriptions that bridge HCI and STS concerns.



Harrison et al. (2007) acknowledge that seeking updated understandings of system efficacy – “asking what it means for a system to be ‘good’ in a particular context” – promotes alternative success criteria indebted to a value-sensitive approach. Third wave HCI’s attention to values in design rightly asks whose values are embedded, and whose values determine success criteria, but rarely traces how those values come to be naturalized – save for when those values are associated with first or second wave HCI (e.g. efficiency). In the case of CAD, we can follow a linear narrative from Coons’s liberatory claims about ending the drudgery of work through to similar ones made by MakerBot’s Bre Pettis a half-century later. There is a language and (expanding) literature for dealing with the value-laden nature of technology development. It can be leveraged not just for showing others what they have done wrong (in the form of critique), but for constructive purposes as well. It should not, however, be limited by concerns related to system efficacy. By using cultural-historical methods, third wave HCI researchers can situate epistemic values historically, enabling a more nuanced picture of where they emerge from, how they create users and audiences, and what influence they have on dynamic use contexts. Once situated historically, though, we must continue to engage values that have been surfaced by recognizing how they influence our own normative assumptions as researchers.

Our argument for denaturalization is linked to additional ideas that have been identified with third wave HCI. In particular, these include the characterization of 3D printing’s design practices as *“messy” rather than formal and principled*; recognition that *no single set of methods* is sufficient for validating and understanding emergent technologies; and the centrality of a phenomenologically situated understanding of reflexive and interventionist 3D printing design, one that is *grounded in the materiality of its context and technologies* (Harrison et al. 2007). We have focused our argument on the importance of tracing historical context as a crucial step toward situating social context. Additionally, we have argued that merely acknowledging values in design is insufficient. Denaturalizing values – epistemic biases, assumptions about use – that are embedded in technologies, interaction paradigms, and social behaviour is also a necessary part of situating context. This includes denaturalizing the values that researchers and developers bring to the design process, a process that the shorthand “reflexivity” doesn’t fully capture. This is a complicated process, given that we are not always aware of the values we bring to the table when we design new technologies.

How can the insights we’ve articulated be leveraged to build more ethical and empirically grounded 3D printing systems? Further, how might they be deployed within third wave HCI methodologies, in spite of the “chaos of multiplicity in terms of technologies, use situations, methods and concepts” the field currently finds itself in (Bødker 2015: 31)? 3D printing appears, on the surface, to be an emergent technology. It invents its own future, we are told, portending not just a revolution in manufacturing, but reimagined human agency. But there are very few *emergent* technologies. The claims made about 3D printing by the startups of a decade ago echo previous claims that need to be illuminated. Desktop 3D printing is, arguably, still in its infancy. Its relationship to HCI has been structured mostly through a patchwork

of software tools that purport to solve various aspects of the challenge of 3D modeling. For third wave HCI to understand and intervene in the 3D printing landscape, its researchers must be clear that hardware cannot be severed from the historical context of design software. The location of creativity cannot be relegated to the rational design space of engineers. And the value statements made about 3D printing must be taken as if they have come from “other worlds,” because, in fact, they do.

## References

- Bardzell J, Bardzell S (2016) Humanistic HCI. *Interactions* 23(2):20–29
- Bødker S (2006) When second wave HCI meets third wave challenges. In: Proceedings of the 4th Nordic Conference on Human-Computer Interaction, Gothenburg, Sweden, October 2006. Association for Computing Machinery, pp 1–8
- Bødker S (2015) Third-wave HCI, 10 years later – participation and sharing. *Interactions* 22(5):24–31
- Bookchin M (1968) Towards a liberatory technology. *Anarchos* 2:43–60
- Coons SA (1966) Computer, art & architecture. *Art Educ* 19(5):9–12
- Coons SA, Mann RW (1960) Computer-aided design related to the engineering design process technical memorandum 8436. MIT Electronic Systems Laboratory, Cambridge
- Frauenberger C (2016) Critical realist HCI. In: Proceedings of the 2016 ACM SIGCHI Conference on Human Factors in Computing Systems, San Jose, USA, May 2016. Association for Computing Machinery, pp 341–351
- Gillespie T (2006) Engineering a principle: ‘end-to-end’ in the design of the internet. *Soc Stud Sci* 36(3):427–457
- Haraway D (1988) Situated knowledges: the science question in feminism and the privilege of partial perspective. *Fem Stud* 14(3):575–599
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: alt.chi Session of the ACM SIGCHI Conference on Human Factors in Computing Systems, San Jose, USA, May 2007. Association for Computing Machinery
- Harrison S, Sengers P, Tatar D (2011) Making epistemological trouble: third-paradigm HCI as successor science. *Interact Comput* 23(5):385–392
- Horkheimer M (1972) *Critical theory: selected essays*. Continuum, New York
- Jasanoff S (2012) *Science and public reason*. Routledge, New York
- Mann RW, Coons SA (1965) *Computer-aided design*, McGraw Hill Yearbook Science & Technology. McGraw Hill, New York
- Nissenbaum H (2001) How computer systems embody values. *Computer* 34(3):120–119
- Noble DF (1984) *Forces of production: a social history of industrial automation*. Knopf, New York
- Parsons Corporation (1952) *Digitron*. Parsons Corporation, Traverse City
- Pease W (1952) An automatic machine tool. *Sci Am* 187(3):101–115
- Ratto M, Qua Hiansen J, Kaweesa M, Taremwa J et al (2017) An international multi-center clinical study: gauging patient experience with digitally designing, 3D trans-tibial prosthetic devices compared to manually produced prosthetic devices. In: International Society for Prosthetics and Orthotics World Congress, Capetown, SA, May 2017
- Record I, Ratto M, Ratelle A, Ieraci A, Czegledy N (2013) DIY Prosthetics Workshops: “Critical Making” for Public Understanding of Human Augmentation. In: IEEE International Symposium on Technology and Society (ISTAS): Social Implications of Wearable Computing and Augmented Reality in Everyday Life, Toronto, Canada, June 2013. Institute of Electrical and Electronics Engineers, pp 117–125

- Reintjes JF (1991) *Numerical control: making a new technology*. Oxford University Press, New York
- Ross DT (1956). Gestalt programming: a new concept in automatic programming. In: Papers presented at joint ACM-AIEE-IRE Western computer conference, San Francisco, USA, February 1956, pp 7–9
- Ross DT (1958) *Papers on automatic programming for numerically controlled machine tools*. Servomechanisms Laboratory, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA
- Ross DT (1965) *The MIT Computer-aided Design Project A Half-Decade Summary*. Computer Applications Group, Electronic Systems Laboratory, Massachusetts Institute of Technology, Cambridge, MA
- Ross DT (1978) Origins of the APT language for automatically programmed tools. *ACM SIGPLAN Not* 13(8):61–99
- Smith DG, Burgess EM (2001) The use of CAD/CAM technology in prosthetics and orthotics – current clinical models and a view to the future. *J Rehabil Res Dev* 38(3):327–334
- Southwick D (Forthcoming) High-noon on the creative frontier: Configuring human machine expertise
- Suchman L (2007) *Human-machine reconfigurations: plans and situated actions*. Cambridge University Press, Cambridge
- Viseu A (2003) Simulation and augmentation: issues of wearable computers. *Ethics Inf Technol* 5(1):17–26
- Winner L (1983) Technologies as forms of life. In: Cohen R, Wartofsky M (eds) *Epistemology, methodology, and the social sciences*, springer. Springer, Dordrecht, pp 249–263

# Chapter 8

## Physical Computing | When Digital Systems Meet the Real World



Alan Dix and Steve Gill

**Abstract** Maker culture, from soldering sensors on an Arduino to 3D printing a prosthetic limb, has established that hobbyist computing is intimately rooted in the physical world. In education, ‘physical computing’ courses have captured this interest, introducing code through its physical interactions. Interpreted more broadly, physical computing sits at the nexus of a number of strands within HCI including tangible interaction, ubiquitous computing, and spatial/mobile systems. Ideas of embodiment and an experiential approach to design are natural frameworks within which to view physical computing and so it is almost tautologically third wave. However, the hidden action of computation in certain kinds of sensor-rich ubicomp and the AI turn in computing calls any simple identification into question. Product design appears to encounter the ‘waves’ in a different order; as its artefacts become more digital, it is having to consider the agency of computing and adopt more analytic approaches in research and design. Physical computing forces us to regard the ‘waves’ less as a teleological progression, and more as complementary approaches addressing different facets of human experience with physically embodied digital technology. Furthermore, it suggests there are new challenges ahead as we seek to find research and design paradigms that use physical objects as part of rich collaborations with active computation.

### 8.1 Introduction

For some, the term ‘physical computing’ conjures up images of heads bent low over soldering irons, Arduino and Raspberry Pi: digital-maker culture for the arts, for DIY domestic automation, and just fun.

---

A. Dix (✉)  
HCI Centre, University of Birmingham, Birmingham, UK  
e-mail: [alan@hcibook.com](mailto:alan@hcibook.com)

S. Gill  
Cardiff Metropolitan University, Cardiff, Wales

However, this hacker image sits at the heart of a matrix of activities in academia and industry. In some areas, such as industrial control, computing has always been very physical, but within HCI the trend to physicality comes after a long period when the focus of design was on the digital aspects and centrally the screen, and where, as far as possible, the physical manifestations in laptop or phone were abstracted away. Even in our mobile-first design era, responsive design often effectively assumes all devices are largely equivalent up to juggling of regions around the display.

In the context of this book, writing about physical computing and third wave HCI seems almost tautological as ubiquitous computing and physical embodiment were two of the driving factors of the formulation of the ‘third wave’ (Harrison et al. 2007).

Embodiment is at the heart of the *application* of physical computing. According to Tom Igoe, one of the authors of the first specific textbooks on physical computing,

[i]n physical computing, we take the human body as a given, and attempt to design within the limits of its expression (Igoe 2004).

There is also clearly an element of ‘*thrownness*’, as Heidegger (1927) would describe it, in the *methods* of physical computing, the deep engagement whether in coding or wiring and the embodied focus on making and playing.

However, the phenomenological language of third wave HCI would sound alien to many engaged in these acts, and indeed for the makers of Raspberry Pi, their purpose is to:

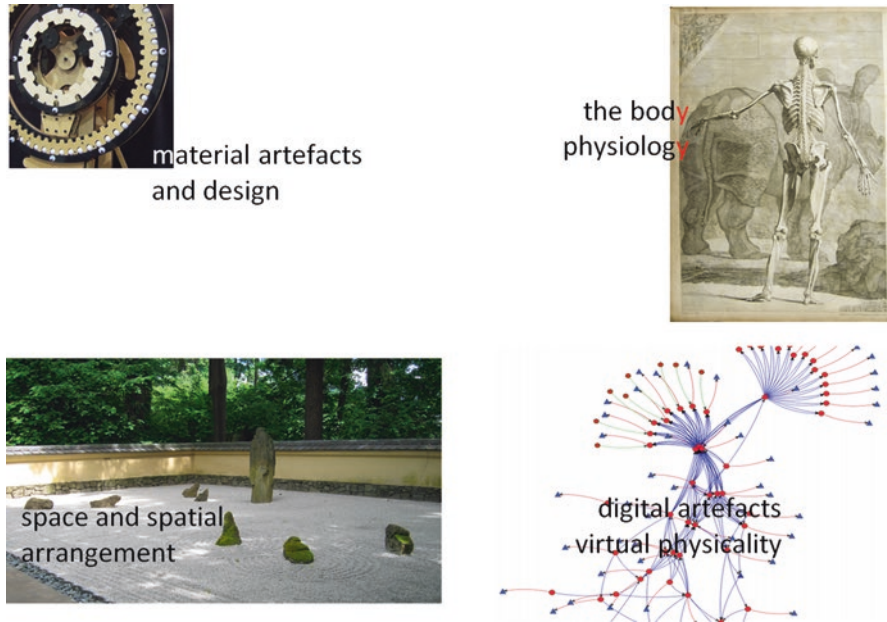
put the power of digital making into the hands of people all over the world, so they are capable of understanding and shaping our increasingly digital world, able to solve the problems that matter to them, and equipped for the jobs of the future (Raspberry Pi Foundation n.d.)

This educational focus on ‘*understanding*’ suggests reflection and breakdown, a deeply cognitive (second wave) activity. Of course this cognitive purpose is achieved through *first* engaging in making and through that concrete action obtaining more abstract understanding. This both defies any simplistic third wave vs. second wave dualism, but also demonstrates a clear break from understand-then-act models of education.

Of course the phenomena of design are different from the phenomena of the designed artifact. The ‘*take the human body as a given*’ approach of physical computing on the surface emphasises embodied interaction. However, the frequent hiddenness of sensors runs counter to Susanne Bødker’s identification of the central role of the “*visibility of meaning and meaning-making*” in her 10 year retrospective on third wave HCI (Bødker 2015).

So the relationships between the physical as a strand in historic HCI research, the HCI implications of physical computing, and the future for physical computing in society are far simpler technologically than they are socially and philosophically.

In the rest of this chapter, we will explore different ways in which physical action in the world relates to HCI. We start with an overview of the background in physical



**Fig. 8.1** Aspects of the physical world affecting design, from TouchIT (Dix et al. 2018)

computing defined both narrowly and more widely in the various strands of research in HCI that take into account physicality. We will then look at product design practice and methodology. As everyday devices in the real world increasingly have digital aspects, HCI and product design have drawn closer; indeed this chapter’s authors have been collaborating on a recent book on physicality and design (Fig. 8.1). Stepping back from these disciplinary roots, we will consider how our design methodologies have to adapt as computing becomes an integral part of the physical world affecting everything from the home to society as a whole and how the physical nature of computation itself is becoming significant. Finally, we will try to draw together the theoretical threads that have emerged, reflecting on the rich interplay between the methodological stances in the three ‘waves’ and some of the challenges that lie ahead as we move from dichotomous paradigms of digital communication and physical manipulation to ones where human and computer agents collaborate on and through a material world.

## 8.2 What Is Physical Computing

The term ‘physical computing’ can sometimes be interpreted quite widely, referring to the broad range of ways computers interact with the physical world through sensors and actuators, and in this sense connects to a wide range of other aspects of this

‘turn to the physical’ including the internet of things, cyber-physical systems, and ubiquitous computing. More narrowly, it refers to the emergent hacker culture using microcontrollers including Arduino and Raspberry Pi to connect with sensors to create interactive installations and DIY automation.

In this chapter we will both take in the wider picture, but also examine the special implications of the more specific hacker-culture.

The roots of the emergence of the narrowly defined physical computing are owed in large part to the Arduino. The Arduino was itself developed at Interaction Design Institute Ivrea (Kushner 2011), so the connections to interaction design are in its roots, and a strong part of its initial growth came from more artistic use rather than mainstream computing.

The success was driven by a number of factors including (relatively) easy to use programming and ease of connection of electronic components, and extensibility. The Arduino’s early adoption of open source hardware was almost certainly also a critical element: this restrained costs, allowed third party additions, and also gave it an ‘alternative’ cachet.

However, perhaps the best illustration of the change in emphasis is in the educational domain. In 1981 the BBC Micro was released to accompany a computer literacy programme (in both senses of the word) launched by the BBC (Garcia 2012). It became the de facto computer used in schools in the UK for the next 10 years. Thirty five years later, in 2016, the BBC micro:bit was launched, as part of a similar schools computer literacy campaign (BBC 2016).

Oddly, despite the 35 years gap, the two are not massively different in crude computing power: the original BBC Micro ran at 2 MHz, with 16 K RAM and 32 K ROM, the BBC micro:bit clocks at 16 MHz with 16 K RAM and 256 K ROM. However, while the original BBC Micro was a desktop computer, approximately a half metre in each direction, the BBC micro:bit is a circuit board 4 cm × 5 cm in size. Furthermore, while the BBC Micro sold at £300–400 (approx £1500 current prices), the BBC micro:bit costs around £10.

Through the eyes of present-day physical computing, the BBC Micro was prescient in that the Model B exposed 8 general-purpose digital I/O pins (GPO) as well as 4 analogue inputs. Some hobbyists, researchers and industrial applications exploited the relative ease of electronic projects on the platform; but the overwhelming use of the BBC Micro was for screen-based applications. In contrast, the use of the BBC micro:bit is focused virtually entirely on hardware electronics.

So the change in 35 years is partly about cost and scale, and this is not insignificant in an educational setting where the cost of a hardware mistake frying a computer becomes manageable. However, the greater difference is the way these and societal changes have changed the conception of computing.

For the 1980s schoolchild the exciting impact of computers was digital: glowing characters on the screen (pixels came later); in the 2010s it is physical: buzzers and LEDs, motion detectors and motors.

### 8.3 Physicality in HCI

Bodily interactions were central to some of the very earliest work in HCI where ergonomics was one of the early contributory disciplines (first wave HCI). At that stage, when data-entry computer operators would sit all day typing at a workstation, safety concerns about posture, finger strain and screen exposure led to detailed studies of different devices and DIN keyboard standards.

This work faded quite quickly from HCI research, partly because its novelty faded and possibly also because commoditization of computer production effectively shifted the responsibilities for safety. However, strands of empirical work around Fitts Law and similar motor tasks do still continue.

Note that it is not that physical ergonomics have ceased to be important, as many who have struggled with low-travel keyboards or maximally RSI-inducing trackpads can attest, but the buying decisions tend to be focused more on aesthetics and immediate experience. One of the lessons of ergonomics research has always been that many damaging effects are not immediately felt as discomfort, but only become apparent after time, and so these rarely form part of individual decision making until too late.

Very quickly these physical roots gave way to several decades of what can be best thought of as virtual physicality. Building on earlier work by Sutherland (1963) and early GUI interfaces such as the Xerox Star (Smith et al. 1982), Shneidermann (1983) coined the term *direct manipulation*, which was also the core of Norman and Draper's influential volume *User-Centered System Design* (1986). At the heart of direct manipulation is a virtual world (typified, although often satirically by the 'desktop' metaphor), where digital objects are operated upon 'directly' as if they were physical objects.

The theoretical foundations of this conception drew partly on more cognitive views (second wave), but also on Gibson's ecological views of perception, his notion of affordance, and Heideggerian ideas of thrownness (third wave). Although the term was not used, effectively DM was about a form of embodiment in the digital world.

Crucially, this 'computer as digital objects' conception was set against earlier command language interfaces, where the computer was effectively anthropomorphised as a form of mediated interaction with the data (objects) through the computer (interactional partner).

Although touch-sensitive screens have allowed the interactions to be more direct, and devices have radically changed from VDU to tablet, phone or public display, this is still the dominant interaction paradigm for day-to-day computer use outside gaming.

Of course, in HCI research, and more gradually in commercial applications, we have seen a move over many years back towards more physical conceptions of



computation, notably in Ishii and Ullmer's tangible bits (1997), which led to the area of tangible computing, and Weiser's coining of ubiquitous computing (1991).

While both of these focused on computation being 'outside the box' and in the world, tangible interaction was about making computation available in physical tokens whereas Weiser's image was of invisibility. However, the distinction is perhaps not as stark as it at first sounds.

Weiser's vision was of a world full of screens (very visible!), but where they *become invisible* through use – indeed this has come to pass, you do not think "I'll interact with the computer to do my washing", but simply set the programme that happens to be displayed on a screen. In Heidegger's terms the displays are precisely '*ready to hand*'. Often this term is assumed to relate to being 'walk up and use', but Heidegger's defining example of the hammer was not of a novice picking up a hammer, but of the experienced craftsman for whom the hammer has become, through use and skill, effectively invisible as attention is focused on the nail in the wood (Dix 2010).

In tangible computing, the physical tokens that denote digital objects are very visible. In one early example, blocks on a table denoted the positioning of machines on a factory floor, through which a camera block could be moved to create fly-throughs (Rauterberg et al. 1998), in other systems the tokens are interaction objects, such as resize or reshape handles (Fitzmaurice et al. 1995).

There are two things to note here. First, just as with Weiser's ubiquitous displays, in a well-designed tangible interaction the blocks become 'invisible', unnoticed, even when they are effectively computational handles. This is perhaps even more clear in touch-table interactions, you do not think "I'll slide my finger to resize the photo", but simply "I'll make this photo bigger". Second, although the tokens themselves are physical and control digital objects' properties and locations in a 'direct' way, they often result in relatively complex computational effects, from simulated urban transport flows to timetabling.

The *physicality of space* has also become a core part of many parts of HCI research and commercial practice. While some mobile applications, such as email, are about making location *not matter*, others, such as maps and location-based social networking, use GPS or other forms of sensing as a central part of the interactive experience. Similarly smart rooms, buildings and cities use sensor-rich environments to craft context-aware interactions, from opening a door to planning traffic light routing. In both cases, it is less the physicality of the computational device itself that is critical, but the way in which it either senses or acts as a proxy to sense the physical location, and environment of individuals.

Of course, this is not to say that the physical design of a mobile device is not also important. Indeed, phones are often marketed less as communication or computational tools than as fashion statements or experiences.

In this and the design of many other devices and appliances, HCI meets product design.

## 8.4 Connecting with Product Design

While HCI sits at the nexus between the digital and the human, product design is at the nexus of the physical and the human.

On the one hand physical considerations are crucial to product design. Designers need to understand the properties of the materials with which they are working, how these properties will impinge on the potential use including structural strength, weight, tactile properties and appearance. Commercial design also has to take into account the physical means of production, for example, a haute couture dress may take many hundreds of hours of hand sewing, but a high-street fashion item needs to be sewn together rapidly by semi-skilled machinists.

On the other hand the product designer has to be aware of human factors in order to ensure that a product can be handled and understood to achieve its purpose. As with interaction design, this usability may not always be right, witness the frequent grazed knuckles from doors with badly designed handles.

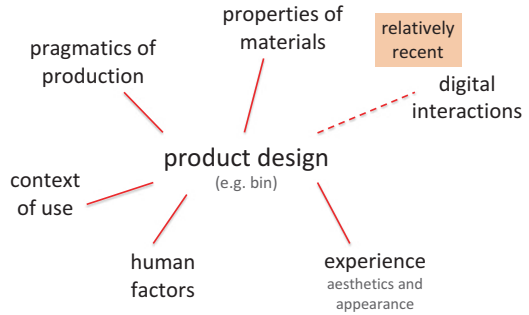
Furthermore, product designers often need to take into account the experiential aspects of their products including aesthetics and appearance. Unlike much of digital user experience design, this typically goes beyond the visual appearance and also has to take in factors such as the tactile qualities of a product, the way it sits comfortably in the hand, and even factors such as the way the combination of feel and weight may give a sense of high (or low!) value. The standard design approach to these tangible but hard to define factors is iterative prototyping, frequently via so-called ‘rigs’. Rigs are designed for testing key dimensions, weighting, comfort and so on. They are frequently designed to be adjustable to enable very rapid iteration and they work well for traditional products. Unfortunately they have proven more problematic to employ for computer-embedded products where the division between interaction with the physical product and interaction with the computer within is somewhat blurred (more below).

Again, just as in interaction design, understanding the context of use is often critical and involves both physical features of the environment and the human usage including commercial and social aspects. The interactions of certain products mean that both designers and HCI professionals have their work cut out to account for the myriad of interactions and contexts the device will be used in. The obvious example is the phone, which might be used as a standalone computer or with others in a series of contexts. On a single working day it could conceivably be used with a projector for a work presentation then to stream music through the car’s sound system on the commute back and then to cast cartoons through the TV for the kids before bedtime. It might even be used to make phone calls at some point!

At its best, product design encompasses a holistic view of the product, its use, its means of production, and increasingly its eventual reuse or recycling.

As an example, one of the authors was once tasked with the design of a bin for city centres (Fig. 8.2). The bin had to resist abuse and so needed to lock, and the

**Fig. 8.2** Factors influencing product design



locks needed to be strong because leverage multiplied the force of a vandal's kick many times. The apertures were designed to accommodate a 12" pizza box, but not large volumes of domestic waste, and rather more seriously, the bin had to have the capacity to be bomb resistant (physical and human *context of use*) while also being easy for municipal employees to remove for emptying and replace (*human factors*). Although throwing things away does not need to be a sublime user experience, the bins were to be in a public place, so elements of *appearance* were also important. The chosen material was a form of plastic chosen for a series of factors including its UV and weather resistance, cost, strength and its high fragmentation and low ballistic properties in an explosion (metal bins create shrapnel – *material properties*). The chosen method of manufacture was rotational moulding (driven partly again by cost and the advantages of that process for products of that scale and form factor – *pragmatics of production*). However, moulded plastics tend to deform slightly (*material properties*) and so the design tolerances needed to take this into account; rotational mouldings typically involve tolerances of  $\pm 1\%$ , meaning that two mating components in a 500 mm diameter product can mismatch by up to 1 cm!

Just as HCI has gradually had to encompass more physical elements into its research and design practice, the ubiquity of computation has meant the product designers have increasingly needed to take into account digital design issues. Indeed, the original creation of the Arduino was precisely to make it easier for design students to be able to experiment with digital design.

In the 1990s when digital appliances were first emerging, it was not uncommon to see product designers attending HCI conferences in an effort to make sense of their new function, and in the years since there has been a growing appreciation of the benefits of learning from design practice within HCI, with many researchers flowing back and forth between computing, design and psychology departments. This period marked the beginning of a period of navel gazing for product design which continues today: the computer-embedded product began a conversation about what a product really is, and today product designers are as likely to design services or even government policy (Policy Lab 2014) as they are injection moulded widgets.

## 8.5 Design Research for Physical Interaction

One of the key methods of design education and research is learning or research by making, the construction of artefacts and through the process of design and construction seeking understanding. Although more traditionally focused on physical form, this tunes well to the methods (if that is an appropriate word to use) of physical computing.

However, a key element of this understanding is experiential: how the artefact feels, the phenomenological impact on designer, critic, and, depending on the branch of design, also user.


This is very similar to the approach taken in technology probes (Hutchinson et al. 2003) and more general intervention-then-comprehension methods within HCI. A technology probe is explicitly created as a putative artefact with sufficient levels of physical and/or digital fidelity that potential users can make sense of the technology, and envisage how it might impact on their work and life – a bootstrapping of potential designs into the space where more incremental and standard user-centred design approaches can work.

Turning to Harrison et al.’s description (Harrison et al. 2007) of the epistemological distinctions between the three paradigms, these practices have elements of both the first “cool hacks; you tried it out and it worked” and third “practice based research; the relationship between your data and what you seek to understand”.

Schön’s (1984) ‘Reflective Practitioner’ has been highly influential; certainly insofar as design thinking has influenced HCI. Through various examples and ‘thick descriptions’ of experienced design professionals (a third wave approach to research *about* designers), Schön concludes that one of the defining features of experienced designers is their ability to reflect upon their own practice both actively during design and more strategically. While the novices appear more lost in their design activity, the expert steps back, in Heidegger’s terms making their design artefacts and even thought processes ‘*present at hand*’. In the end use of the designed artefact *unintended breakdown* is typically a sign of a failure that has brought the device or tool, rather than the ultimate purpose, into attention. In contrast, during the process of design *intentional breakdown* is a means to obtain a higher order understanding that can then be reapplied in new situations – certainly not context free, but creating transferable knowledge.

In some ways this more structured, more generalizable knowledge has elements that are reminiscent of the descriptions of second wave HCI. However, the grounds of that knowledge are quite different (Fig. 8.3), instead of empirical data based on

**Fig. 8.3** Sources of knowledge



physically observable phenomena and objective measurement, this reflexive design practice is based on professional experience and subjective introspection.

Of course, even this dichotomy is a simplification as the professional's experience is based on seeing the outcomes of their own and other people's previous concrete design embodiments, and the empiricist's measurement choices and interpretations of data are influenced by previous experience.


Interestingly as design researchers have engaged with the digital facets of their products, they have often employed more empirical methods. For example, as part of work to understand the validity of a prototyping tool (the IE unit), the authors and colleagues ran experiments varying the level of physical fidelity of physical device mock-ups attached to digital prototypes of the device's screen contents (Gill et al. 2008). The second author's research group conducted substantial research in this vein creating tools such as IRIS (Zampelis et al. 2012) – an augmented reality method, and StickIT (Culverhouse and Gill 2009) – an ultra-rapid method based on RFID tags. Their work was greatly influenced by a number of designer-friendly tools being developed by computing researchers at around the same time. These included DTools (Hartmann et al. 2006) which included both a state transition derived programming interface and physical components driven by it, and Switcheroo (Avrahami and Hudson 2002) – an RFID-based solution, which, together with VoodooIO (Villar and Gellerson 2007) were fundamental building blocks for StickIT.

While the metaphor of 'intellectual waves' suggests a level of progression, considering both HCI and product design suggests that we instead have a number of complementary approaches or viewpoints that together create methodologically rich disciplines. The order in which they appear or are more or less dominant depends on the pragmatics of the underlying phenomena and the development of the fields.

As another example, as part of a project addressing hybrid physical–digital systems, the authors and others created a design notation *physigrams*, based on state–transition networks, but with variations to be able to talk about the complex ways in which simple buttons, switches and knobs behave (Dix et al. 2009, 2010). Of course, the team included HCI researchers, so this was, in part, due to their influence. However, when the designers on the team independently used physigrams as part of an early design project, it helped uncover subtle, but crucial physical interaction differences between superficially similar design choices, thus demonstrating its utility in product design (Fig. 8.4).

Crucially, physigrams did not simply attempt to fit the physical phenomena into a pre-existing notation, but instead explored the felt experience of different controls and then, based on this experience, attempted to capture critical aspects and to establish guidelines on appropriate use. Furthermore the notation has an open-ended nature that allowed designers to tweak and modify to ensure it adequately represented the tactile and experiential qualities of the devices.

This is in part about an open attitude towards the more formalised parts of the notation (e.g. extending the set of allowable transitions), but partly because there are



**Fig. 8.4** Physigrams as used by designers

elements of physigrams, in particular spatial layout, which are not interpreted – that is the notation satisfies one of the heuristics for appropriation (Dix 2007) – “*allow for interpretation*” to leave aspects uninterpreted by the notation/system and thus open to interpretation by the user (in this case the user is a designer).

In terms again of the descriptors of the three paradigms (Harrison et al. 2007), physigrams are certainly focused on “interaction as phenomenologically situated”, but then use the knowledge from this to inform “structured design” and obtain “statements (not necessarily objective) with general (although not universal or context free) applicability”.

One of the authors’ colleagues, Jo Hare, built on the physigrams concept to propose a way of considering the physicality of a prototype in terms of *active* and *passive* physicality (Hare et al. 2014). Passive physicality is the perceived affordance based on the visual appearance and tangibility of the prototype, and active physicality the perceptible experience of interacting with the prototype.

## 8.6 Action!

Much of HCI is focused on closed world interactions. An individual or group of people is interacting with a system running on a device or collection of devices. As part of our design we take into account the “existing situated activities in the world” and “what goes on around systems”, but the system itself does not have any physical impact on the world except through the actions of the users.

Of course, even traditional desktop systems often have printers, which create physical change on paper. However, by and large, even where the devices we deal with have important aspects of physical form, the majority of the effects we create are digital. In addition, the closed-ness of a system depends on where we draw the boundaries, in particular communication technologies from email to social networks allow non-local effects, but again largely limited to digital influences mediated by other people.

This is less true for product design, where, for example, a washing machine cleans real clothes. Traditionally these effects have been physically local: the clothes washed in the machine, the water heated in the kettle, the hole drilled in the wall. The user's conceptualization of the effects is also typically instrumental, the drill is a tool that *you use* to make the hole. Of course, digital technology has changed this, and certainly it may seem that we coax the washing machine to do what we want rather than simply using it.

However, from industrial robotics to voice based home control systems, drones to autonomous cars, transport logistics to digital fabrication, HCI and product design are being thrown into a world where computers are active: they have physical effects on the world; and autonomous: they appear less as tools and more like actors. Furthermore, the effects may be non-local in space and time, breaking the illusion of direct manipulation or instrumental interaction.

Elsewhere, we have discussed three properties of physical things:

- *Directness of effort* – Small effort produces small effects, large effort produces large effects.
- *Locality of effect* – The effects of actions occur where and when you physically initiate the action.
- *Visibility of state* – Physical objects have complex shape and texture, but this is largely static.

Computation systematically breaks these: allowing small effort to achieve large changes, enabling non-local effects, and hiding complex state in a few square millimetres of silicon. The direct manipulation paradigm effectively fights these, attempting to create a simulation of a semi-physical world. This certainly helps us to recruit our natural human skills at dealing with the physical world, but at the cost of either ignoring the benefits of computation, or making them difficult to access (hence the developer's love of the Apple Mac's Unix terminal).

Third wave HCI encouraged us to take the social and political context of interaction seriously. The AI-fuelled growth of active physical computation means that we not only have to take seriously the physical context but also computation itself.

Many in HCI have taken up this challenge. Some have worked to allow direct manipulation to continue through digital/material boundaries, for example Constructable allowing direct 'drawing' on material with a laser cutter (Mueller et al. 2012), or MixFab combining additive and subtractive 3D fabrication (Weichel et al. 2014). However, even here the computational role is not fully hidden, for example, the creation of supports for 3D printing, or allowing 'undo' of laser cuts through rapid re-creation.

In other areas, such as autonomous vehicles, the paradigms have had to change much more, with a sense of mediated interaction: asking or telling an active, sometimes embodied, partner to perform the task rather than directly doing it oneself. For those who have been in the discipline for some years there is a sense of partial déjà vu of the command line vs direct manipulation discussions of the 1980s, but with potentially even less predictable outcomes of one's 'commands'.

Indeed unintended outcomes have become one of the hallmarks of many of the new wave of intelligent systems whether largely in the digital domain (e.g. privacy and bias in algorithms, see Dix 1992), in the physical domain (e.g. vehicles that obey the rules of the road too well), or at the boundaries (e.g. denial of service attacks on a building heating system). In the UK the Human-Like Computing programme is addressing these issues of comprehensible and explainable intelligence, and there are similar initiatives elsewhere.

In general the digital effects on society may permeate far wider than the immediate effects of specific software or apps. For example, the so called 'sharing economy' (although it often seems anything but) such as Uber or Airbnb have 'disrupted' the direct industries they address, but also challenged existing frameworks and legislation for taxation, employment, urban planning, and health and safety. Arguably these wider effects are the job of other disciplines such as law or political science. However, the position of HCI at the boundary between technology and people suggests that just as aspects of ergonomics, psychology, sociology and anthropology have become accepted parts of the discipline, perhaps we need to expand our scope to rise to these challenges also. This will certainly expand further the notions of 'context' and methodology more broadly still than those of the third wave.

The combinations of digital fabrication at the material level and cloud-based service models in the digital realm offer the potential to radically reconceptualize large areas of physical production. Here it is not so much the direct impact of the digital technologies, but the digital 'eye' that can conceive of logistics without centralisation and volume production without mass production. This will undoubtedly require a rich cross-disciplinary understanding of the world that encompasses traditional HCI, industrial design, business models and legal and regulatory frameworks.

## 8.7 Physicality of Computation

The term 'physical computing' assumes that there is computing that is not physical. In fact *all* computation has physical effect and takes place in a physical form with physical constraints whether a chip embedded in a credit-card, the laptop on your desk, or a data centre in Iceland. Ultimately all computing is physical computing.

One way this is evident in interfaces is when networking is required. This may be obvious when there is an explicit act of communication, such as email, but as services become cloud-based the implications may be more wide reaching, for example, if there is no internet connectivity you won't be able to access your Google



Docs, or use Siri. In the case of devices making use of the Internet of Things this may mean you are no longer able to control your home heating or boil a kettle.

To some extent this is not a new problem, indeed the first author wrote about these issues more than 20 years ago (Dix 1995), but they have become more critical as systems that rely on networking have become more ubiquitous. This is partly a HCI/usability issue, but also a political one as access to connectivity is far more limited for those who are already socio-economically disadvantaged (RSE 2013) – in other words the digital divide deepens the social divide (Morgan et al. 2014). Again this suggests a widening of the scope of our conception of HCI. However, even with a narrow scope we need to understand that the design choices in an interface, in particular its resilience to poor connectivity, have major social ramifications. The accessibility community has shown that progress can be made for physical, perceptual and cognitive differences; a similar effort is needed for social disadvantage.

From the hearts of black holes to the end of the universe, physicists often revert to entropy and energy to take a ‘big picture’ view of systems. This is no less true of computation.

In 2015, a report commissioned by the Semiconductor Industry Association and Semiconductor Research Corporation projected that by 2040 the envisaged level of computing would require more power than the total global energy capacity (SIA and SRC 2015). In 2017, it was Bitcoin in the news when Digiconomist’s ‘Bitcoin Energy Consumption Index’ (Digiconomist 2017) showed that Bitcoin was consuming more energy than the majority of the world’s countries, and that the energy needed for each transaction could power a US household for 9 days (Bridge 2017). Furthermore, e-waste is “the fastest growing part of the world’s domestic waste stream”, currently running at about 20 kg per person per year in the US (Leahy 2017; Baldé et al. 2017).

The calculations for Bitcoin are disputed (Schroeder 2017), and improvements to Bitcoin’s algorithms will in the long run reduce resource usage (Ethereum only requires about 1/6 of the energy although still thousands of times more than a conventional monetary transfer system such as Visa). Similarly the semiconductor industry report used the energy figures as part of an argument to fund research to reduce the energy consumption of chips. So, the picture may be a little less apocalyptic than first appears; however, even with the best outcomes, it is clear that, just as with the oil economy that preceded it, the digital economy is not costless in terms of either money, energy or environmental impact.

## 8.8 Theoretical Implications

### 8.8.1 Recap

We will recap some of the theoretical lessons we have learnt in this chapter.

First, we saw how the central mode of interaction in HCI research and design has shifted back and forth between interpersonal and instrumental interaction. The early

command line interfaces were largely based on language and a sense of mediated interaction where the operator asked the computer to perform some action on its data. In direct manipulation, the paradigm was much more one of simulated physical interaction on the data, albeit via a mouse.

This has been the dominant paradigm for 30 years, with touch-based interfaces meaning now it is merely a glass screen that separates hand and data. However, the ‘AI turn’ in recent years has started to shift the emphasis back to more intra-personal modes of interaction where the computer is seen as having agency. The UK human-like computing programme and similar initiatives elsewhere are starting to take the ‘other’ seriously in human-computer interaction.

This said, in UX practice the dominant focus appears to still often be on the individual screen and between screen interactions, with popular design and prototyping tools heavily screen- rather than context-centric.


We have also seen an expansion in the scope of HCI. In the earliest days the focus could be seen as simply action: for example the choice between keyboard and mouse; however the dialogue and flow of interaction were recognised as critical from an early stage, not least in the Seeheim model (Pfaff and Hagen 1985). In 1987 Suchman’s *Plans and Situated Action* (1987) and Winograd and Flores’ *Understanding Computers and Cognition* (1986) widened the focus to the physical and social context of the interaction, introducing ethnography and effectively starting third-wave HCI (even if it was not named until 20 years later). This was significant across HCI, but particularly within CSCW, and inter alia, brought inspiration and established collaboration with anthropology and sociology to join the previous psychology-computing disciplines.

However, the increasing ubiquity of computing in the world has led to a changing level of choice from discretionary user to digitally essential citizenry. This means that our focus as a discipline needs to encompass the impacts of and design of digital technology on communities and society as a whole. Our physical focus has likewise shifted from computer on desk (1980s) to workplace (1990s) to home (2000s) and now to neighbourhood and city scale.

We also saw how HCI and product design have encountered in different orders the factors that are described by the three ‘waves’ of HCI: cognitive/pragmatic – analytic/cognitive and experiential/phenomenological.

### 8.8.2 *Waves, Phases or Complementary Approaches*

HCI has progressed from more pragmatic 1st wave, to a more analytic/metric oriented second wave and then a more phenomenological third wave. In contrast product design has traditionally been engaged in the interplay between pragmatic and phenomenological approaches, driven not least by the nature of physical products and the ways in which we engage in physical interactions (build it and feel it). However, as product design has encompassed more digital elements into the products it considers, it has adopted more analytic practices (Fig. 8.5a).



**Fig. 8.5** Arcs of methodological foci in HCI and product design: (a) Modes of study; (b) Modes of interaction

Rather than teleological waves or even fashion-like phases, perhaps it is better to see these as complementary approaches. It also is reminiscent of the craft vs. engineering vs science discussions in the early years of HCI (Long and Dowell 1989).

Until recently, the principal material of HCI has been digital. We have seen how this moved from more intrapersonal agent-mediated modes of interaction to more instrumental object manipulation. In contrast, as a physically focused discipline product design has always been based on direct manipulation (in a physical sense) of objects. However, as product design has become more digital it has had to take the agency of digital content seriously (Fig. 8.5b). If, as suggested, HCI is also needing to move back in this direction, this suggests a convergence – of course, not ‘back’ to the command line interface, but towards design philosophies that take both physical interaction and the computer as active partner seriously.

### 8.8.3 *The Nature of Human (Physical) Experience*

This does take us back squarely to the fundamental nature of human experience with physical technology.

This clearly involves *physical bodily actions on the world*, the realm of muscle, bones and physiological study. There is also evidently an *interior intellectual and emotional life*, the study of cognitive science, counselling and the Cartesian perspective. However, we have repeatedly referred to philosophical concepts of *embodiment*, which encourages us to see perception, cognition and action not as separated stages in a pipeline, but as a single process, with a rich intertwining of self and world. This is now a popular viewpoint, which owes its origins to the early twentieth century writings of Heidegger (1927) and Merleau-Ponty (1945), and in HCI was first really encountered in the concepts of distributed cognition (Hutchins 1994) and the influence of Gibson’s ecological approach to perception (1979).

To some extent these ways of looking at individual experience correspond to the three waves:

- physical body – first wave – human factors
- interior cognition – second wave – measurable behaviour, user testing, interaction design
- embodiment – third wave – phenomenology, experience design

Again it is tempting to see this as a teleological progression, but the picture is in fact a little more complicated and there are clear limits to embodiment as an explanatory framework:

First *evolutionary design* has constructed a means for conscious linear thought that sits uncomfortably on a largely parallel, subconscious ‘infrastructure’. The costs of this emphasise the importance of this form of less immediately engaged thinking, even in pre-literate and possibly pre-verbal societies.

Second *empirical evidence* points to a complex interplay between internal and external representation. Andy Clark, one of the key philosophers of embodiment, talks about the 007 principle:

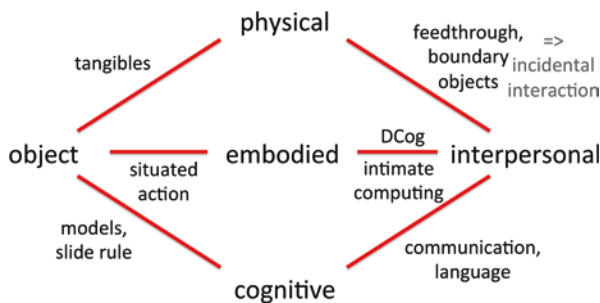
In general evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations on it as a convenient stand-in for the information-processing operations concerned (Clark 1989; Clark 1998).

However, this cuts both ways, there will be occasions when it is more efficient to consult internal models of the world rather than engage in *epistemic actions* (such as turning one’s head) to read this information from the world itself. Indeed, this is precisely what experiments have found, for example Gray and Fu (2001) found that in certain circumstances users effectively unconsciously consult their internal models when it would be more efficient and accurate to do otherwise.

Finally *externalisation* is an incredibly powerful form of expert behaviour as we saw with the reflective introspective practices of Schön’s subjects (1984). Indeed, while Heidegger wrote *about* skilful unconsidered actions, he also *thought* about this, *wrote* about it and we then *read* and *reflect* on it (and maybe puzzle over it a little).

Of course human–human connections are also a crucial part of our physical experience of the world, and are arguably the heir of being truly human. Partly this emphasises again that mediated forms of human–computer interaction are quite natural. However, more crucially it calls into attention the ways in which we engage in these human–human contacts within the physical world.

Sometimes these connections are directly physical, through joint action (lifting a heavy weight together) or in personal intimacy. In the context of physical computing there have been numerous systems which seek to use computational means, especially to enable levels of remote intimacy (Bell et al. 2003).



**Fig. 8.6** A rich world of interactions

However, often the physical world is used as a medium for communication whether through vibrating air molecules, flying photons, or ink on paper. This is the origins of language, and possibly also logic and maybe even the sense of self (Dix 2017).

### 8.8.4 Ongoing Directions

The summary of the progression of HCI in Fig. 8.5b almost suggests a regression or return to more mediated forms of interaction. However, this is not so much a regression as an opportunity for radically new forms of physical and digital interactions.

We have looked several times at the object/instrumental vs mediated/intra-personal dynamic in HCI and Fig. 8.6 summarises the different ways these interact with our physical, cognitive and embodied natures. It is perhaps easier for digital interactions to be first rooted in our cognitive and linguistic nature, hence the mediated to direct manipulation progression in HCI can be seen as a digital domain move from bottom-right to the left of Fig. 8.6.

In contrast, the AI turn, context-aware and ubiquitous computing are taking us back into more interpersonal territory; that is the top right, rather than the bottom right. Note that there is no direct communication with the computer or each other, but physical computing opens up a design space or physically mediated interaction with digital agents.

We already have understanding of digitally mediated human–human communications in this area from the CSCW literature. Although boundary objects (Star and Griesemer 1989) may be cognitive or physical, ethnographers have often found shared physical representations are crucial, and this is a key aspect emphasised by Bødker’s third wave retrospective (Bødker 2015). *Feedthrough* is also a key concept in the CSCW Framework (Dix 1994), that is the way in which one person’s actions on physical (or digital) objects are observed visually, aurally, or tangibly, and this forms a direct path of communication through the shared objects.

**Table 8.1** Material interaction

Singular	Multitude	Lumpen
sticks and stones manipulation	herds and large building command and control	fluids, earth constrain and channel
high school physics	complexity	summary statistics flows and fields
embodied action	?	JCB, slide rule
direct manipulation	programs, (TOM) Simula (OO)	Mathlab

These human–physical–human interactions can form the start of a potential patterning space for future human–physical–computer interactions.

However, there are also key differences: many intelligent interfaces are effectively disembodied, and in incidental interaction, subtle computer interventions based on unintentional user actions may give a sense of a ghost in the walls, odd changes in the physical environment, which are helpful, but hard to comprehend.

As well as a change in the ‘mediated’ end of the modes of interaction (Fig. 8.5b), there is also a subtle change in the ‘object’ end. Direct manipulation interfaces and also many of the artefacts created by product design are used largely in the singular: you drag a single file, or lift a single litterbin. Occasionally one may lift a small pile of objects, or digitally select several items, but the number is typically small.

In the physical world we also deal with lumpen materials: such as soil or water where the individual particles are so numerous we effectively treat them as a single malleable whole (Table 8.1). Here we may use tools to shape, mould or constrain the lumpen material. Intellectually we reason about the lumpen as a whole using summary statistics (the total weight, average flow speed) and maybe use mathematical tools to model the flows and fields.

Between these extremes are large collections of objects where the separate objects are still important, but are too numerous to easily deal with as an individual person. This area includes large herds of animals, crowds of people, and large construction projects. In the intellectual domain, this is the topic of complexity studies, in the physical world we typically move from individual direct manipulation to complex organisational command and control structures: such as teams of hunters and construction companies. That is, as we seek to manipulate the multitude physical actions on objects naturally require interpersonal collaboration.

In computing, the multitude is the focus of programming, in the terms of an early analogy telling the ‘totally obedient moron’ (TOM) what to do. The origins of object-oriented programming come also from this area in Simula – a language designed for simulating ensembles of physical things.

As we consider big data and physical computing, we may simply look to physical embodiments of agents, as found in household controllers and say, as in Star Trek, “computer, please turn off all the radiators in the building”. Alternatively we may seek to use forms of tangible interaction but where the physical objects are not tokens of individual digital objects, but more physical manifestations of workflows or other forms of programming-like constructs (Wang et al. 2012; Turchi et al. 2015).

However, perhaps there is something that bridges these, where the physical tokens are not simply an embodiment of instructions or code, but boundary objects in a collaboration with active computation.

## References

- Avrahami D, Hudson S (2002) Forming interactivity: a tool for rapid prototyping of physical interactive products. In: Ballagas et al (eds) Proceedings of DIS'02. ACM, New York, pp 141–146
- Baldé CP, Forti V, Gray, V et al (2017) The global e-waste monitor. <http://collections.unu.edu/view/UNU:6341>. Accessed 2 Jan 2018
- BBC (2016) BBC micro: bit launches to a generation of UK students. <http://www.bbc.co.uk/mediacentre/latestnews/2016/bbc-micro-bit-schools-launch>. Accessed 13 Dec 2017
- Bell G, Brooke T, Churchill E et al (2003) Intimate ubiquitous computing. In: Proceedings of the International Conference on Ubiquitous Computing (UbiComp'03). ACM, New York, pp 3–6
- Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. *Interactions* 22(5):24–31. <https://doi.org/10.1145/2804405>
- Bridge M (2017) Bitcoin network 'is using more energy than world can sustain'. *The Times* <https://www.thetimes.co.uk/article/bitcoin-network-is-using-more-energy-than-world-can-sustain-sqks6rpk0>. Accessed 6 Dec 2017
- Clark A (1989) *Microcognition: philosophy, cognitive science and parallel processing*. MIT Press, Cambridge, MA
- Clark A (1998) *Being there: putting brain, body and the world together again*. MIT Press, Cambridge, MA
- Culverhouse I, Gill S (2009) Bringing concepts to life: introducing a rapid interactive sketch modelling toolkit for industrial designers. In: Proceedings of the 3rd international conference on tangible and embedded interaction, pp 359–362
- Digiconomist (2017) Bitcoin energy consumption index. <https://digiconomist.net/bitcoin-energy-consumption>. Accessed 14 Dec 2017
- Dix A (1992) Human issues in the use of pattern recognition techniques. In: Beale R, Finlay J (eds) *Neural networks and pattern recognition in human computer interaction*. Ellis Horwood, Amsterdam, pp 429–451
- Dix A (1994) Computer-supported cooperative work – a framework. In: Rosenburg D, Hutchison C (eds) *Design issues in CSCW*. Springer, Berlin, pp 23–37
- Dix A (1995) Cooperation without (reliable) communication: interfaces for mobile applications. *Distrib Syst Eng* 2(3):171–181
- Dix A (2007) Designing for appropriation. In: Proceedings of the 21st British HCI group annual conference on people and computers: HCI...but not as we know it – volume 2. BCS Learning & Development Ltd, Swindon, pp 27–30
- Dix A (2010) Struggling with Heidegger. <http://alandix.com/blog/2010/08/12/struggling-with-heidegger/>. Accessed 13 Dec 2017
- Dix A (2017) I in an other's eye. *AI & Society* (in press). <https://doi.org/10.1007/s00146-017-0694-7>
- Dix A, Ghazali M, Gill S et al (2009) Physigrams: modelling devices for natural interaction. In: *Formal aspects of computing*. Springer, 21(6): 613–641
- Dix A, Gill S, Ramduny-Ellis D et al (2010) Design and physicality – towards an understanding of physicality in design and use. In: *Designing for the 21st century: interdisciplinary methods and findings*. Gower, London, pp 172–189
- Dix A, Gill S, Ramduny-Ellis D et al (2018) *TouchIT*. Oxford University Press, Oxford. <http://physicality.org/TouchIT/>

- Fitzmaurice GW, Ishii H, Buxton WAS (1995) Bricks: laying the foundations for graspable user interfaces. In: Proceedings of the SIGCHI conference on human factors in computing systems (CHI '95). ACM, New York, pp 442–449. <https://doi.org/10.1145/223904.223964>
- Garcia C (2012) The BBC micro. <http://www.computerhistory.org/atcm/the-bbc-micro/>. Accessed 13 Dec 2017
- Gibson J (1979) The ecological approach to visual perception. Houghton Mifflin Company, Boston
- Gill S, Walker D, Loudon G et al (2008) Rapid development of tangible interactive appliances: achieving the fidelity/time balance. In: Hornecker E et al (eds) Special issue on tangible and embedded interaction. *Int J Arts Technol* 1(3/4): 309–331
- Gray W, Fu W (2001) Ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. In: Proceedings of CHI '01. ACM, Seattle, pp 112–119
- Hare J, Gill S, Loudon G et al (2014) Active and passive physicality: making the most of low-fidelity physical interactive prototypes. *Int J Des Res*. Inderscience Publishing Ltd
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: Alt. Chi. Session at the SIGCHI Conference on Human Factors in Computing Systems. ACM, San Jose, pp 1–18
- Hartmann B, Klemmer SR, Bernstein M et al (2006) Reflective physical prototyping through integrated design, test, and analysis. In: Proceedings of UIST. ACM, New York, pp 299–308
- Heidegger M (1927) *Sein und Zeit*. (English translation: *Being and Time*. Harper, 2008)
- Hutchins E (1994) *Cognition in the wild*. MIT Press, Cambridge, MA
- Hutchinson H, Mackay W, Westerlund B et al (2003) Technology probes: inspiring design for and with families. In: Proceedings of the SIGCHI conference on human factors in computing systems (CHI '03). ACM, New York, pp 17–24. <https://doi.org/10.1145/642611.642616>
- Igoe T (2004) What is physical computing?. <http://www.tigoe.net/blog/what-is-physical-computing/>. Accessed 7 Dec 2017
- Ishii H, Ullmer B (1997) Tangible bits: towards seamless interfaces between people, bits and atoms. In: Proceedings of the ACM SIGCHI conference on human factors in computing systems. ACM, Atlanta, pp 234–241. <https://doi.org/10.1145/258549.258715>
- Kushner D (2011) The making of Arduino. In: *IEEE Spectrum*. <https://spectrum.ieee.org/geek-life/hands-on/the-making-of-arduino>. Accessed 7 Dec 2017
- Leahy S (2017) Each U.S. family trashes 400 iPhones' worth of e-waste a year. *National Geographic*. <https://news.nationalgeographic.com/2017/12/e-waste-monitor-report-glut/>. Accessed 13 Dec 2017
- Long J, Dowell J (1989) Conceptions of the discipline of HCI: craft, applied science, and engineering. In: Sutcliffe V, Macaulay L (eds) Proceedings of the fifth conference of the british computer society, human-computer interaction specialist group on people and computers. Cambridge University Press, Cambridge, UK, pp 9–32
- Merleau-Ponty M (1945) *Phénoménologie de la Perception* (English translation: *Phenomenology of Perception*, Routledge, 1958)
- Morgan A et al (2014) Blue sky thinking meets green field usability: can mobile internet software engineering bridge the rural divide? *Local Econ* 29(6–7):750–761. <https://doi.org/10.1177/0269094214548399>
- Mueller S, Dix A, Phillips M et al (2012) Interactive construction: interactive fabrication of functional mechanical devices. In: Proceedings of UIST '12. ACM, New York, pp 599–606. <https://doi.org/10.1145/2380116.2380191>
- Norman D, Draper S (1986) *User centered system design: new perspectives on human-computer interaction*. Lawrence Erlbaum Associates Inc., Hillsdale
- O'Sullivan D, Igoe T (2004) *Physical computing: sensing and controlling the physical world with computers*, 1st edn. Thomson, Boston. ISBN:159200346X
- Pfaff G, ten Hagen PJW (1985) *Seeheim workshop on user Interface management systems*. Springer, Berlin
- Policy Lab (2014) About us. <https://openpolicy.blog.gov.uk/about/>. Accessed 7 Dec 2017
- Raspberry Pi Foundation (n.d.) About us. <https://www.raspberrypi.org/about/>. Accessed 7 Dec 2017



- Rauterberg M, Fjeld M, Krueger H et al (1998) BUILD-IT: a planning tool for construction and design. In: CHI 98 conference summary on human factors in computing systems (CHI '98). ACM, New York, pp 177–178. <https://doi.org/10.1145/286498.286657>
- RSE (2013) Spreading the benefits of digital participation: an interim report for consultation. Edinburgh, Royal Society of Edinburgh, pp 22–24
- Schön D (1984) *The reflective practitioner*. Basic Books, London
- Schroeder S (2017) Bitcoin consumes a ton of energy, but it's not as bad as you've heard. Mashable UK. <http://mashable.com/2017/12/01/bitcoin-energy/>. Accessed 14 Dec 2017
- Shneiderman B (1983) Direct manipulation: a step beyond programming languages. *IEEE Comput* 16(8):57–69
- SIA, SRC (2015) Rebooting the IT revolution: a call to action. Semiconductor Industry Association and Semiconductor Research Corporation. <http://www.semiconductors.org/clientuploads/Resources/RITR%20WEB%20version%20FINAL.pdf>. Accessed 7 Dec 2017
- Smith DC, Irby C, Kimball R et al (1982) Designing the star user interface. *BYTE* 7(4):242–282
- Star SL, Griesemer JR (1989) Institutional ecology, translations and boundary objects: amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–1939. *Soc Stud Sci* 19(3):387–420
- Suchman L (1987) *Plans and situated actions: the problem of human-machine interaction*. Cambridge University Press, Cambridge
- Sutherland I (1963) SketchPad: a man-machine graphical communication system. In: AFIPS conference proceedings, 1963, p 323–328
- Turchi T, Malizia A, Dix A (2015) Fostering the adoption of pervasive displays in public spaces using tangible end-user programming. In: IEEE symposium on visual languages and human-centric computing, Atlanta, Georgia, US, p 18–22
- Villar N, Gellerson H (2007) A malleable control structure for softwired user interfaces. In: Proceedings of tangible and embedded interaction conference (TEI '07). ACM, New York, pp 49–56
- Wang D, Zhang Y, Gu T et al (2012) E-Block: a tangible programming tool for children. In: Adjunct proceedings of the 25th annual ACM symposium on user interface software and technology, 2012
- Weichel C, Lau M, Kim D et al (2014) MixFab: a mixed-reality environment for personal fabrication. In: Proceedings of the SIGCHI conference on human factors in computing systems (CHI '14). ACM, New York, pp 3855–3864. <https://doi.org/10.1145/2556288.2557090>
- Weiser M (1991) The computer for the 21st century. *Sci Am* 265(3):94–104
- Winograd T, Flores F (1986) *Understanding computers and cognition: a new foundation for design*. Addison-Wesley, New York
- Zampelis D, Loudon G, Gill S et al (2012) IRIS: augmented reality in rapid prototyping. In: Proceedings of the 13th rapid design, prototyping & manufacturing conference. Lancaster University, Lancaster

# Chapter 9

## Third-Wave HCI Perspectives on the Internet of Things



Tom Jenkins

**Abstract** This chapter draws out three threads from contemporary HCI, design research, and design practice to consider the Internet of Things from a third wave HCI perspective. These stances towards the IoT emphasize an agentic thing in itself as component of a broader system; the network that a system resides inside of as members of a broader social or cultural context; and the role of a system to articulate and maintain sites of contestation around public issues. These perspectives both build on and react to a set of categories that describe how the existing Internet of Things has been approached from a second-wave HCI orientation, as primarily providing opportunities for command and control, making technologies more efficient, and consuming products and services. This third-wave perspective hopes to broaden the conversation around the potentials for networked technologies that operate inside of rich cultural and social context.

### 9.1 Introduction

“The Internet of Things” (IoT) describes a trend advocating that all sorts of physical artifacts become connected to and controllable from the Internet. According to it, a coffeepot might be controlled alongside a thermostat to have a home warm and the coffee on when a person wakes up in the morning; or sensors in the basement might email you if your basement is flooding. While this all seems very sanguine, current IoT technologies rely on centralized servers, well-defined APIs, and black-boxed electronics for the end-user, and are built only to be used in specific, condoned ways. While future visions of computationally-augmented domestic life have been present for decades (cf. the “kitchen of tomorrow”, smart homes, and the field of ubiquitous computing), only very recently mainstream technology companies

---

T. Jenkins (✉)  
Georgia Institute of Technology, Atlanta, GA, USA  
e-mail: [tom.jenkins@gatech.edu](mailto:tom.jenkins@gatech.edu)

have begun to manufacture and market smart devices to be used in real homes by real people, instead as technical demonstrations used for marketing purposes or as standalone design concepts.

In some ways, an Internet of Things has been with us for a long time. In addition to end-user devices and products, there are things like algorithmic traders—programs that make thousands of stock transactions per second to obtain a market edge by any means possible. These are things on the Internet, but they are used solely for pragmatic purposes and for financial gain. Likewise, massive and often overlooked examples of the IoT take place in shipment tracking, with pallets of merchandise being tracked by RFID tag from one side of the Earth to the other. The vision of everyday objects that are each individually connected to and addressable from the Internet provides a colossal opportunity to produce new devices replacing almost everything in a home. The utopian vision from the consumer's perspective is one where the world is a seamless collection of devices that together “just work”—Networked devices that operate in concert with finely-tuned algorithms and descriptive big data to predict needs before you know you have them. From the corporate perspective, the story of IoT is just as rosy: access to fine-grained information about habits in the home, household purchases and their frequency, social graph information, and more, all baked into hardware that needs occasional replacement and is locked into proprietary software. As it's commonly used, however, the term is relatively torpid: it refers mainly to marketing materials that describe how devices and systems in the future will become smarter, more convenient, and more responsive to human needs. This IoT is a set of rhetorical devices that speak to the promises of the near future. Implicitly, the Internet of Things exists as a site for industrial opportunity, a way of selling individually Internet-addressable objects *en masse* and in concert to homes around the world—or at least selling the vision of those things.

The emerging Internet of Things is of natural interest to HCI researchers. In many ways, the IoT is a fulfillment of the promise of truly ubiquitous computing as advanced by Mark Weiser (1999). Beyond the interactive tabs, pads, and screens of vision of seamless and calm computing for the end user, the Internet of Things begins to articulate the promise that computational technologies can find residence in all kinds of everyday object, opening the door to a world of interactive objects that mediate and support human experience in all sorts of ways. With respect to both timing and its nature, the dawn of the Internet of Things aligns itself naturally to the third wave of HCI. Historically, there have been three waves (Bødker 2006) or paradigms (Harrison et al. 2007) to HCI. Broadly, these waves correspond to different questions and modes of inquiry that each foreground. The first wave is purely technologically focused, and emphasizes a cognitive science approach to solve issues that together have been called “human factors.” Here, understanding interaction is related to the most fundamental use of input devices and the efficacy of performing simple tasks. This first wave generally uses systems building as a methodology, and creates knowledge about how people use systems formally through user testing. The second wave extends the first wave through a broader frame of sociability. In the second wave, the focus of attention shifted from individual users towards groups of users collaborating across applications within communities of practice (Bødker

2006). In this paradigm, what counts as interaction changes from being direct manipulation of input to understanding interaction as communicating information within this group. Second-wave HCI concerns itself with questions regarding efficiency, and producing design strategies for better software, with a goal towards optimization that can be applied regardless of context. Finally, the third wave or paradigm of HCI takes interaction as situated as part of a broader cultural context. Instead of being clustered tightly inside the workplace, computing becomes intertwined more closely with our homes, everyday lives, and overall culture (Bødker 2006; Harrison et al. 2007). Likewise, the kinds of questions being asked and examined became similarly open-ended. Third wave HCI research asks questions about the experiences that interactions can engender and support, and how cultural meaning can emerge through design and use of computing technologies.

One way to look at the Internet of Things is as a fundamentally third wave concept that is understood—so far—from a perspective that is firmly rooted in second wave HCI. While it is true that the Internet of Things extends far beyond the workplace—from homes, to warehouses, ships, and automobiles—the way that it is presented and discussed remains fundamentally service-oriented. Whitepapers describing the business potential of the Internet of Things are primarily understanding the nature of this field as being Business to Business (B2B) or Machine to Machine (M2M), and note that the major task for researchers and designers is to understand how to leverage the coming explosion of IoT-generated data.<sup>1</sup> At the other, more consumer-oriented end of the spectrum, the systems and devices offered as members of this emerging revolution in computing technology fit generally into three categories: devices for security, monitoring, and control; systems that promote efficiency; and those built primarily for entertainment and consumption.

The first and largest category of devices are for security, monitoring, and control. One system in this category includes the Samsung SmartThings hub and its various SmartSense modules that provide different kinds of sensors to deploy in a home. These include motion sensors, moisture sensors, temperature and humidity sensors, smart power outlets, an open/closed sensor, and so on. Together, these sensors are meant to instrument the home and provide total knowledge of its condition. On the other hand, the Philips Hue is a “smart bulb” that lets a resident set exactly the color and brightness that they want for a space from their phone, or even to program different settings based on particular conditions. The Hue offers their owners fine-grained control over the feel of the home, and can offer an endpoint to visualize a host of different information sources. Products like these illustrate one of the major promises of the Internet of Things. The IoT offers a way to make individual objects addressable and controllable, while simultaneously reporting that information to a resident via the Internet.

The second category of contemporary IoT devices are designed to make everyday life more efficient. Google’s Nest Thermostat, for example (Fig. 9.1), promises to help homeowners reduce their energy use over time, saving both the planet as

---

<sup>1</sup> <https://community.arm.com/iot/b/blog/posts/internet-of-things-business-index-a-quiet-revolution-gathers-pace>



**Fig. 9.1** Google's Nest Thermostat

**Fig. 9.2** Amazon Echo



well as on their energy bills. It does this by learning resident's daily patterns and schedules over time in order to build a model of their lifestyle, and operates more efficiently by coupling heating and cooling changes to these patterns more closely than a person could do (or, perhaps, would want to). The “smartness” of the Nest and other IoT devices like it comes from sets of algorithms that operate in concert to develop rules to describe larger events in the world. As it “learns” the behavior of a home's residents, the Nest and other IoT devices in this category exemplify the promise of ever smarter algorithmic ways to make everyday life easier. This perspective, building on a similar rhetoric of more perfect knowledge of many conditions through “big data” is another way that the Internet of Things is being positioned as a way to participate in this information revolution.

Finally, the third class of devices emphasize entertainment and consumption. The Amazon Echo (Fig. 9.2) is, at its core, a computerized Bluetooth speaker coupled with a conversation-based interface that lets it both respond to and answer short verbal commands or queries. Its conversational agent, called Alexa, uses the internet

to provide weather updates, news, and music (among others), while providing hands-free kitchen timers, grocery lists, to-dos, and more. The real appeal and value to the Echo, however, is its access to the massive product and service infrastructure that Amazon operates. Amazon Prime members have access to a colossal library of music that can be played directly from the speaker, and can order products from Amazon directly using voice commands. The Amazon Dash goes even further to make ordering products easier. These are small buttons that connect to a home Wi-Fi system and have product names and logos emblazoned on them. They are placed where the product is used, and by pressing the button, replacements can be ordered directly from Amazon. The Echo functions admirably as an entertainment device, but as an endpoint for Amazon's shipping infrastructure, it makes it clear that much of the promise of the Internet of Things, at least for now, rely on creating new opportunities for traditional commerce.

Of course, these categories are not hard and fast. The Nest's broad history and data screens mean that it is also a monitoring system, even as rules set up in a SmartSense Hub might mean that it can automate everyday tasks, helping them become more efficient. At their core, the entire class of aftermarket smart home technologies may mean that they are all about entertainment and consumption as much as they are meant to solve any real problem. These categories—devices that establish monitoring, make homes more efficient via judiciously applied computer science, and together support new kinds of entertainment and consumption—make it clear that the way that IoT technologies and devices are commonly framed is decidedly retrograde. Each of these modes reflect certain assumptions around computing technology based on early-wave HCI concepts like information transmission fidelity, efficiency, and direct user input. At the same time, the interests of developers seem to primarily be concerned with collecting, managing, and monetizing user data. What kinds of perspectives could be brought to bear on the Internet of Things to focus on broader issues of context, content, and values?

Considering these perspectives leads to a fruitful reframing of what the Internet of Things could be. The rest of this chapter describes three perspectives on third-wave IoT. The first of these examines the IoT from the level of the things themselves as agentic entities. The second takes the IoT as a system of connections and networks that offer opportunities for new kinds of collaboration. Finally, the third perspective from the effects that these devices can play as members of public social life. It's important to note that these perspectives are not exclusive categories, but instead are each a different way of examining the Internet of Things that draw attention to the role that the IoT can play at different levels.

## 9.2 An Internet of Agentic Things

The first way that we can begin to unpack how the Internet of Things might operate in third-wave HCI is by renewing our focus on the things itself. For much of HCI's history, emphasis has been placed on how users approach technical challenges, or

how technical implementation might be helpful to solve problems among and between users. The shift towards a richer understanding of context that is characteristic of HCI's third wave, however, offers an opportunity to reframe how we understand interacting with technology as being more collaborative and agentic—on the part of the technology itself—than we may have otherwise imagined.

This perspective is driven theoretically by historical understandings of the role of the *thing*. If the promise of the Internet of Things is that new kinds of things can be present and active on the internet, then it makes sense to unpack things themselves. The “thing” is one way that designers and theorists have contemplated the role of objects in sociotechnical systems. Historically, this perspective on the role of the thing comes from the term *ding*, an ancient Germanic word for political gatherings. As described by Heidegger: *the Old High German word thing means a gathering, and specifically a gathering to deliberate on a matter under discussion, a contested matter. In consequence, the Old German words ‘thing’ and ‘ding’ become the names for an affair or matter of pertinence. They denote anything that in any way hears upon men, concerns them, and that accordingly is a matter for discourse.* (2009).

Understanding things as being constituted by both matter and discourse, Studio Atelier have expanded the role of design in creating *design things*, materials emerging from the process of design that enact multiple transformations (Atelier et al. 2011). Unlike design projects which have well defined boundaries, design things are messy, supporting many different values and viewpoints. Borrowing from Latour, the concept of actants working inside a network (Latour 1993), the design thing “aligns human and nonhuman resources to move the object of design forward, to support the emergence, translation, and performance” of the design object through “participation, intervention, and performance in this sociotechnical thing” (ibid).

Design things refer to the inherent ability of the object to align various interests. The designed object becomes a point of contention and contestation between many different factions during the course of the design process as well as afterwards. On the one hand, the design object is part of a lasting record of process: the designers build into the object a history of decisions and compromises that each play a role in the creation of that thing. At the same time, the designed object is still an active space for current controversy and consequently, future decision-making. The device itself also plays a role in how interaction is construed. The design choices that are embedded into a thing generates constrains and opportunities for how it should be understood and used. In this way, the design thing can be understood as having agency in the world, by both creating a site for issues to be contested as well as offering a perspective on how contestable issues should be understood. The Nest Thermostat, for example, participates in multiple discourses around climate change, energy efficiency, and home automation, and might make a claim that its material arrangement is the right one to engage with these ideas. Material agency can go deeper than considering the role that, however, and considering how objects may understand their position in the world can build productively on ideas of devices solely operating in the service of human interests.

Ian Bogost's concept of *alien phenomenology* offers another way to approach how things might understand the world on their own terms (Bogost 2012). His proj-

ect of speculative realism seeks to understand how non-humans perceive the world, and considers the role of a philosopher in articulating and considering that perception:

*The true alien recedes interminably even as it surrounds us completely. It is not hidden in the darkness of the outer cosmos or in the deep-sea shelf but in plain sight, everywhere, in everything. Mountain summits and gypsum beds, chile roasters and buckshot, microprocessors and ROM chips can no more communicate with us and one another than can [an] extraterrestrial. It is an instructive and humbling sign. Speculative realism really does require speculation: benighted meandering in an exotic world of utterly incomprehensible objects. As philosophers, our job is to amplify the black noise of objects to make the resonant frequencies of the stuffs inside them hum in credibly satisfying ways. Our job is to write the speculative fictions of their processes, of their unit operations. Our job is to get our hands dirty with grease, juice, gunpowder, and gypsum. Our job is to go where everyone has gone before, but where few have bothered to linger. (Bogost 2012)*

One project in HCI and design that has begun to speculate toward understanding and interpreting the perspective of alien things is Giaccardi et al.'s pictorial *Thing Ethnography* (Giaccardi et al. 2016a). Thing ethnography equips everyday objects like kettles and mugs with “autographers:” logging sensor systems that let everyday things capture and convey the social practices that they are members of as well as the patterns of use that they take part in Fig. 9.3. This project meant to understand the evolving use and applications of things in the world, what has been called “design after design” (Ehn 2008). Taken as a whole, this set of data becomes what Giaccardi et al. call *data worlds*. The autographers record data from accelerometers, color sensors, magnetometer, thermometer, and proximity sensors, and are not triggered by the researcher exploring practices in the home, but instead by the devices themselves. The autographers are meant to access the things’ perspective. “Autographer data illuminated unexpected and otherwise invisible relationships among objects—a perspective that we could not elicit through observation and interviews alone” (Giaccardi et al. 2016b).

Through thing ethnography, devices are able to offer their perspectives on our world, and contribute a new, relevant viewpoint to the design of interactive systems that need not be anthropocentric. This work attempts to reimagine the ontological relationship between people and things, and provides a method to put each actor on a more level playing field. Images and other sensor become the substance to transfer information between these incompatible phenomenologies, and slot in to existing design strategies in ways that are more intuitive than raw data. Understanding things as agentic devices that may have their own wants and needs provides a new light on how the Internet of Things might operate. Taking these things as representative of other, more alien perspectives and ways of knowing might take place in networked futures. This turn is especially useful to help defamiliarize what visions of the connected future might be like (Bell et al. 2005).

The Internet of Things offers a means to examine how cumulative computational relationships between things operate to produce social effects. Taking the devices, platforms, and systems of the IoT as agentic broadens what it may mean to create IoT for HCI designers. Considering the roles of devices in themselves, as agentic






**Fig. 9.3** Images from the thing ethnographer's perspective

things with their own phenomenological understandings and relationships to both humans and the world at large offers a richer design space to think of emerging landscapes of people and things.

### 9.3 An Internet of Social Things

The second way we can understand the Internet of Things from a third-wave perspective is through the relationships that these networked things create and sustain. Ultimately, the fundamental conceit of the Internet of Things is that by granting



**Fig. 9.4** Morse Things

access to the internet for everyday devices and systems, things are made to be more than they were—the network elevates them. Ultimately, this means that one of the aims in building relationships among and between agentic things is to consider the connections themselves, and what kinds of values and practices they might support in the future.

One example of a network-driven research project that uses networks as a design material are *Morse Things* (Wakkary et al. 2017). Morse Things are sets of ceramic flatware—a large bowl, small bowl, and cup, marked red, yellow, and blue—that communicate with one another over a home’s Wi-Fi network (Fig. 9.4). Embedded with electronics, they are simultaneously things of home and domestic life as well as network-enabled computational objects. These things use their network connection to communicate with other Morse things as well as their human roommates. Sleeping for at least 8 h a day, the things wake up at random times and both send and receive messages to each other. These messages are in Morse code and are both audible and broadcast on individual, device-specific twitter accounts.

Unlike more traditional IoT artifacts, the Morse Things are built to serve their own interests: “*Their human functionality is of an everyday nature that already exists in homes today shifting the question from what they do to how they are in our homes. The Morse Things embody the proposition that our relationships with Internet enabled things is a matter of negotiation over time rather than predefined or prescribed as a service or functionality*” (Wakkary et al. 2017). They send dots and dashes to each other to know that the other is there, using a human-created network as a substrate for their own social needs and wants. While the bowls and cups can be used to eat or drink with, as with any other bowl or cup, they are at the same time participating in a networked life that doesn’t really have anything to do with their owner. The things extend the thing ethnography of Giaccardi et al in a completely different direction. Instead of looking to build a kind of empathy with the sensed world of a thing in itself, *Morse Things* forces us to understand and respect that things have their own inner worlds or motivations that we cannot access. These things’ Twitter account reflects the intentions of the devices, but are still devoid of meaning. The messages are present: “calling anyone this is Blue” as well as a Morse equivalent, but any understanding of why the message was sent at a particular time, or for what end, remains elusive.

*Morse Things* uses a material speculation approach (Wakkary et al. 2015) to begin to consider what agentic, networked devices might be when they are more interested in serving their own needs than the needs of their human hosts. In this project, pushing the boundaries of what IoT might be is necessary. The devices need to feel familiar-but-strange, be both legible as an Internet of Things object, but also keep some sense of what they might really be like reserved. *Morse Things* is a present-day project that speculates as to what a domestic future might be like when different kinds of participants in a home come together and begin to learn how to collaborate.

*Bots*<sup>2</sup> is a design fiction project by Kevin Gaunt that considers the future of agentic objects, speculating on how AI, IoT might cohabitate in the future, and where human and object autonomy might mingle or become more complex. Bots emphasizes a future where the crux of *Morse Things* has been negotiated, and mutual intelligibility secured. This is not necessarily a utopian vision of the domestic future, however: “*What are we left to do when technology just happens to do everything for us? Future visions depict life in smart-homes where technology deals with all our chores and reads our deepest wishes before we are even aware of them. But is living in such an environment actually desirable? How might people appropriate their smart-home technology to regain control?*”

*Bots* is a system of robotic home assistants that collaborate with each other and their resident to make living alone “just a little” more interesting. There are two main ways to interact with the system. The “Brain” unit is a video display with speaker and microphone that has slots in it for up to 16 smaller Bot units. The “Sense” is a remote speaker and microphone unit that offers a way to interact with the bots while away from the panel, in another room, for example. The Bots themselves are modular artificial intelligences that each focus on a single task. For example, the website includes images of memory bot, emergency bot, bank bot, recipe bot, security bot, social bot, revenge bot, challenge bot, gym bot, surprise bot, diet bot, meditation bot, diva bot, puzzle bot, transport bot, etiquette bot, care bot, random, preference, shop, neighbor, lottery, skate, tech support, and so in, and it is not at all difficult to imagine many other kinds of single-serving bot identities that each focus on a particular aspect of life.

Each of these bots can be combined and asked to work together to introduce certain kinds of serendipity or emergent effects into home life. In the video that illustrates how this system might be used,<sup>3</sup> an older woman who lives alone is bored and asks for something to do. Surprise, shop, skate, and bank work together to purchase a skateboard and have it delivered to her home. As each bot has a single passion and personality, which bots are currently placed in the brain influences what kinds of outcomes are possible without quite being deterministic. Gaunt described it this way “*The group dynamics that arise from the bots placed in the main unit might change our relationship with the technology to something more like having a*

---

<sup>2</sup><https://www.kevingaunt.com/#bots-section>

<sup>3</sup><https://vimeo.com/149985577>

*group of pets: never entirely predictable but always succeeding or failing with the best of intentions.”*

Agentic things of the future offer one vision of how to approach a social networked effect of the Internet of Things. If we remain in the present, we can also think of the networks themselves as a material for understanding how the Internet of Things makes social and cultural values manifest. What kind of networks matter to the Internet of Things? Networks themselves can be considered as a kind of design resource that illustrates how there is more than one kind of connectivity.

*Foxes Like Beacons* is an exploratory project using open data of public radio stations with inexpensive, low-power signal detection in order to create an open location-positioning system.<sup>4</sup> In this project, the material of the network itself is under scrutiny. Here, the designers seek to understand both what it means to understand finding location and what kinds of resources might be available in that project. Even further, the global standard for positioning—GPS, or global positioning system—is a black box tightly controlled by governments, militaries, and similar-scale actors. *“Since GPS ubiquitously affect our interactions and experiences with our environment, economy and privacy, Foxes Like Beacons questions this present model, thus opening up space for speculations about alternative navigation systems and new models for interaction.”*

There are three devices in *Foxes Like Beacons*. Device No.1 measures the signal strengths of radio towers and decodes their station. Onboard data about transceiver locations and transmission power on the device is then used to calculate the distance to each station. That data is used to triangulate the position of the user, finally providing their geolocation. Adding more stations results in higher resolution and more accurate geolocations. Device No.2 enables users to navigate space using sound instead of geographic features, as users following “geo-acoustic” maps. These maps emerge as the device automatically adjusts the volumes of simultaneously played radio stations according to their proximity. Device No.3 speculates about further applications by analyzing and interpreting signal modulations which occur due to factors like electromagnetic radiation, weather, and geographic conditions. These modulate radio signals in characteristic ways which can be interpreted and incorporated in navigation or exploration.

*Foxes Like Beacons* is part of an ongoing project called “Stupid, Messy Networks”. The larger project investigates the process of digital networks becoming ubiquitous infrastructure, and moreover, how these new infrastructures empower or constrain various members of those networks: *Digital networks are based on protocols which can be seen as sets of technological standards, building the overall network-architecture. Just like the architecture of a house constrains our physical movements walls, windows, doors and locks, network architectures can constrain us in everything we do within them by certain protocols and technical standards. As we increasingly transfer critical infrastructures, like communication, identification, currencies or even education to networks like these, it is only appropriate to ask who the architects of these networks are. What are their individual interests? How*

---

<sup>4</sup><http://stupidmessy.net>

*might they implement their own interests in network architectures or even utilize network architectures to achieve idiosyncratic goals?*

In exploring questions like these, *Stupid Messy Networks* becomes a tool to find out insights from technology critique, in order to design and develop network architectures and technologies that are of public interest. At first blush, the projects in *Foxes Like Beacons* may not be legible as members of the Internet of Things. But at their core, these are projects that are in some ways prototypical IoT. They are devices that rely on networked, technological infrastructures even as they probe and push on what the material capabilities of those networks might be. These are far removed from devices that make everyday life more efficient or offer new opportunities to consume, however. These are devices that engage with the ideas of the internet of Things in order to push on the edges of what might be out there: novel networks, novel contexts, and site-specific exploration of data worlds. In these projects, networks are not always well-controlled spaces for sending and transmitting cultured, well-formed packets to friendly ports. They are that, to be sure, but the projects in this section reveal that networks are also wild, untamed spaces that leave themselves open to opportunities for feral objects or parasitic uses.

## 9.4 The Internet of Things in Action

One final way that the internet of things might be understood, and building on ideas put forward in the *Stupid, Messy Networks* project, is to consider the purpose and outcome of networked systems in practice. What kinds of values and practices are being constructed into the IoT? Frequently, home IoT platforms are hermetic, relating primarily to practices within the home and relaying information about them to a homeowner or resident. The effects of networked technologies at the end-user level should not be limited to the just domestic or industrial spheres. The Internet of Things can also play a role in public and social life, and become part of political issues and broader decision-making processes. While the social sphere from the last section is essential in considering how IoT devices and platforms can produce public effects, this builds most concretely from the definition of “thing” as a site for issues and controversies that take place from the first section.

That understanding of things—as a shift from an object in space to a site for contested political action—is expanded by Bruno Latour in *From Realpolitik to Dingpolitik* (Latour 2005). This is Latour’s proposal of thing-centered or “object oriented” democracy. Objects become battlegrounds for differing perspectives on matters of concern: “*every one of these objects, you see spewing out of them a different set of passions, indignations, opinions, as well as a different set of interested parties and different ways of carrying out their partial resolution.*” These contested objects are things in themselves—in that they are materialized issues that simultaneously embody multiple political perspectives—but also provide another kind of assembly for those who have a vested interest in that thing and those issues. One way to approach design as performing this kind of assembly is through the lens of

public design. Public design is design work that articulates particular issues in order to provide infrastructure for collective debate or action. The goal of public design work is to make the factors which inform and constitute an issue and the consequences of that issue able to be experienced directly (DiSalvo et al. 2014). Public design objects express matters of concern by articulating an issue and giving form to problematic situations, and following Latour, provide a venue for that debate to become tangible. At the core of public design is the idea that design can organize groups of people invested in an issue—a public (after Dewey 1954)—in concrete, material ways.

This concept of designing for publics can be extended to the Internet of Things (Jenkins 2015). What might a “Public IoT” be like? Here, the term Public IoT serves as shorthand for an assemblage of technologies and services that comprise the Internet of Things, and as a description of the context and purpose of that assemblage (DiSalvo and Jenkins 2017). More specifically, what is of interest to us is how IoT technologies and services work to both serve public interest and contribute to the formation of publics. How things behave and affect public and political life is becoming increasingly important as networked devices become more and more ubiquitous. What role does networked devices play in political issues and conversations? What might the role of networked things be in articulating and responding to public concerns?

The Public Design Workshop at Georgia Tech has been designing and prototyping public IoT systems in response to questions like these, primarily to support urban foraging groups in Atlanta, Georgia. Atlanta has an unusually dense canopy cover for a large city, and many of the trees produce fruit that, left unattended, would fall to the ground and rot. Concrete Jungle is a non-profit organization that finds these fallow, untended trees in the city, organizes volunteer-staffed picks when the fruit is ripe, and collects and distributes the fruit to food banks and homeless shelters throughout Atlanta.

Keeping track of a city’s worth of fruit-bearing trees is not a simple task. Concrete Jungle depends on volunteer labor for most of what it does in the city. Volunteers are organized as a group to harvest apples, pears, plums, and peaches from trees, or service berries from a bush. Larger trees require more volunteers and more time. On top of the physical labor of organizing picks, there is the work of monitoring when trees should be picked. Knowing what trees are ripe is an essential part of being able to mobilize volunteers without wasting effort. The group has a detailed map of fruiting plants that they use to keep track of approximate ripening based on previous years’ data, but even knowing timing to a set of weeks requires large amounts of manual attention, and organizers or other volunteers may spend hours each week cycling through the city to get a sense of what trees are ripening when.

In order to help them orient their limited volunteer resources towards the more manually labor-intensive fruit picks, our studio worked with Concrete Jungle to build technologies that could monitor ripeness remotely, automatically, and ideally, at scale. We explored two design avenues. The first was using drones to detect ripeness using a camera, while the second used a sensor platform mounted to a tree.



**Fig. 9.5** Foraging drone, fruit-vision

*Drones for Foraging* (Fig. 9.5) is a research through design project that explores the use of hobbyist drones in support of urban foraging by developing use-cases that make strong claims around the current and future uses of drone technologies, including bottom-up food finding, prototyping software and user interfaces for drone navigation, helping build a broader community of practice around these complex issues, and exploring the use of open-source technologies for image capture and analysis (DiSalvo and Jenkins 2015). The project imagines, investigates, and prototypes uses of drones for foraging as an alternative agricultural practice, and documents and shares that research-through-design so that it is available for others to build upon.

*Drones for Foraging* suggests how collectives of people and objects come together in inventive ways to address an issue and its conditions: helping to foster and maintain resilient food systems. Through design, this issue and conditions are then brought together with a technological platform and its affordances: hobbyist drone technologies that provide semi-autonomous remote media capabilities. The purpose is not to make the use of drone foraging possible, but rather to articulate the potential, and to have that articulation serve as the basis for a considered reflection. As a design thing, this reflection includes the proposed product or service, and it extends to also include a consideration of the social and political context in which this product or service might exist in the future (Halse et al. 2010; DiSalvo 2016).

*Fruit Are Heavy* (Fig. 9.6) is a sensor system that measures the bend of fruit tree branches as a proxy for tracking the ripeness of the fruit on those branches, in order to help Concrete Jungle (DiSalvo and Jenkins 2017). The system is comprised of three components. First is the platform hardware and software, which utilizes a



**Fig. 9.6** Fruit sensor, sensor installed on tree, being taken down

custom board and software to monitor the sensor. The second component are the sensors themselves, one of either a bend sensor or stretch sensor. The third component is the housing sensor platform's housing, which both provides protection from the elements as well as attaches the platform to a tree. It has backpack-like form with clips that allow for a range of adjustments to accommodate a variety of thicknesses of tree trunks and branches. The two straps wrap around a sturdy branch, tightening around both the sensor box and the tree. The sensor platform needs to be deployed for a month in order to detect the progressing bend in the branch corresponding to the fruit's growth.

The *Fruit Are Heavy* project allows us to explore one way of using IoT technologies to support a collective and public endeavor. Working together with community partners, we were able to design and prototype a system that serves a public interest and contributes to the work of a public in addressing its concerns. In this case, this happened to be a public of urban foragers broadly concerned with providing food for those in need, and more narrowly concerned with mitigating the work of that endeavor. Regarding the work of fruit monitoring, the prototype platform suggests that a finished platform based on similar engineering and design principles could do the monitoring work necessary to alleviate the need for foragers to check fruit trees regularly. This could reduce their total labor, enabling foragers to direct those efforts elsewhere in the organization, such as increasing the number of trees that they monitor—thereby increasing the amount of fruit donated—or simply put, taking less time and effort to achieve the main goal of providing for others.

Foraging exemplifies the relationship between economies and civics. Providing food to the needy, and more broadly, contributing to the resiliency of the local food system to improve food security, is a civic endeavor. Foraging is an alternative way to do this—it operates in parallel to existing systems of food provisioning, intersecting at some points (such as distribution channels), while diverging at other points (such as how the fruit is collected). Foraging, then, is suggestive of a new



kind of civics that operates in conjunction with governmental and non-governmental organizations and services, but is not in any official manner part of those systems (DiSalvo et al. 2016).

The drones and sensors for foraging show another side of what the Internet of Things could be. They are semi-autonomous agents that are demonstrating the capability of design to engage with issues of civics and infrastructure. While one reading of both the drones' activity for finding ripened fruit or the sensor boxes reporting when fruit on a limb is heavy enough to be picked reflect the second-wave IoT perspectives of command and control or simply making existing foraging processes more efficient, there is also something else happening here. These technologies are making social and public effects manifest, and the devices themselves are articulating issues that are important to particular groups and communities.

## 9.5 Conclusion

This chapter draws out three threads from contemporary HCI, design research, and design practice to consider the Internet of Things from a third wave perspective. These stances toward the IoT emphasize an agentic thing in itself as component of a broader system; the network that a system resides inside of as members of a broader social or cultural context; and the role of a system to articulate and maintain sites of contestation around public issues. These perspectives both build on and react to a set of categories that describe how existing Internet of Things platforms have been approached from a second-wave HCI orientation. Instead of understanding the IoT as providing opportunities for command and control, making technologies more efficient, and consuming products and services, this third-wave perspective hopes to broaden the conversation around the potentials for networked technologies that operate inside of rich cultural and social context.

These perspectives are meant to provide a way of both critically engaging with contemporary IoT systems and trends as well as to offer techniques and angles on designing and developing new systems. From this third-wave viewpoint, the Internet of Things shifts from being a suite of consumer products, systems that exist to be used for monitoring, efficiency, and consumption into novel configurations of materials and publics. This Internet of Things proposes systems devices and platforms that creates speculative encounters between people, objects, and infrastructures.

## References

- Arm Community. <https://community.arm.com/iot/b/blog/posts/internet-of-things-business-index-a-quiet-revolution-gathers-pace>. Accessed 25 Oct 2017
- Atelier BT, de Michelis GD et al (2011) *Design things*. MIT Press, Cambridge, MA
- Bell G, Blythe M, Sengers P (2005) Making by making strange: defamiliarization and the design of domestic technologies. *ACM Trans Comput Hum Interact* 12:149–173. <https://doi.org/10.1145/1067860.1067862>

- Bødker S (2006) When second wave HCI meets third wave challenges. In: Proceedings of the 4th Nordic conference on human-computer interaction: changing roles. ACM, New York, NY, USA, pp 1–8
- Bogost I (2012) Alien phenomenology, or, what it's like to be a thing. University of Minnesota Press, Minneapolis
- Dewey J (1954) The public and its problems. Swallow Press, Athens
- DiSalvo C, Jenkins T (2015) Drones for foraging. In: 2nd Conference on research though design. Cambridge, UK
- DiSalvo C, Jenkins T (2017) Fruit are heavy: a prototype public IoT system to support urban foraging. In: Proceedings of the 2017 Conference on designing interactive systems. ACM, New York, NY, USA, pp 541–553
- DiSalvo C, Lukens J, Lodato T, Jenkins T, Kim T (2014) Making public things: how HCI design can express matters of concern. In: Proceedings of the SIGCHI conference on human factors in computing systems. CHI '14. ACM, New York. <https://doi.org/10.1145/2556288.2557359>
- DiSalvo C, Jenkins T, Lodato T (2016) Designing speculative civics. In: Proceedings of the 2016 CHI Conference on human factors in computing systems. ACM, New York, NY, USA, pp 4979–4990
- Ehn P (2008) Participation in design things. In: Proceedings of the tenth anniversary conference on participatory design 2008. Indiana University, Indianapolis, IN, USA, pp 92–101
- Giaccardi E, Cila N, Speed C, Caldwell M (2016a) Thing ethnography: doing design research with non-humans. In: Proceedings of the 2016 ACM Conference on designing interactive systems. ACM, New York, NY, USA, pp 377–387
- Giaccardi E, Speed C, Cila N, Caldwell M (2016b) Things as co-ethnographers: implications of a thing perspective for design and anthropology. In: Smith RC, Vangkilde KT, Kjaersgaard MG et al (eds) Design anthropological futures. Bloomsbury Academic, London/New York, pp 235–248
- Halse J, Brandt E, Clark B, Binder T (2010) Rehearsing the future. The Danish Design School Press, Copenhagen
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: CHI '07 Extended abstracts on human factors in computing systems. ACM Press, NY, USA
- Heidegger M (2009) The thing. In: The object reader. Routledge, London/New York
- Jenkins T (2015) Designing the “things” of the IoT. In: Proceedings of the ninth International conference on tangible, embedded, and embodied interaction ACM, New York, NY, USA, pp 449–452
- Latour B (1993) We have never been modern. Harvard University Press, Cambridge, MA
- Latour B (2005) From realpolitik to dingpolitik. In: Making Things Public atmospheres democrat, pp 14–44
- Wakkary R, Odom W, Hauser S et al (2015) Material speculation: actual artifacts for critical inquiry. Aarhus Ser Hum Centered Comput 1:12. <https://doi.org/10.7146/aaahcc.v1i1.21299>
- Wakkary R, Oogjes D, Hauser S, et al (2017) Morse things: a design inquiry into the gap between things and us. In: Proceedings of the 2017 conference on designing interactive systems. ACM, New York, NY, USA, pp 503–514
- Weiser M (1999) The computer for the 21st century. SIGMOBILE mob Comput. Commun Rev 3:3–11. <https://doi.org/10.1145/329124.329126>

# Chapter 10

## Inclusion in the Third Wave: Access to Experience



Christopher Power, Paul Cairns, and Mark Barlet

**Abstract** In this chapter, we examine inclusive design of technology for people with disabilities in the context of the Third Wave HCI. As technology becomes more integrated into our lives beyond work, there are increasing opportunities for people with disabilities to have new experiences through technology. However, we argue design knowledge and practice in inclusive design has lagged behind the broader HCI field in two different, but related, ways. First, when new technology is released, an implementation lag in designs for access and enablement invariably lead to late adoption of technology for people with disabilities. Secondly, this implementation lag has resulted in a conceptual lag, where to solve these problems the research field remains grounded in HCI methodologies from First and Second Waves. This results in a reliance in checklist style engineering approaches that are unable to properly support user experience design. We explore these ideas in the two examples of the web and digital games, and argue that while we must not supplant previous approaches, we need to decouple the implementation lag from the conceptual lag to change inclusive design research and practice. We argue that we must not only plan for accessibility, but instead adopt pluralistic approaches that recognise the diversity of lived experiences of people with disabilities, and use them to design options for people to customise their own inclusive experiences.

---

C. Power (✉) · P. Cairns  
Department of Computer Science, University of York, York, UK

The AbleGamers Charity, Kearneysville, WV, USA  
e-mail: [christopher.power@york.ac.uk](mailto:christopher.power@york.ac.uk)

M. Barlet  
The AbleGamers Charity, Kearneysville, WV, USA

## 10.1 Introduction

Human-Computer Interaction has, on occasion, been described as the conscience of computing:<sup>1</sup> it reminds researchers and developers to remember the people who use their systems and to try to do good for those users. Accessibility is arguably the conscience of HCI. The field regularly gets excited about new technologies like Twitter, live-streaming and public interactive art, but it is Accessibility that reminds us to not forget people with disabilities for whom such technologies present new barriers to full participation in modern life. However, while mainstream HCI, through the Third Wave, has embraced understanding the unique, situated and value-driven approach of individual people using interactive systems (Bødker 2006, 2015), we make the case that Accessibility research has lagged behind the mainstream and, in many ways, has yet to enter the Third Wave. The reasons for an implementation lag in access that necessarily follows the appearance of any new technology are many and complex, ranging from economic factors, to societal attitudes, to designer training. However, we make the case that, as technology enters every corner of our lives, the conceptual lag cannot continue. If it is the case that leisure technology can have intrinsic benefits on our quality of life (Danilina et al. 2017), it is no longer enough to discuss inclusion in terms of how successful or efficient people are with systems. Inclusion needs to address the full range of outcomes, by recognising cultural, societal and physical facets of users that shape the design and evaluation of interactive technologies (Harrison et al. 2011).

In this chapter, we examine the current state of inclusion research for digital technology in terms of the first two waves of HCI research. First Wave approaches focus on “the user” as an information processor that needs to get information in and out of the machine. This serves to provide basic access to technology by addressing specific perceptual, cognitive or motor mismatches between the system and the individual. Once we have these basic access needs addressed, Second Wave approaches provide framing of people with disabilities as actors in the system, who use the system to achieve specific goals. These approaches tend to focus on task success rates and efficiency of operation, with these traditional usability measures encapsulating whether an individual is enabled to complete their goals.

However, in conceptualising these different levels of inclusion this way, a cycle of repeating exclusion begins to appear in the technology landscape. This cycle usually begins with a new technology being introduced without basic access needs of people with disabilities being considered. Then, in response to this, a series of initiatives start to bolt on alternatives and enhancements which allow people to get information in and out. Once we have this basic access, design then moves to try to enable people with disabilities to accomplish their goals. As we begin to understand how people with disabilities interact with the technology, we begin to encode the problems they encounter in sets of rules that designers should avoid, which then moves us to checklists for designers to refer to when they are doing their work.

---

<sup>1</sup>Alan Dix, <http://alandix.com/academic/papers/thirty-years-of-HCI-2014/>, Retrieved December, 2017

However, all too often, this is a never ending process. As problems present themselves, new rules are codified, and new checklists made. Further, when a new innovative piece of technology comes out, we start the cycle again because designers struggle to transition the rules from the previous generation into new designs.

This cycle defines the implementation lag in Accessibility: there is (and always will be) a lag from the introduction of a new technology to the time at which the technology is made accessible to users with disabilities. However, in examining the field of Accessibility, a conceptual lag is also seen in the progression of research approaches through the First, Second and Third waves, where Third Wave approaches are sparse. To illustrate this, this chapter delves deeply into two specific domains. The first is the web, where Second Wave approaches have used the above cycle to make real progress in inclusion on the web; however, only recently has that domain begun to consider the experiences of users in situ. Web accessibility is only now beginning to adopt Third Wave thinking. This is a conceptual lag of Accessibility research behind that of mainstream HCI.

In contrast, our second domain is that of digital games, which has come to maturity during the Third Wave of HCI, where understanding experience and the impact it can have on our lives is the key outcome of research. This domain requires Third Wave thinking from the outset even though the implementation lag of the First and Second Waves also must be addressed for each new game and game technology.

This chapter proposes decoupling the (inevitable) implementation lag of access from the conceptual lag of current research approaches. We close the chapter with a discussion of what the Third Wave of experiential research might look like for increasing inclusion in digital games. We relate it back to where we have come from as a field of passionate Accessibility researchers and practitioners. We acknowledge the contributions that have brought us this far in reducing the exclusion of people with disabilities from our digital society, while challenging the field to tackle new problems from a more broad set of approaches from across the three waves of HCI research.

## 10.2 First Wave Inclusion

A lot of our early work in inclusion with people with disabilities is framed around the idea that people should have access to technology. This framing sets up a very binary notion: is it accessible or is it not accessible? This framing originates in a way that many of our concepts in interactive systems are established: by using metaphor from the physical world into the digital. We understandably took the notion of access, enshrined in law in many developed and developing economies before the personal computing revolution, and brought it into the virtual world. The idea of physical access to buildings, for example, was well understood, so naturally we took this notion of access to the digital world.

As a result of this framing, much of the early work in inclusion is situated in the First Wave. The focus of early accessibility work is framed in issues relating to

translating information into an alternate modality, replacing a sense with technology, or in technology that allows input to the machine, so that users can get information in and out of the system. Situated in the heart of the medical model of disability, much of the early work was not driven around what users wanted to do, but instead about overcoming impairments. Due to the fact that users were often unable to even take basic actions in the system, looking at any interaction beyond the fit between human and machine was not possible. It was not until the early 1980s that the social models of accessibility began to take hold in areas of research and practice, where we shifted to think of disabilities as a mismatch between a person's abilities their environment (or system). As a result, this is where the first conceptual lag occurred. Mainstream HCI had been considering the interaction between a person and the system well before this became an important discussion in inclusive design (Card et al. 1986).

Due to this framing, we see cycles of inclusion and exclusion as technology paradigms shifted. Taking one example, if we consider the case of people who are blind, there have been several paradigm shifts in computing that have led to a cycle of inclusion and exclusion. Blind programmers were explicitly recruited to work on punch card systems in the 1960s (Pullin 2009). When visual terminals were introduced, and punch cards gradually fell out of use, blind programmers found themselves excluded because they could not get at program information. As Bach-y-Rita and others begin work on sensory substitution using brain plasticity to replace vision with touch (Bach-y-Rita and Kercel 2003) and the first embossers are introduced to give access to large amounts of text previously available only in traditional books, we had display terminals see widespread use in computing labs. Then, just as the first screen readers are introduced to provide access to digital text (Adams et al. 1989), graphical user interfaces (GUI) are established as the dominant form of interaction for the next 20 years. Finally, as the first screen readers purpose built to navigate GUIs enter the research space, we see a shift to the web for digital transactions (Mynatt and Weber 1994; Petrie and Gill 1993; Petrie et al. 1995).

In this example, the conceptual lag is understandable. It is not possible to apply Second Wave thinking about tasks until First Wave engineering of access has been established. Invariably, over the last 50 years, behind the vanguard of change and innovation, there is a ground swell of researchers, practitioners, tinkerers and users who, out of necessity, revert back to trying to solve the First Wave human factors problems of getting information out of the machine to users, or for users to put instructions into the machine. Consequently, the problem of providing access remains an active and necessary part of inclusion research today.

### 10.3 Second Wave Inclusion

Towards the end of the twentieth century, there were a number of changes in the way that we conceptualised inclusion, that coalesced into a new framing for inclusion research and design practice. First was the framing of disability not just as medical

conditions, but instead as a social construct. From this perspective, disability is something that can happen to anyone wherever there is a mismatch between the designs present in society and the abilities of individuals which prevents them from achieving their goals. This shift in framing co-occurred with strong civil rights movements across developed economies. A large number of legal frameworks were introduced throughout the 1990s that enshrined the rights of people with disabilities to not be excluded from society, for example: the Disability Discrimination Act of 1995 in the UK (DDA), and the Americans with Disabilities Act of 1990 in the USA (ADA). In 2000, Section 508 of the Rehabilitation Act of 1973 (Section 508) had its first set of implementation guidance notes published by the access board. Twenty five years on, many of these laws have been amended or wholly replaced to include access to digital services and information as part of their coverage (Clarkson and Coleman 2015), with the Equality Act 2010 in the UK and updates to Section 508 guidance in 2017 in the USA.

At the same time, we had number of factors emerge that necessitated the change in the way we talked about design in digital systems. First was an ageing population resulting in more people with a diversity of needs in terms of sensory, cognitive and mobility support, and this was combined a growing movement of including people with disability in society (Clarkson and Coleman 2015). At that time, we also saw personal computing technology become dramatically cheaper and more available. Machines became more powerful, including the necessary processing crunch to run a variety of assistive technologies. Further, they included robust hardware for generating sound and graphics, meaning that more options were available for developing assistive technologies for a wider variety of people. And within HCI, we had a shift away from humans as being information processors, and instead being actors in an active dialogue with systems to achieve goals, what Bødker (2006, 2015) and others identify as the Second Wave of HCI. For inclusion, this meant looking at how do we enable people with disabilities to achieve their goals.

As a result of all of these factors, we saw the emergence of three distinct movements about designing for diversity. The first is the Universal Design movement that appeared in the USA circa the mid-1990s. Mace (1988) bridged the gap between First Wave and Second Wave, relating issues of cognitive psychology to engineering in the way that Norman (1983), Nielsen and Molich (1990), and others did for mainstream usability. The team at North Carolina State University issued a set of guidelines that took many of the key concepts of usable design and related them to the challenges encountered with different motor, cognitive and sensory disabilities. Emerging around the same time were the design philosophies of Design for All (Europe) (Bühler and Stephanidis 2004) and Inclusive Design (UK) (Clarkson et al. 2013) which placed consideration of diversity as a component of design life cycles. This firmly moved inclusion into the Second Wave because now designing for user tasks, not just the design of input and output, was important to achieve inclusion. While work on access to different specific technologies continued (Fraser and Gutwin 2000; Brewster 2002; Wobbrock et al. 2005), around this time we began to see more of a discussion around technology as barriers to people participating in society (Gregor and Newell 2001; Gregor et al. 2005; Jacko and Hanson 2002;

Stephanidis 2001), and more publications talking about Accessibility using more traditional usability criteria of effectiveness and efficiency (Jacko et al. 2002). However, once again the field was conceptually lagging behind the practices in mainstream HCI. As the rest of the HCI world started to talk about user experience, many Accessibility researchers were only getting to grips with what inclusion really meant as a term in design practices for users being enabled in systems, and in particular the adoption of designing to checklists and heuristics.

### 10.3.1 *An Example: Web Accessibility*

There is perhaps no better example of the process of transitioning inclusive design through the first two waves of HCI than work on inclusion on the web. During the middle of the 1990s there was increasing concern about people with disabilities, in particular motor and sensory disabilities, being left behind by the shift of services and information to the web. In 1994, in a now famous keynote in accessibility circles, Tim Berners-Lee first mentioned disability access<sup>2</sup> and kicked off what would be a flurry of activity around the sphere which culminated in the publishing of the Web Content Accessibility Guidelines 1.0 (WCAG 1.0) (Chisholm et al. 1999).

Early problems on the web were often conflicts between assistive technology and the technology on which the web ran. These problems are firmly situated within the First Wave, and these mismatches in technology can still occur, but in lesser numbers than the early days of web accessibility. WCAG 1.0 had a number of guidelines that were intended to help alleviate these problems, and many of the guidelines were grounded at the level of web code, with checklists supporting developers by indicating what their code should and should not include.

As people with disabilities gained more access, we were able to ask different questions about how people with disabilities use the web, and try to approach things in a more Second Wave, usability oriented, point of view (Iwarsson and Stahl 2003; Shneiderman 2000; Petrie and Kheir 2007). For example, early work on strategies of blind screen reader users and low-vision magnifier users subsequently influenced the creation of new assistive technologies (Theofanos and Redish 2003).

However, there were also reports that pointed to problems that were beyond what existing guidelines covered. One of the largest ever conducted, the Disability Rights Commission report of 2004 (Disability Rights Commission 2004), pointed to a number of problems that looked distinctly like usability problems encountered by non-disabled users. For example, unclear and confusing navigation mechanisms that were not only a problem for blind screen-reader users but also cross-cutting with all groups with disabilities who were engaged in the studies detailed in that report.

---

<sup>2</sup>It seems to be a mandatory requirement for all web accessibility articles to include the precise text, so: “The power of the Web is in its universality. Access by everyone regardless of disability is an essential aspect.” – Tim Berners-Lee.



One way to look at these problems are that they are the first problems to be encountered after we had succeeded in providing users with basic *access*. Disabled users were able to identify that there was a navigation bar and that it was very confusing, and so users were unable to complete their goals. Consequently, this should not be seen as a condemnation of WCAG but instead an important milestone. The maturity of access had reached a point that these very thorny, difficult, design problems could be discussed and debated within communities of users and practitioners.

In 2008, WCAG was refined, updated and re-arranged (Caldwell et al. 2008). Guidelines were linked with overarching principles tied to core HCI concepts, such as making it so users with disabilities could perceive, understand and operate the web. The checklist approach was maintained, and success criteria were phrased carefully to be testable, and techniques for implementation were linked to the success criteria. Within WCAG 2.0, we saw an evolution of a set of guidelines that commuted the First Wave into the Second Wave. No longer were we trying to simply “create a logical tab order through links” (WCAG 1.0), but instead we were trying to make content understandable to people with disabilities by conveying a “Meaningful Sequence” through our content and seeking to provide “Information and Relationships” (WCAG 2.0) to users trying to link content together content within a page to complete their goals.

### 10.3.2 *Emergent Experiences*

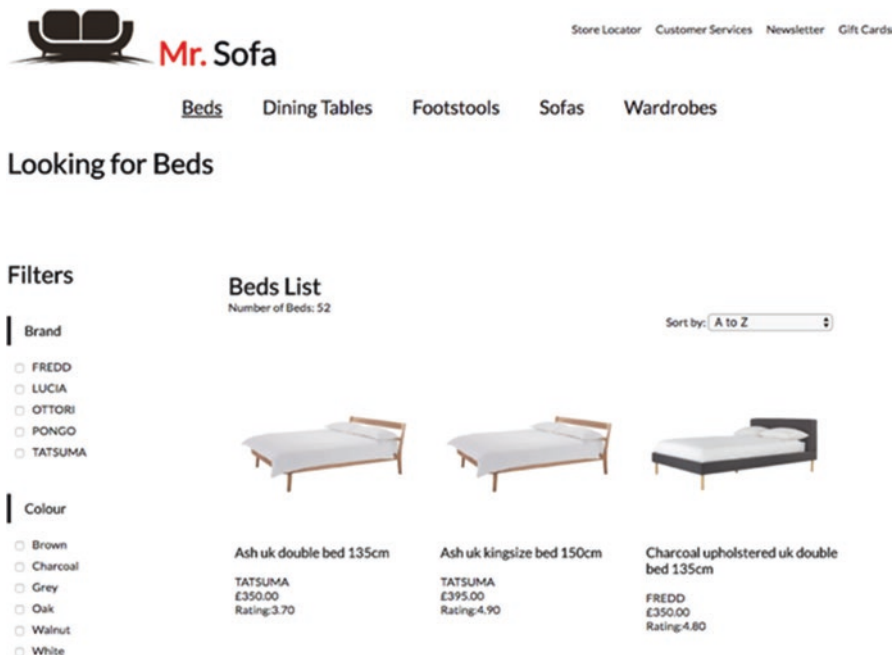
After this shift of WCAG 2.0, several studies began to collect and classify the problems of users with disabilities on the web (Power et al. 2012; Rømen and Svanæs 2012). Users were found to have problems with information overload, or not finding what they were looking for, or being confused by irrelevant information, and these problems were not directly covered by guidelines in WCAG 2.0.

However, we should not expect them to be covered. Many of the problems described above only occur when we have solved some of the navigation problems found during the transition between WCAG 1.0 and WCAG 2.0, and many of the problems are conditional on the users being able to more effectively take action in web systems. As a result, they are new and different problems than those covered by WCAG 2.0.

This is where things get complicated: do we now amend the guidelines again to capture these problems and recommend ways to solve them? This could be quite hard given the variety of designs now on the web. Even more problematic is the question: where do we stop writing guidelines? If we are forever finding new problems, we will reach a point where we no longer have a checklist but a catalogue!

For a more clear demonstration of this point, we present the following description of some recent user studies recently completed at the University of York (Savva 2018).

Assume that we have a shopping website laid out in the modern fashion for online shopping as shown in Fig. 10.1. It has all the standard design elements: a



Mr. Sofa

Store Locator Customer Services Newsletter Gift Cards

Beds Dining Tables Footstools Sofas Wardrobes

## Looking for Beds

**Filters**

**Brand**




- FREDDO
- LUCIA
- OTTORI
- PONGO
- TATSUMA

**Colour**

- Brown
- Charcoal
- Grey
- Oak
- Walnut
- White

**Beds List**  
Number of Beds: 52

Sort by: A to Z

		
Ash uk double bed 135cm TATSUMA £350.00 Rating: 3.70	Ash uk kingsize bed 150cm TATSUMA £395.00 Rating: 4.90	Charcoal upholstered uk double bed 135cm FREDDO £350.00 Rating: 4.80

**Fig. 10.1** An example furniture shopping website with a search and filter interface from (Savva 2018)

heading area, a navigation bar, a breadcrumb, some filters, and a content area where products are displayed and refreshed based on the filters selected.

Now, we know that one key element of design for blind screen reader users is to have headings that describe sections of the page (Theofanos and Redish 2003; Watanabe 2009; Power et al. 2013). If a first level heading (element h1) is placed at the top of the page, then the page will come back as being technically accessible, by which we mean that an automated accessibility tool will tell the developer that the heading tests have been passed. However, we also know that the effectiveness, of blind screen reader users will be low in completing their goals with only a single heading. In fact, they will almost certainly report a number of the problems identical to those reported circa 2004 (Disability Rights Commission 2004; Power et al. 2012).

To address (some of) these issues, we could add new headings to each visual section of the page so that blind users can find their way around the different sections. When evaluated by screen reader users, we will find that new problems are reported. Among them will be that users are unable to understand what happened on the page in relation to the filters they selected. While we have higher success rates climbing into the 85–90 percentile range, the user problems have morphed, and are clearly still having a detrimental effect on users.

To address those emerging problems, we could again add new headings, this time into the filters to help users find their way around within the filter sets. This time,

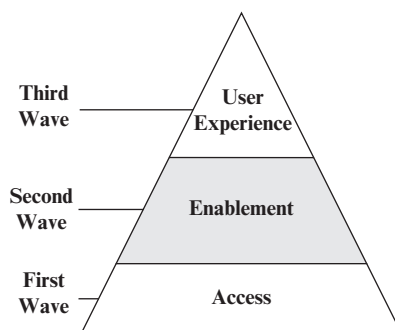
success rates are between 95 and 100% success rate. However, in evaluations users will report issues of insufficient feedback that prevents them from knowing whether or not the content has updated the way they expect. Further, they report the refresh of the page puts the cursor of their screen reader back to the top of the page, and while they can find their way back to the filters or the content, there is large amounts of extra effort involved, leading to frustration.

Looking at the above example, even if all of the above designs were encapsulated into a checklist, not only would the checklist become unwieldy, there would still be new problems that have to be addressed. Even though we might be able to say that structure is present in a meaningful way for users to navigate via headings, the situated nature of users' tasks means that there are differences in the experiences and expectation of users that we cannot account for a checklist based approach. The epistemological orientation of checklist based approaches is rooted in that of planned accessibility, that of codifying and testing for a set of properties to which the designer will potentially provide access (Hedvall 2009). In this regard, a well constructed checklist can ensure properties are present or absent, but it is impossible to predict the actual experience users have from those properties alone. When technology is put to use by people with disabilities, they bring with them all of their goals, expectations, and all of their own experiences with them, which shapes the experience into what Hedvall refers to as lived accessibility (Hedvall 2009). This lived experience, which will change over time, based on both the individual using the technology and the contextual factors around them, needs to be accounted for an understood in our design processes. This is true of all users, not just users with disabilities, and is at the heart of the move of HCI into the Third Wave (Harrison et al. 2011).

## 10.4 From Access to Experiences

The example of the web brings to the fore what it means to design something inclusively. When grounded in past successes and challenges, we start to see an overall picture emerge depicted in Fig. 10.2.

**Fig. 10.2** The layers of inclusion, from basic access where users can perceive and operate aspects of the system, to enablement where they can achieve goals, which then yield different experiences



First, users need access. They need to be able to operate systems, and consume information that is presented to them. In order to address these needs, we need to apply First Wave approaches where we build and test pieces of technology to ensure they are of appropriate fidelity for users. Once those needs have been addressed, users can be enabled to do things in the system through Second Wave approaches that focus on ensuring users can meaningfully act in the system. While succeeding, or sometimes failing, in their goals, people with disabilities will have situated experiences, that are shaped by the technology they are using, the options available to them, their own self-efficacy and competence, and many other facets of their inter-action.

As a result, we can see access and enablement as precursors to any type of experience, and as such it is important to not ignore work that has come before. Perhaps more than any other field in HCI, we must not supplant that which has come before. We must remember what we have learned about access and enablement, and where and when to use different types of techniques.

However, in regards to capturing and understanding the experiences of users with disabilities in interactive systems, we have very few tools and techniques. Worsening the situation, with the advent of technology that is wholly about hedonic experiences, and not about pragmatics, there is a lack of criteria that we can measure in regards to “task success.” In that design space, we need new ways to conceptualise how we measure success in inclusion, and develop design thinking approaches that take into account the situated actions, the experiences, and the values of people with disabilities with whom we are designing. We need to decouple the conceptual lag from the implementation lag, and start thinking ahead of how we design with experiences of users with disabilities in mind. Only then, will see truly inclusive technology.

## 10.5 Inclusion in Digital Games

When thinking about what inclusion in the Third Wave of HCI looks like, it is tempting to use the web as an example. After all, it has dominated much of the research landscape in inclusion over the last 20 years, and the web is still an area of expanding influence in our lives. People spend time on it not only to pay their taxes, but also to pass the time watching funny cat videos, browse the daily news, read Wikipedia, book a hotel, and so on. Similar to other technologies, and perhaps magnified by the flexibility of the web, there has been a conceptual lag in uptake of different types of approaches, but we are now starting to catch up in thinking about what does an user experience on the web look like for people with disabilities (Horton and Quesenbery 2014; Horton and Sloan 2015; Aizpurua et al. 2016). There is also an encouraging increase in individuals who are interested in examining our methodologies around web accessibility (Savva et al. 2015; Brajnik et al. 2016; Bigham et al. 2017). Indeed, early in the Second Wave of inclusion research, Newell and Gregor were pushing the boundaries of Third Wave approaches with their work

in understanding the needs of older adults on the web (Newell and Gregor 2000; Newell et al. 2011). However, even with these promising signs, we would argue that most things on the web are goal driven and enablement tends to dominate the discussions. While many of the activities above may be considered fun (and others clearly not), only the funny cat videos appear at first glance to be solely about leisure, and the rest are measured primarily by whether or not someone can do the thing for which the site was designed.

In order to avoid the task-based bias of web research, therefore, we adopt digital games as the place to explore aspects of the future of inclusion in the Third Wave. Digital games (games hereafter) represent probably one of the most extreme examples of where technology is used by people almost exclusively simply to provide some kind of experience. Games elicit a variety of different experiences from players and at their heart are about play, representing perhaps the quintessential example of leisure technologies discussed by Bødker (2006, 2015), representing an evolution of technologies to meet rest-of-life needs. Whether a player is playing hide-and-seek alone against the undead, raiding dragon hordes with guildmates, or crushing rows of candies, there is the potential for everyone to find a game they want to play.

### *10.5.1 Access in Games*

Much like other interactive systems, access in games is primarily about taking action within the game, or consuming information about the state of the game through different modalities.

Players need to control different aspects of games, but what they need to control will vary wildly based on the type of game it is and the platform on which it runs. For example, consider *Monument Valley*. This beautiful game uses Escher inspired impossible objects to create puzzles that players need to solve to advance a princess through the levels. In order to do this, players must tap quite precisely on the screen where they want the princess to move. Compare this to a game like *Dragon Age* on a games console using a controller. In *Dragon Age*, players navigate in three dimensions around an immense world, while also managing a variety of combat oriented spells and attacks to defeat enemies, all through the combinations of 2 analog sticks, 4 bumpers and 4 buttons. Further compare these two sets of controls to those of the popular *Starcraft* series on a personal computer (PC), where dozens of mapped keys, mouse movements and clicks are required to gather resources, produce units, and then send forces to pile-on your enemy. Similarly, when we look at the presentation of information in the above games, we find a similar variety in types of information is provided (e.g. health bars, status of attacks, what players can interact with) and the modalities they are presented in (e.g. visual movement of surroundings, sounds of footsteps behind the heroes, text streams of unit statuses).

Clearly, there are many different ways that the above games could exclude players with disabilities. Players are diverse, and in many cases there will be co-occurrence of a number of different mismatches between players and the game

controls and presentation. We could prescribe specific solutions, but there is no guarantee that it would meet the needs of even a small number of players, and indeed we would need to do it for every different type of game.

The reality is, each player will encounter a different set of problems within a game that need to be addressed, and the best people to tell us what works for them is the players themselves. As opposed to prescribing designs to game developers, and in particular petitioning to remove variety from the gaming space, in our experience it is better to advocate for options be provided to people with disabilities to customise their controls and their presentation. For controls, some of these options will be of the form of allowing alternate controllers or providing ways to remap buttons. For presentation, options can include allowing alteration of colour or size of things on the screen, resizing elements of the user interface, or hiding unnecessary details. Providing the option for players to do these different things, allows players to customise things to meet their own needs and moving them forward to enablement within the game space.

### ***10.5.2 Enablement in Games***

Where games become really interesting in inclusive design, and where we start to bump up against the limits of Second Wave approaches, is when we start trying to identify what it means for players to be enabled in the game space. While it is easy to say that if a player can play the game, then they are enabled, that may not be the end of the story for that player. What if the game isn't fun?

For instance, let's assume that a player with mobility disabilities sets up a controller to play *Dragon Age* such that they can reach all of the controls to control their party, trigger spells and abilities, and make dialogue choices in the narrative. Is that enough? If the game operates at a speed where the player's reactions cannot keep up, and their performative uncertainty (Costikyan 2013) is too high, then it is likely that they will quit the game before long and try something else. Luckily, *Dragon Age* provides mechanisms for pausing the game, reducing the level of challenge to a point where it is balanced against the player's experiences in the game so that the game remains fun. Our player can still play the game, and the developers get the money from a sale to a now loyal customer. Everyone wins.

When we unpick this example, the remapping of the controls removed the fundamental barriers that kept our player from playing the game. Beyond that barrier, there were further barriers introduced by the level of challenge presented by the game. The challenges in the game itself were likely balanced with an average player in mind, and tested with average players in the gaming chair. For our disabled player, their situated context is very different from the average. Some of that context might be fixed over time, such as the range of motion the player has in their hands and arms. On the other hand, some of the context, like our player's natural reaction time, might improve over time, similar to any other player who is new to a particular game or genre of game. If we provide ways and means for players to change the

**Fig. 10.3** The layers of inclusion in games, starting with basic controls and presentation, providing access to challenges which may further need to be customised by players, so that they can have fun and other accessible player experiences (APX) in games



game, so that the challenge is better suited to their current context, then the result is another game that is now including those players. We end up with our levels of inclusion for games looking more like Fig. 10.3.

From the analysis of this example, enablement in games can be conceptualised as being able to have options around the challenges presented in games. After gaining access, players need to be able to shape the challenge so that it better suits them. Due to the range of different challenges presented in games it becomes impossible to write a checklist to meet them all, nor would we expect to do so. Where it is appropriate, we can provide a checklist at the level of whether or not the options have been included, which is what AbleGamers has done with their Includification guidelines (Barlet and Spohn 2012). These guidelines ensure that developers can question whether or not they have found alternatives that players can use, and whether they have considered adding different features to their games. However, it cannot tell them how to design those options, or if the game will be fun.

### 10.5.3 *Experience in Games*

However, the diversity in games highlights another key aspect that makes them distinctive for exploring inclusion from a Third Wave perspective. Games are intricately connected to players and their standpoint as to why they are playing. When someone goes to pay their taxes on the web, there is basically one reason they are doing that: to pay their taxes. When people play a game, their individual reasons for play are as varied as the players. In some cases, people may be looking to simply de-stress and unwind after a busy day (Collins and Cox 2014). Alternatively, players may want to earn achievements, explore a world, simply exist in the world to spend time with friends (Bartle 1996), or something else entirely. Even when goals are shared by players, such as “to win”, what it means “to win” may vary from player to player. One person might see being at the top of the leaderboard as being the only acceptable winning condition, while another may just want to take the other person

down with them. In any of these cases, the standpoint of the players frames and contextualises the challenges presented within the game in different ways. Even when players all experience the same game, they may all have different takeaways from that game in terms of experience.

Indeed, this is what makes digital games perhaps the most important technology for improving opportunities for diverse experiences for people with disabilities. Digital games themselves expand the range of experiences that people with disabilities can have in our society. It allows players get out beyond their four walls, to connect with other people, and to engage in a shared experience that can be related to and discussed well beyond any individual play session. People discuss their play sessions, share stories, avoid spoilers, and can forge lifelong friendships that extend out into the “real” world. Games provide not only a place where people can connect, but something people can connect over.

So what are the ways that we can help games developers design new games that are inclusive? It seems to us that when it comes to designing for experiences with people with disabilities, First Wave access and Second Wave enablement are necessary but insufficient. To achieve the experiences that designers and developers want for their users, disabled or not, requires Third Wave thinking.

## 10.6 Third Wave Inclusion: Inclusive Experiences

Despite the extensive research in accessibility and inclusion, and the notable improvements in accessibility of many different interactive devices, to achieve inclusive experiences requires a further evolution in design research and practice.

We need a step change away from the idea that we are planning for accessibility, and instead focussing on what are the lived experiences of players when they encounter our games. In what Hedvall (2009) referred to as epiaccessibility, we need to consider not only what types of options we are providing, but we need to consider the range of ways people will use those options, how they will opportunistically use technology, the expectations they bring playing games, and many other facets users’ situated contexts.

While we still have a long way to go in providing access and enablement in the game space, we now have enough players playing games that we can begin to understand how the myriad of factors that impact their play and their subsequent experience. Further, we need to acknowledge that while different players have different standpoints, in the end everyone wants to have a “good game.”

For research, this means that we need to go back to first principles, and begin understanding what players with disabilities expect and want from their experiences in games. Some of these, such as a desire to socialise with friends, or to rise to meet a challenge in game, or to engage with an emotionally satisfying story, are shared among all players, but will be shaped by the cumulative experiences each different person has had in their lives. In other cases, some experiences will be very specific



to individual players with disabilities, such as the level of comfort they have using a piece of assistive technology, or how they use games to manage pain in their day-to-day lives. Capturing, conceptualising and understanding these accessible player experiences (APX) will push beyond the types of traditional approaches where we have users undertake a task and count the number of times they get something “right”. It will require a variety of methods drawn from participatory design, contextual design, ethnography and beyond, to understand the different facets of these experiences. We will need not only large studies with many players, but also more focused, intimate case studies with players or groups of players so we can get rich descriptions and depictions of the experiences of players with disabilities and their peers. In places where players can identify and describe the experiences of players, we can begin to build measures for testing different designs.

However, with any measure that we define, we need to be aware when interpreting the outcomes that the experience may have a substantially different meaning for different players. For example, consider the situation where a player has all of their access needs met in a first person shooter (FPS) game. They can play the game, but don’t find it particularly fun because they die within seconds of their respawn point. Through the use of target assist, they reduce the challenge to the point where they are able to tag their opponents. They still get knocked out, but they have enough of a chance to take the opponent down that they find it fun. Another player might be an avid FPS player, and they might consider that they don’t need target assist and resist turning it on even though it would allow them to compete on a more even footing. Both of these players would probably rate the challenge as being particularly high, but that does not necessarily mean that either is a negative experience.

For game designers, we need to identify new ways of conceptualising the experience of players with disabilities in terms of the goals of the designer. While checklists of accessibility options provide a means of challenging assumptions in design, as demonstrated in the web, it is not enough to enable designers to deliver experiences. We need to change the paradigm. Designers want to build games that deliver some kind of experience to their players, so we need to shift to providing designers insight into what are the physical, cognitive and emotional challenges of games (Denisova et al. 2017) and the ways players will want and need to customise that challenge. For example, if a designer wants to tell a moving story about swashbuckling space cowboys, they want all players to be able to take in that story. However, for a player who may struggle to read dialogue options, with the cognitive challenge being too high, we want designers to think of ways to deliver that story without compromising the emotional resonance of it in the player. Similarly, if a player finds it difficult to sit through large cutscenes due to their attention preventing them from engaging in the emotional challenge the story offers, we want designers to be thinking of ways for players to skip the story, yet still have the narrative feed into the game’s main theme.

With these examples, it is easy to see that we cannot prescribe lists of do’s and do not’s to game design. We need to be developing deep understanding of how players experience games, generated from a wide variety of methods, and deliver it into

the hands of designers. We need mechanisms by which designers can integrate this knowledge into their own practices, and situate it in the way they design, as it is starting to be done on the web. Further, we need deep, contextual knowledge about how games, or other systems, are created so that we can graft these understandings onto the language and goals of designers. By doing this, we can provide new approaches. Only by understanding these different standpoints will we be able to make inclusive design something new and supportive, and possibly reach the point where games are designed to be inclusive simply because it is part of how we do things.

## 10.7 Decoupling the Lags

The implementation lag of providing access to games seems unavoidable. New games use different modalities, such as haptic or VR interactions, in new ways and new platforms offer new controllers, like the Nintendo Switch. However, it is already being recognised by games designers that there are solutions, such as providing options in games, that make engineering access easier (Barlet and Spohn 2012). This is reducing the implementation lag from new games and games technology to players with disabilities being able to play. However, there will always be some degree of implementation lag because developers cannot easily account for the diversity of players with quick fixes. There will also be a need for further work as each new game technology is produced.

However, what need not happen is the conceptual lag. If we remain with First and Second Wave thinking in game design, games cannot deliver to all the players the experiences that are, in essence, their purpose. Third Wave thinking is necessary from the outset. If designers think of a diversity of players with their own goals, values and contexts of play (Harrison et al. 2011), they can be thinking as much about diversity in terms of the disabilities of their players as much as they think about their capabilities as gamers. No one player or type of player needs to be privileged in terms of the experiential goals that the designers want to deliver. If this conceptual lag is removed, then there is no need for further work once the implementation catches up: the diverse experiences of a game will be ready to whomsoever is enabled to play.

Games are in some sense pure experience (Cairns 2016), and as such, they provide an ideal domain in which to explore Third Wave thinking for inclusion. Indeed, this is perhaps one of the largest opportunities that will be provided by pursuing this pluralistic approach in games. By shifting the epistemological approach we use to design games, to be situated on experiences and outcomes for players with disabilities, we will undoubtedly learn new methods that can be transferred back to the web and other systems.

## References

- Adams FR, Crepy H, Jameson D, Thatcher J (1989) IBM products for persons with disabilities. In: IEEE Global Telecommunications Conference and Exhibition Communications Technology for the 1990s and Beyond' (GLOBECOM), 1989, IEEE, pp 980–984
- Aizpurua A, Harper S, Vigo M (2016) Exploring the relationship between web accessibility and user experience. *Int J Hum Comput Stud* 91:13–23
- Bach-y-Rita P, Kercel SW (2003) Sensory substitution and the human–machine interface. *Trends Cogn Sci* 7(12):541–546
- Barlet MC, Spohn SD (2012) *Includification: a practical guide to game accessibility*. The Ablegamers Foundation, Charles Town
- Bartle R (1996) Hearts, clubs, diamonds, spades: players who suit muds. *J MUD Res* 1(1):19
- Bigham J, Lin I, Savage S (2017) The effects of “Not Knowing What You Don’t Know” on web accessibility for blind web users. In: Proceedings of 19th International ACM SIGACCESS Conference on Computers and Accessibility (AS-SETS 2017). ACM, New York, pp 101–109
- Bødker S (2006) When second wave hci meets third wave challenges. In: Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles. ACM, New York, pp 1–8
- Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. *Interactions* 22(5):24–31
- Brajnik G, Vigo M, Yesilada Y, Harper S (2016) Group vs individual web accessibility evaluations: effects with novice evaluators. *Interact Comput* 28(6):843–861
- Brewster S (2002) Visualization tools for blind people using multiple modalities. *Disabil Rehabil* 24(11–12):613–621
- Bühler C, Stephanidis C (2004) European co-operation activities promoting design for all in information society technologies. In: Computers helping people with special needs. ICCHP, Paris, pp 625–625
- Cairns P (2016) Engagement in digital games. In: Why engagement matters. Springer, pp 81–104
- Caldwell B, Cooper M, Reid LG, Vanderheiden G, Chisholm W, Slatin J, White J (2008) *Web Content Accessibility Guidelines WCAG 2.0*. Web Accessibility Initiative, W3C
- Card S, Moran T, Newell A (1986) The model human processor—an engineering model of human performance. *Handbook of perception and human performance* 2:45–1
- Chisholm W, Vanderheiden G, Jacobs I (1999) *Web content accessibility guidelines WCAG 1.0*. Web Accessibility Initiative, W3C
- Clarkson PJ, Coleman R (2015) History of inclusive design in the UK. *Appl Ergon* 46:235–247
- Clarkson PJ, Coleman R, Keates S, Lebbon C (2013) *Inclusive design: design for the whole population*. Springer, Heidelberg
- Collins E, Cox AL (2014) Switch on to games: can digital games aid post-work recovery? *Int J Hum Comput Stud* 72(8):654–662
- Costikyan G (2013) *Uncertainty in games*. MIT Press, Cambridge, MA
- Danilina O, Cox A, Fonseca A, Johnson S (2017) The effectiveness of serious video games on mental health related outcomes: systematic review. In: Proceedings of CHI play 2017. ACM, New York
- Denisova A, Guckelsberger C, Zendle D (2017) Challenge in digital games: towards developing a measurement tool. In: Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems. ACM, New York, pp 2511–2519
- Disability Rights Commission (2004) *The web access and inclusion for disabled people: a formal investigation conducted by the Disability Rights Commission*. The Stationary Office
- Fraser J, Gutwin C (2000) A framework of assistive pointers for low vision users. In: Proceedings of the fourth international ACM conference on Assistive technologies. ACM, New York, pp 9–16
- Gregor P, Newell AF (2001) Designing for dynamic diversity: making accessible interfaces for older people. In: Proceedings of the 2001 EC/NSF workshop on Universal accessibility of ubiquitous computing: providing for the elderly. ACM, New York, pp 90–92

- Gregor P, Sloan D, Newell AF (2005) Disability and technology: building barriers or creating opportunities? In: *Adv Comput*, vol 64, Elsevier, p 283–346
- Harrison S, Sengers P, Tatar D (2011) Making epistemological trouble: third-paradigm HCI as successor science. *Interact Comput* 23(5):385–392
- Hedvall PO (2009) Towards the era of mixed reality: accessibility meets three waves of HCI. In: *Symposium of the Austrian HCI and Usability Engineering Group*. Springer, Heidelberg, pp 264–278
- Horton S, Quesenbery W (2014) *A web for everyone: designing accessible user experiences*. Rosenfeld Media, Brooklyn
- Horton S, Sloan D (2015) Accessibility for business and pleasure. *Interactions* 23(1):80–84. <https://doi.org/10.1145/2843590>. <http://doi.acm.org/10.1145/2843590>
- Iwarsson S, Stahl A (2003) Accessibility, usability and universal design—positioning and definition of concepts describing person-environment relationships. *Disabil Rehabil* 25(2):57–66
- Jacko JA, Hanson VL (2002) Universal access and inclusion in design. *Univ Access Inf Soc* 2(1):1–2
- Jacko JA, Scott IU, Sainfort F, Moloney KP, Kongnakorn T, Zorich BS, Emery VK (2002) Effects of multimodal feedback on the performance of older adults with normal and impaired vision. In: *ERCIM workshop on user interfaces for all*. Springer, Berlin, pp 3–22
- Mace R (1988) *Universal design: housing for the lifespan of all people*. The Center for Universal Design, North Carolina State University, Washington, DC
- Mynatt ED, Weber G (1994) Nonvisual presentation of graphical user interfaces: contrasting two approaches. In: *Proceedings of the SIGCHI conference on human factors in computing systems 1994 (CHI'94)*. ACM, New York, pp 166–172
- Newell AF, Gregor P (2000) “User sensitive inclusive design” – in search of a new paradigm. In: *Proceedings of the 2000 Conference on Universal Usability*. ACM, New York, pp 39–44
- Newell AF, Gregor P, Morgan M, Pullin G, Macaulay C (2011) User-sensitive inclusive design. *Univ Access Inf Soc* 10(3):235–243
- Nielsen J, Molich R (1990) Heuristic evaluation of user interfaces. In: *Proceedings of the SIGCHI conference on human factors in computing systems 1990 (CHI'90)*. ACM, New York, pp 249–256
- Norman DA (1983) Design principles for human-computer interfaces. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 1983*. ACM, New York, pp 1–10
- Petrie H, Gill J (1993) Current research on access to graphical user interfaces for visually disabled computer users. *Eur J Special Needs Edu* 8(2):153–157
- Petrie H, Kheir O (2007) The relationship between accessibility and usability of websites. In: *Proceedings of the SIGCHI conference on human factors in computing systems 2007 (CHI 2007)*. ACM, New York, pp 397–406
- Petrie H, Morley S, Weber G (1995) Tactile-based direct manipulation in GUIs for blind users. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 1995 (CHI'95)*. ACM Press, New York
- Power C, Freire A, Petrie H, Swallow D (2012) Guidelines are only half of the story: accessibility problems encountered by blind users on the web. In: *Proceedings of the SIGCHI conference on human factors in computing systems 2012 (CHI 2012)*. ACM, New York, pp 433–442
- Power C, Petrie H, Swallow D, Murphy E, Gallagher B, Velasco CA (2013) Navigating, discovering and exploring the web: strategies used by people with print disabilities on interactive websites. In: *IFIP conference on human-computer interaction*. Springer, Heidelberg, pp 667–684
- Pullin G (2009) *Design meets disability*. MIT press, Cambridge, MA
- Rømen D, Svanæs D (2012) Validating wcag versions 1.0 and 2.0 through usability testing with disabled users. *Univ Access Inf Soc* 11(4):375–385
- Savva A (2018) *Understanding accessibility problems of blind users on the web*. PhD thesis, University of York
- Savva A, Petrie H, Power C (2015) Comparing concurrent and retrospective verbal protocols for blind and sighted users. In: *Human-Computer Interaction*. Springer, Heidelberg, pp 55–71

- Shneiderman B (2000) Universal Usability. *Communications* 43(5):84–91
- Stephanidis C (2001) User interfaces for all: new perspectives into human-computer interaction. In: *User interfaces for all—concepts, methods, and tools*, vol 1. Lawrence Erlbaum Associates Inc, Mahwah, pp 3–17
- Theofanos MF, Redish JG (2003) Bridging the gap: between accessibility and usability. *interactions* 10(6):36–51
- Watanabe T (2009) Experimental evaluation of usability and accessibility of heading elements. *Disabil Rehabil Assist Technol* 4(4):236–247
- Wobbrock JO, Aung HH, Myers BA, Lopresti EF (2005) Integrated text entry from power wheelchairs. *Behav Inform Technol* 24(3):187–203

# Chapter 11

## Deep Subjectivity and Empathy in Virtual Reality: A Case Study on the Autism TMI Virtual Reality Experience



Jonathan Weinel, Stuart Cunningham, and Jennifer Pickles

**Abstract** The Autism Too Much Information (TMI) Virtual Reality Experience is a virtual reality (VR) application produced by The National Autistic Society (NAS) as part of an awareness campaign. The design of the application creates a short narrative simulation from a first-person perspective, which conveys aspects of what it may be like for a child on the autistic spectrum to experience a stressful situation precipitated by environments with ‘too much information’. The application is part of a recent trend in VR and 360-degrees video, to create simulations of subjective experience, as a means to generate empathy. Yet the success of such tools depends significantly on how well sound and graphics can be used to communicate such experiences in a meaningful way. In this article, we provide a case study of the Autism TMI Virtual Reality Experience, as a means to unpack design issues for these simulations. Through an expert analysis and pilot study of user experience, we propose three distinct forms of subjective first-person simulation that may be produced in virtual reality. We argue that the Autism TMI Virtual Reality Experience exemplifies the third of these: ‘deep subjectivity’, which may lead to an improved sense of empathy by representing various aspects of multimodal perception and emotion. However, our study also suggests that VR may offer limited benefits over 360-video for generating a sense of empathy.

---

J. Weinel (✉)  
Aalborg University, Aalborg, Denmark  
e-mail: [weinel@hum.aau.dk](mailto:weinel@hum.aau.dk)

S. Cunningham  
Wrexham Glyndŵr University, Wrexham, UK

J. Pickles  
University of Southampton, Southampton, UK

## 11.1 Introduction

Autism is a “lifelong development disability that affects how a person communicates with and relates to other people, and how they experience the world around them” (NAS 2016f). Among the features of the condition, autistic people<sup>1</sup> may experience increased or reduced sensitivity to certain sensory stimuli, which can cause everyday situations and public spaces to become extremely stressful and hard to deal with. ‘Too Much Information’ (TMI) is an autism awareness campaign produced by The National Autistic Society (NAS; a UK-based autism charity), which aims to highlight this issue, and the challenges faced by autistic people in public spaces (NAS 2016e).

As part of the TMI campaign, the Autism TMI Virtual Reality Experience is an application that includes a short virtual reality (VR) film, which seeks to convey the experience of someone on the autistic spectrum in a shopping centre (NAS 2016c). The application aims to give an impression of the overwhelming sensory experience that environments such as this can induce for an autistic person, while also highlighting the negative reactions that an autistic person having a ‘meltdown’ or ‘shut-down’ can produce from the general public (NAS 2016b). The application provides this impression through means of a first-person point-of-view (POV) narrative, which makes use of graphical and sound design techniques to convey the subjective experience of an autistic child. In doing so, the film aims to promote an awareness of autism, and foster a greater understanding from the public, in order to reduce the negative reactions that further confound such situations.

The Autism TMI VR Experience can be seen as part of an emerging field of work in audio-visual media and VR, which seeks to use first-person POV as a means to simulate consciousness states or generate empathic tools. For example, the 360-degrees film *Clouds Over Sidra* (Arora and Pousman 2015) tells the story of Sidra, a 12-year old girl who has spent the last 18 months at the refugee camp Zaatari of Jordan. Other examples of work in this field include those discussed at Virtual Futures Salon: Electronic Empathy (Mason et al. 2016), such as *The Machine to Be Another* (BeAnotherLab 2016): a VR experiment in which two participants swap perspectives while performing the same physical actions, leading to a sense of exchanged embodiment; or Jane Gauntlett’s *In My Shoes* (Gauntlett 2016), an arts project aiming to produce empathy by recreating life experiences. Electronic simulations of auditory hallucinations are also being developed for the purposes of treating schizophrenia (Craig et al. 2015); while in entertainment, simulations of altered states of consciousness produced by drug experiences or psychosis are being used in an increasing number of video games (Cunningham et al. 2016). These works are united through the use of audio-visual media and technologies such as VR as a means to convey conscious states in the first-person, using forms of telepresence to place the user inside the sensorium of another. The Autism TMI VR Experience can be seen in the context of such work, and seems to be one of the first examples of VR

---

<sup>1</sup>There is some debate regarding terminology to describe autism; see Kenny et al. (2015).

being used to communicate the experience of autism; though other interactive systems do exist, such as Hermann, Yang and Nagai's *EmoSonics* (2016), which seeks to communicate the emotions of autistic people through electronic sound synthesis.

Applications such as the Autism TMI VR Experience may provide powerful forms of communication. Yet the success of such media rests significantly on the extent to which they can make use of graphics and sound to provide effective and meaningful forms of communication. In attempting to provide simulations of consciousness, designers must overcome challenges in the use of graphics, sound and interactivity to convey the conscious experience of the subject. Since the 1980s, discussions of realism in computer graphics have tended to focus on the photo-realistic accuracy of graphical representations that reproduce shading, shadows and lighting in ways that approximate patterns of light entering the eye (Amanatides 1987). In more recent studies such as Borg et al. (2014), these features continue to be used as benchmarks when assessing realism. Graphics and sound are key features of a VR system that may contribute towards its immersive capabilities, by generating a virtual environment that engulfs the senses, leading to provide the user with a sense of 'presence' within it.<sup>2</sup> Other related work such as Slater (2009) have sought to expand upon this area, by isolating 'place illusion' (PI) as presence, or the feeling of 'being there'; and 'plausibility illusion' (Psi) for the belief that what is occurring is actually happening. However, it seems that reproducing incoming visual and auditory stimuli that a person would see and hear from an external environment, if they were 'in the shoes' of another person, may not necessarily be sufficient to meaningfully communicate the subjective experience of what it is actually like to be that person. For instance, with regards to autism, the perceptual experience of incoming sensory information is significantly different than it is for someone who is not autistic, and thus designers must find ways to communicate this. As we shall see in this chapter, this may involve the use of various audio-visual techniques to suggest selective attention, non-aural or non-visual sensory experience, or emotional states. In this way, designers can not only create an impression of receiving similar visual or auditory stimuli as another person in a given location, but also begin to relate aspects of subjective perception that such stimuli may invoke. Recent applications such as the Autism TMI VR Experience seek to do this, yet to date we are not aware of any studies that evaluate the user experience of these designs.

In order to further our understanding of how VR applications may communicate subjective experiences, this chapter presents a case study on the Autism TMI VR Experience. We begin with an introduction to autism, which explores how the condition affects sensory experience and may cause high levels of stress in certain environments. Following this, we explore how the Autism TMI VR Experience communicates such experiences. First, we provide an independent 'expert analysis' of the application, which describes the narrative, graphical and sound design

---

<sup>2</sup>The definitions of 'presence' and 'immersion' used here follow those described by Slater and Wilbur (1997), where presence describes the feeling of 'being there' facilitated by digital media that has 'immersive' technical features. For a further discussion of presence in VR see also Sanchez-Vives and Slater (2005).



techniques that are used in detail.<sup>3</sup> We then discuss a quantitative pilot study that was undertaken to evaluate the user experience of the Autism TMI VR Experience. This study explores the extent to which selective attention, sensory experience and emotion are communicated effectively by the video, comparing both the VR and 360-degrees video versions. Through this, we reveal how participants may gain a sense of what it is like to be a person on the autism spectrum, and the extent to which VR supports this when compared with a 360-degrees video viewed from a standard LCD computer screen. Building upon the outcomes of this study, we propose three main types of first-person POV simulation, including simulations of ‘deep subjectivity’ that take into account multi-modal perception. We propose that ‘deep subjectivity’ more effectively communicates the perceptual experience of what it is like to be another person, and leads to more effective empathic tools. Finally, we conclude that the principles explored in this chapter are applicable not only for autism, but also other forms of first-person POV simulation.

## 11.2 Autism and Sensory Experience

### 11.2.1 *What Is Autism?*

Autism Spectrum Disorder (ASD) as defined by the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) (American Psychiatric Association 2013) is a developmental disorder. It is characterised by impairments in social communication and interaction, and restricted, repetitive behaviours, interests and activities, which includes sensory sensitivities. Autism is a lifelong condition and people diagnosed as autistic will have presented with difficulties in these areas since early childhood. The condition is being increasingly recognised, and it is now thought that more than 1 in 100 people may have autism (Brugha and Spiers 2012). It is a spectrum disorder, and the presentation is highly individualised in each area of impairment. Among the features that autism presents, differences in experiencing and processing sensory stimuli have been consistently recognised and reported since the condition was first described (Kanner 1943; Asperger 1944). Indeed, some of the earliest research stated that this should be a diagnostic feature of autism (Creak 1961; Wing 1969), however sensory sensitivity was not included among the diagnostic criteria until the DSM-5 in 2013. Recently research has also considered emotional response and regulation in people on the autism spectrum (Mazefsky et al. 2013; Berkovits et al. 2017). Despite this being a key area, particularly when designing intervention practices for behavioural responses, this area of research is in its relative infancy and further work is needed to determine the link between emotional regulation and behavioural response in autism. Increased awareness of

---

<sup>3</sup>Note that since our analysis is carried out independently from the design process, we provide a critical analysis of the artefact, but do not provide further commentary on the design rational, prototypes, or other aspects of its creation.

autism is important within society, as this may lead to improved diagnosis and support, and hence tools such as the Autism TMI VR Experience are provided in an attempt to highlight features of the condition to the general population.

### ***11.2.2 Sensory Sensitivities***

Traditionally research has focused on the impact of autism on seven senses: gustation (taste), olfaction (smell), tactility (touch), proprioception (body awareness), the vestibular system (balance), vision and hearing (Bogdashina 2010). Through first-hand accounts from autistic authors (Williams 1994; Grandin 2009) it has become evident that these sensory differences are apparent from the early years and can result in an altered perception of an environment from that of a neuro-typical person (Bogdashina 2016). Sensory differences include both hypersensitivity and hyposensitivity to certain stimuli. Hypersensitivity may result in strong adverse reactions to certain sensory information, for example, finding certain sounds, such as a clock ticking, too loud to bear. Alternatively, an individual may experience hyposensitivity to stimuli, leading to dampened responses or not attending to certain sensory information (Pellicano 2013). These differences in sensory perception are not equivalent across individuals, and may also vary over time; in different environments; and due to increased or decreased stress levels.

Oscillations between hypersensitivity and hyposensitivity, across any or all of the senses, can have a distressing or overwhelming impact upon the individual and can lead to feelings of disorientation and confusion. The constant bombardment of environmental sensory stimuli can be difficult to process, and may cause ‘melt-downs’ or ‘shutdowns’. Such sensitivity may also cause the individual to exhibit sensory-seeking behaviours, often referred to as self-stimulating behaviours, or ‘stimming’. These may include hand-flapping, rocking and humming, among others (NAS 2016d). In some cases, these behaviours can be harmful; for example, head banging to seek an increased tactile response can obviously have negative consequences. Additionally, these behaviours can cause consternation amongst the general public, which in turn can create stressful social situations. Meanwhile, sensory stimuli can also have a negative impact on an autistic person’s ability to focus (Baranek 2002), and thus careful consideration of an environment must be taken to adapt it in accordance with the needs of an individual (Pfeiffer et al. 2005; Kinnealey et al. 2012; Guldberg 2010).

In response to the need to design environments that address the needs of autistic people, the NAS’s SPELL (Structure, Positivity, Empathy, Low arousal and Links) framework emphasises the need to create spaces that encourage low arousal levels to mitigate the possibility of sensory overstimulation, which can result in melt-downs or shutdowns. In some autism specific support services, environments such as sensory rooms have been specifically designed to deprive certain senses and enhance others, in order to reduce the anxiety that can occur from experiencing too much sensory information. For instance, a sensory hub may be soundproofed but

filled with lights that change colour and brightness; or they may be completely dark but contain strong smelling materials. These sensory hubs are adapted to suit individual requirements, based on an individual's sensory profile. Recently some research has been carried out exploring the use of interactive audio-visual technologies to aid the design of these environments and create multi-sensory instruments (Capellen and Andersson 2016).

### **11.3 An Analysis of the Autism TMI Virtual Reality Experience**

In order to increase awareness of sensory sensitivities that may be experienced by autistic people within public spaces, the NAS has created the TMI awareness campaign. As part of this, the Autism TMI VR Experience was created in collaboration with the design company Happy Finish. The application aims to demonstrate to users how a fairly standard environmental setting may appear to a person who has hypersensitivities to certain stimuli. It also demonstrates how the effects of this can build over time when appropriate support or reasonable accommodations are not put into place. The design of the video is informed by first-hand accounts from autistic people and individuals affected by autism such as parents and carers, and demonstrates the cumulative effect of adverse sensory information and the resulting 'melt-down' or 'shutdown'. The report the video was based on utilised collective responses from focus groups, online surveys, polls, and qualitative analysis of personal accounts. This research collated the views of over 7500 people and was conducted on behalf of the NAS by Breathe Research, YouGov and nfpSynergy (NAS 2015). The video is available to view through both a mobile application for use with VR-cardboard; and as a 360-degrees YouTube video (NAS, 2016a), the latter of which received nearly 200,000 views in the first 3 months since its launch.

In order to consider the design of Autism TMI VR Experience, the following subsections give a detailed account of the structural design and narrative; visual design; and sound design techniques that are used. This provides our 'expert analysis' of how these can be understood to relate to the conscious experience of a person on the autistic spectrum, and the type of effect that the audience may experience. This appraisal was carried out independently from the design process, and utilises an approach of critical analysis to evaluate the salient features of the digital media artefact as an artistic composition.

#### ***11.3.1 Narrative and Structure***

For the VR-cardboard version of the Autism TMI VR Experience, the menu page states:

For autistic people, the world can be a scary place. Everyday sounds, lights, colours, they can all be too overwhelming. We invite you to experience this for yourself and enter a world of Too Much Information. #AutismTMI.

A menu then allows the user to either visit the NAS website, or to view the video. The 360 YouTube version does not include these menus, but goes straight into the video part, which is the same on both.

The narrative of the video begins in a shopping centre, as a woman purchases a ticket. The film takes a first-person POV from a child (Alex's) perspective, as he waits while the woman (who we may presume is his mother) operates a ticketing machine. While he waits, the film shows us various sensory inputs that are a product of the environment. First we notice the glare of flickering lights, background noise and the muzak of the shopping centre. Next, we hear footsteps as a woman passes by. As another woman passes, she appears to step in something sticky that is spilt on the floor, and wipes her foot for a few moments before moving on. The scene begins to intensify as more people walk past, including a man holding a large collection of bright balloons; while in the background two ladies appear to be testing out some perfume. Alex's breathing quickens as he begins to panic, while his mother tells him to calm down as she finishes operating the ticketing machine, which appears not to be functioning as intended. Alex begins to panic as more sensory events unfold, such as a woman cleaning the floor. His stressed mother attempts to calm him, but by this point Alex is overcome by the situation and unable even to hear her clearly, as the scene overwhelms him. His breathing intensifies and the scene turns to black as he experiences a 'meltdown' or 'shutdown'. In the next scene, as he begins to calm down, we see a view of his mother and the car park, as his mother tells him to keep breathing, that everything is fine and that it is time to go home.

In the closing of the film, Alex delivers the message "I'm not naughty, I'm just autistic, and I get too much information", while text displays the web address of the NAS, encouraging the viewer to learn more and sign up to charity.

### ***11.3.2 Visual Design***

As noted, the visual design of Autism TMI VR Experience assumes a first-person POV. This is practically achieved using a head-mounted VR camera, which presents the field of vision available from Alex's eye-view. This allows the visual material to provide stereoscopic 3D, which includes a sense of depth; and allows 360-degree head tracking, allowing the viewer to look around from Alex's viewpoint. However, the film not only presents a first-person POV from the location Alex is in, but also does so in a way that seeks to describe the subjective experience of Alex using various visual techniques. As Alex blinks, the screen flickers black, and various visual effects such as graphical filters are used to describe his experience of the environment, or highlight specific sensory elements which capture his attention. This technique of subjective, first-person POV, which includes blinking and the use of digital

effects processes, is notably used elsewhere in feature films such as *Enter the Void* (Noé 2009) to describe the subjective experiences and hallucinations of a protagonist. However, in this case, visual effects are utilised specifically in order to give an impression of an autistic person in a situation where too much sensory information is causing panic.

Considering the visual effects used to create this impression in more detail, brightness parameters are temporally boosted to create an oversaturated flickering image, which suggests the sensory overload caused by the lights (0:14).<sup>4</sup> Meanwhile, a temporal blurring filter effect suggests the discomfort or jarring effect that the environment is causing Alex to experience (0:19). As a woman walks across the scene, we see bright ripples of light emerging from her footsteps, giving an impression of how prominent and powerful the sound of her heels are for Alex (0:33). As people walk past, such as the man with the balloons, a filter is used to create an echo of the image, lending an impression of disorientation (0:42). Following this, a cloud of purple smoke is emitted from the characters in the background, who seem to be testing some perfume; the purple particle effect providing a visual means through which to describe the overpowering smells that Alex may detect (0:48). Alex attracts an accusing glance from a young man passing by in a red coat (0:55). Throughout this section we can also see Alex playing with a dinosaur toy in his hands (visible by looking downwards in VR), reflecting the need for tactile sensory stimulation.

It is notable that some of these effects provide a synaesthetic visual impression of the sensory experience across the modalities, since light is used to highlight footsteps, and purple smoke to draw our attention to powerful experiences of scent. Some of these effects are provided using ‘keyframing’, a visual effects technique that enables the animation of graphical parameters. As the film progresses, keyframed effects such as these are used with increasing frequency and intensity, to give a visual impression of the rising sensory overload that Alex is experiencing. As the film continues to show Alex’s increasing sensory overload, a keyframed filter is ultimately used to blur the visual field and cause a ‘tunnel-vision’ effect that gradually closes in upon him as he enters a ‘meltdown’ or ‘shutdown’ (1:02–1:27). This technique is suggestive of the sensory overload Alex is experiencing, where he is completely overwhelmed by ‘too much information’, and is unable to cope with any individual elements successfully; instead he effectively closes down. Notably in the 360-degree presentations of *Autism TMI VR Experience*, this visual effect constrains the viewer’s field-of-vision, as darkness closes in, ultimately plunging the audience into blackness in a manner that some may find scary.

In the final section (1:35), a transition effect returns the image to a normal view of a car park outside the shopping centre, which is unaltered by the various graphical filters. The absence of visual effects, taken in contrast with the previous scene of the shopping mall suggests a calmer environment, which is less overwhelming for Alex.

---

<sup>4</sup>For convenience, the time references used here and throughout relate to the YouTube version of *Autism TMI VR Experience* (NAS 2016a).

### 11.3.3 *Sound Design*

The soundtrack of Autism TMI VR Experience describes the subjective aural experience of Alex, with the additional use of music to metaphorically relate his mood. At the beginning of the film we hear the background noise of the shopping centre, the footsteps of a person off-screen, and the voice of Alex's mother. As the flickering lights cause a swell in brightness, we hear a whooshing sound (0:14), while the diegetic background muzak<sup>5</sup> of the shopping centre swells. We also hear the sounds of Alex's quickening breathing, and the rustling as he plays with the toy dinosaur that he holds in his hands. As the woman enters the scene, who steps in something that is spilt on the floor, we hear a prominent squelching sound (0:24). The sound of Alex's breathing continues to quicken indicating his increasing stress, while more synaesthetic whooshing sounds, footsteps, his mother's voice, and the swinging of a metal door are heard (0:40). As an alarm sound is heard (0:43), the noise levels rapidly intensify, as we hear a loud spraying sound that accompanies the use of perfume in the background, and a high frequency ringing sound. Further clattering is heard in the background and more squelching sounds as a woman begins to clean up the floor spillage, and Alex begins to gasp as he becomes overwhelmed (1:00).

At this stage, the dialogue from Alex's mother become muffled; an effect that is achieved with a low-pass filter (1:00). As she drops her keys, the sound is extended using a time-stretching effect, creating a sustained high frequency drone (1:05). During the phase where Alex's tunnel vision begins to close in, we continue to hear his panicked breathing, amidst the background cacophony of dissonant noise textures<sup>6</sup> and high frequency drones; the noise becoming louder as the amplitude of his mother's voice decreases (1:02–1:27). In this way, the audio track reflects the shift of attention from the discernible voice of his mother as she attempts to calm him, to the wall of noise that gives an impression of Alex's anxious mood and experience of sensory overload.

During the transition phase where the visual field goes dark, we hear only Alex's stressed breathing, which gradually subsides as the scene fades into the car park (1:37). The ambient sounds of the car park fade in, together with the voice of Alex's mother, creating a mix of amplitude between parts that seems relatively normal, reflecting Alex's experience of sound, which is no longer overwhelming, but has now become manageable. As end credits and message are displayed, we hear Alex's voice as the narrator, who delivers the message "I'm not naughty, I'm autistic, and I just get too much information".

---

<sup>5</sup>Film sound can be considered in terms of diegetic sounds that relate directly to the world presented within the story, and non-diegetic audio such as music that has no explicit basis within that world (Sonnenschein 2001). Here 'diegetic background muzak' refers to the music of the shopping centre, in contrast with the non-diegetic music that is also used to reflect Alex's mood during the video.

<sup>6</sup>These 'noise textures' could be considered as a form of music, and provide a means through which to communicate emotion. For a further discussion of how dissonant musical features may represent unpleasant emotions, see (Gabrielsson and Lindstrom 2012).

## 11.4 A User Experience Study of the Autism TMI Virtual Reality Experience

### 11.4.1 Pilot Study Design and Methodology

A pilot study was conducted to investigate perceptions of the VR-cardboard and the 360 YouTube video versions. A convenience sample of twenty ( $n = 20$ ) participants were recruited from a University campus and consisted of undergraduate and post-graduate students, faculty staff and administrative staff. Of the 20 participants, 1 was female and 19 were male. Each participant was briefed and randomly assigned to either view the VR video ( $n_{VR} = 9$ ) or the 360 video ( $n_{360} = 11$ ). Both groups of participants wore headphones during the experiment.

Participants experiencing the 360 version of the video were sat at a desk in front of a computer (an Apple iMac with a built-in 21-inch display screen, wireless keyboard and wireless mouse), where the YouTube video was maximised to fill the screen. They were then briefly introduced to the 360 video system by one of the researchers, and given instruction of how to change their viewpoint in the video should they wish to do so.

Participants viewing the VR version of the video were asked to stand to provide them space to move freely and instructed that they were required to hold the headset to their eyes for the duration of the video. As with participants in the 360 group, they were briefed about the use of the VR-cardboard headset and the ability to move their head to change their viewpoint in the video.

An online, web-based, survey was used to collect participants' responses, which was completed immediately after watching the video. The total process took approximately 10 min per participant. The principal aims of the pilot study were to determine:

1. The efficacy of both videos in portraying Alex's subjective experiences to the viewer.
2. If there is a difference in the viewer perception of Alex's experience in the video between the VR and 360 formats.

The survey consisted of 9 questions relating to sensory features and 3 questions relating to their overall impressions of the video and their sense of relationship to Alex. These 12 questions were mandatory and presented the participant with a statement, to which they were asked to indicate their level of agreement using a Likert scale with the following points: 1. Strongly disagree; 2. Disagree; 3. Neither disagree or agree; 4. Agree; and 5. Strongly Agree. In addition, there were two open-ended questions at the end of the survey where participants could provide free text responses, which were optional.

The sensory aspects of the video section of the survey explored how effective participants found the video was at highlighting various aspects of Alex's subjective experience. The questions also correspond with the seven sensory aspects considered important in autism research: gustation (taste), olfaction (smell), tactility (touch), proprioception (body awareness), the vestibular system (balance), vision and hearing.

**Table 11.1** Sensory concepts and corresponding survey statements

Sensory concept	Statement
Attention	The video highlighted to me what caught the boy (Alex’s) attention
Gustation	The video highlighted to me what the boy (Alex) was tasting
Olfaction	The video highlighted to me what the boy (Alex) was smelling
Tactility	The video highlighted to me what the boy (Alex) was touching
Proprioception	The video highlighted to me the boy’s (Alex’s) awareness of his body
Vestibular	The video highlighted to me the boy’s (Alex’s) sense of physical balance
Vision	The video highlighted to me what the boy (Alex) was seeing
Hearing	The video highlighted to me what the boy (Alex) was hearing
Emotion	The video highlighted to me what the boy (Alex) was feeling

In addition, since selective attention mechanisms towards sensory inputs are important, we look at attention. Also, as sensory experience leads Alex to have an emotional response of high arousal and distress, we also explore emotion specifically.

Considering the sensory aspects of the video, the following null and alternate hypotheses were investigated:

H<sub>0</sub>: There is no difference in viewer perception of sensory aspects in the video between the VR-cardboard and 360 YouTube formats.

H<sub>1</sub>: There is a significant difference in viewer perception of sensory aspects in the video between the VR-cardboard and 360 YouTube formats.

Sensory concepts were explored with participants in both groups by asking for their level of agreement with the corresponding statement as shown in Table 11.1.

In addition, we also seek to extract participants’ overall impressions of the video, which explore whether the video makes the user feel as if they are in the same location as Alex (what it is like to be somewhere) and/or if they gain a sense of what it is like to be Alex (what it is like to be someone). The distinction between these two is important, since the video seeks to convey Alex’s subjective experience of the shopping centre, beyond simply replicating the environment.

Regarding the overall impressions of the video and the participants’ sense of relationship towards Alex, the following null and alternate hypotheses are investigated:

H<sub>2</sub>: There is no difference in overall impressions of the video between the VR-cardboard and 360 YouTube formats.

H<sub>3</sub>: There is a significant difference in overall impressions of the video between the VR-cardboard and 360 YouTube formats.



**Table 11.2** Overall impression concepts and corresponding survey statements

Impression concept	Statement
Somewhere	The video made me feel that I was in the shopping centre
Someone	I felt like I was the boy (Alex) in the video
Empathy	Overall, the video made me feel empathy towards the boy (Alex)

The overall impressions concepts were collected via the study, through asking participants to indicate their level of agreement with the statements shown in Table 11.2.

Finally, at the conclusion of the Likert response questions, qualitative data was collected using open responses, to ask which features participants found most effective in conveying an experience of autism, or to allow them to express other comments about the video they watched. The two questions asked of participants were:

- What features (if any) of the video were most effective in giving you an experience of Autism?
- Do you have any other thoughts, impressions or comments about the video?

## 11.4.2 Results and Analysis

### 11.4.2.1 Sensory Concepts

The results obtained from the sensory concepts questions are presented using a polar plot, as shown in Fig. 11.1. It is seen from this set of responses that there is a general consistency between the two types of video. Emotion, hearing, vision, and attention all receive a consensus of agreement responses, with the mean of each group's response being  $\geq 4$  (agree). Proprioception receives the closest to a neutral response, with the mean response from the VR-cardboard group being 3.22 and the 360 YouTube group being 3.00, followed by vestibular, which has a mean rating of 3.00 from the VR-cardboard group and 3.45 from the 360 YouTube group. Olfaction and tactility indicate responses towards the disagreement end of the Likert scale, with the respective means for the VR-cardboard group being 1.89 and 2.22, and the 360 YouTube being 2.27 and 1.91. The remaining sensory concept of gustation has the most consistent disagreement concept, with both ratings being  $< 2$  (Disagree).

A one-way MANOVA test showed that there was no statistically significant difference in sensory responses based on the type of video (VR or 360),  $F(9, 10) = 0.29$ ,  $p > 0.05$ ; Wilk's  $\Lambda = 0.793$ , partial  $\eta^2 = 0.21$ . From these results, we maintain the null hypothesis  $H_0$ : There is no difference in viewer perception of sensory aspects in the video between the VR-cardboard and 360 YouTube formats.




Fig. 11.1 Sensory concepts between VR-cardboard and 360 YouTube Videos




Fig. 11.2 Overall impression concepts between the VR-cardboard and 360 YouTube Videos

### 11.4.2.2 Overall Impression Concepts

These are investigated to determine which features came across to participants viewing the video. Results for these three questions are presented in the bar graph shown in Fig. 11.2. Both of the experimental groups (360 YouTube and

VR-cardboard) show strong levels of agreement with respect to the somewhere and empathy concepts, with mean responses across both groups being  $>4$  (Agree): the somewhere mean responses were 4.22 for the VR-cardboard and 4.18 for the 360 YouTube; the empathy mean responses were 4.56 for the VR-cardboard and 4.55 for the 360 YouTube. The responses for the someone concept are weaker, means for VR-cardboard of 4.00 and 360 YouTube of 3.82, but are towards the Agree end of the scale.

A one-way MANOVA test showed that there was no statistically significant difference in the sense of being another person experienced, based on the type of video (VR or 360),  $F(3, 16) = 0.57, p > 0.05$ ; Wilk's  $\Lambda = 0.99$ , partial  $\eta^2 = 0.011$ . From these results, we maintain the null hypothesis  $H_2$ : There is no difference in overall impressions of the video between the VR-cardboard and 360 YouTube formats.

### 11.4.2.3 Qualitative Responses

Finally, we consider the qualitative, open-ended feedback that was provided in the questionnaire.<sup>7</sup> In response to the query regarding features of the video that were effective in giving an experience of autism, participants referred to a variety of audio-visual features in relation to the communication of sensory experience. For instance, 11 of 18 responses commented on audio features, highlighting the breathing sound effects in particular, while 5 of 18 responses discussed visual features. The communication of emotions such as stress or anxiety was described by 5 of 18 responses. Meanwhile, in response to the invitation for other general comments, participants noted the possible application of the video in educational settings, technical challenges they faced in using the VR cardboard, and their own emotional reactions to the video.

## 11.5 Unpacking Subjectivity in First-Person POV Audio-Visual Media

The results of the user experience study suggest that the participants found the Autism TMI VR Experience provides an effective means through which to communicate the subjective experience of an autistic person, leading to a sense of empathy: both groups in the study indicated high levels of agreement with the empathy overall impression  $\bar{x}_{VR} = 4.56$  and  $\bar{x}_{360} = 4.55$ . This is achieved through the use of a first-person POV representation. However, this does not only communicate Alex's incoming visual and aural sensory inputs through graphics and sound, since it also

---

<sup>7</sup>In the interests of brevity, the description here is provided as a concise overview. In a future development of this study, a more in-depth qualitative analysis could be included utilizing formal methods of thematic analysis.


uses these to provide an impression of other multimodal senses and emotions: both groups in the study indicated high levels of agreement with the attention sensory concept,  $\bar{x}_{VR} = 4.00$  and  $\bar{x}_{360} = 4.09$ , and the emotion sensory concept,  $\bar{x}_{VR} = 4.56$  and  $\bar{x}_{360} = 4.73$ . Through this, the first-person POV simulation does not only show what it is like to be somewhere else,  $\bar{x}_{VR} = 4.22$  and  $\bar{x}_{360} = 4.18$ , but also to be someone else,  $\bar{x}_{VR} = 4.00$  and  $\bar{x}_{360} = 3.82$ . This is an important distinction, which we will now further explore by unpacking the notion of first-person POV representations, in order to reveal three possible design approaches. Of these three varieties of first-person POV, the Autism TMI VR Experience utilises the third, leading to simulations of ‘deep subjectivity’ that may yield improved levels of empathy.

### 11.5.1 Layers of Subjectivity

In order to unpack different forms of first-person POV, let us consider the simplified model of sensory perception shown in Fig. 11.3. During visual and auditory perception, incoming sensory information is received via an external environment (stage 1). In accordance with Broadbent (1958) and Treisman’s (1960) theories of attention, cognitive processes of attention filter or attenuate incoming sources of information (stage 2). Attended sources reach perceptual awareness and are passed for higher level processing; while unattended sources are attenuated, and receive more limited processing (stage 3). Research regarding the ‘McGurk effect’ (McGurk and MacDonald 1976) and neuroimaging studies of the auditory cortex (Callan et al. 2003; Calvert et al. 1997) show that in the second stage auditory cortex, higher-level processing and the multimodal integration of visual and auditory information occurs. Similarly, visual and auditory perception may invoke memories and emotions. Here these associative, multi-modal processes, which involve other cognitive systems, are collectively summarised as ‘multimodal integration’ (stage 4).

### 11.5.2 Varieties of First-Person POV

The four stages of the simplified model of audio-visual perception shown in Fig. 11.3 can be used to suggest multiple approaches for the design of first-person POV media that represents subjective experience. In what follows, we identify



**Fig. 11.3** Simplified stages of perception, which suggest corresponding categories of first-person POV audio-visual media

**Table 11.3** Representational content included in three varieties of simulation

Representational features	Varieties of first-person POV simulation		
	Environmental reconstruction	Subjective attention	Deep subjectivity
Incoming external audio-visual information	✓	✓	✓
Attentional filters		✓	✓
Sonification/visualisation of other sensory modalities: gustation, olfaction, tactility, proprioception and the vestibular system			✓
Sonification/visualisation of emotion			✓

three distinct categories of first-person POV media that correspond with these perceptual stages, describing simulations that take into account progressively higher or ‘deeper’ levels of processing. The features of each of these varieties are summarised in Table 11.3.

### 11.5.2.1 Environmental Reconstruction

First-person POV simulations that provide ‘environmental reconstructions’ replicate stimuli similar to that experienced by the subject, without taking account of subjective attention or perception. These use ‘Incoming Sensory Information’ (stage 1) as a basis for representation (Fig. 11.3). For example, we may conceive of videos that use a head-mounted camera and microphones, to reproduce the incoming visual and auditory stimuli that a person may receive in a given situation. For audio-visual media, these simulations are usually limited to reproducing sensory information from the visual and auditory domains, although in principal artificial stimulation of olfactory (Spencer 2006; Murray et al. 2016) or gustatory (Ranasinghe et al. 2012) senses could also be used. We see examples of these videos regularly in the form of the Internet videos produced by rock climbers and mountain bikers, who show us what they see and hear as they engage in their respective sports. These simulations do not necessarily show us which visual or auditory entities within the scene the subject is directing his or her attention towards, except perhaps through physical motion where the subject turns his or her head to face the object of attention.

### 11.5.2.2 Subjective Attention

First-person POV simulations that provide ‘subjective attention’ utilise visual or auditory techniques to reflect the perceptual mechanism of attention, by attenuating unattended sources. These simulations use ‘Attention Filters’ (stage 2) and ‘Perceptual Awareness’ (stage 3) as their basis (Fig. 11.3), since they represent the sensory experience as it is perceived after attentional filtering. In sound this can be achieved by adjusting the amplitude or applying filters to unattended sources, to attenuate their levels and reduce their prominence in the mix. Conversely, attended

sources can have their amplitude levels boosted, or frequency equalisation processes can be used to make them seem more clear or bright, so that they ‘stand out’ more clearly. We see this approach in aspects of the Autism TMI VR Experience, since certain elements, such as the sound of shoes squelching through the floor spillage, are given relatively higher amplitude levels. Conversely, towards the end of the film, the sound of Alex’s mother’s voice is processed with a low-pass filter, creating a muffled sound that reflects reduced awareness. Using techniques such as these, it is possible for simulations to take the perceptual mechanism of attention into account, thus providing simulations of attended sources.

### 11.5.2.3 Deep Subjectivity

First-person POV simulations that provide ‘deep subjectivity’ utilise the audio-visual medium to reflect not only the perceptual experience of visual and auditory information, but also aspects of perception that are non-aural or non-visual. As described by ‘Multimodal Perception’ (stage 4), higher level processing involves the integration of visual and auditory information, and invokes other cognitive systems such as memory (Atkinson and Shiffrin 1971) and emotions (Russell 1980) (Fig. 11.3). Sensory experience of gustation, olfaction, tactility, proprioception, and the vestibular system are not primarily visual or auditory. Nonetheless, it is possible to represent these using techniques, such as sonic or graphical metaphors. Indeed, our ‘expert analysis’ noted that this is used in various aspects of the Autism TMI VR Experience, although our user study also revealed that participants did not strongly detect these features. Similarly, our ‘expert analysis’ suggested that the video might communicate a sense of emotion through features such as its narrative and the use of dissonant noise textures that sonify Alex’s sense of emotional anxiety. Since emotion was strongly recognised by the participants in both the quantitative results, ( $\bar{x}_{VR} = 4.56$ ;  $\bar{x}_{360} = 4.73$ ), and the qualitative results, these features seem to have been more effective. In general, these poetic or metaphorical techniques give an impression of other aspects of perception, emotions, tastes and smells.

### 11.5.3 Discussion

By proposing three distinct categories of first-person POV simulation, we have unpacked the distinction between simulations that show what it is like to be somewhere and what it is like to be someone. Although often an ‘environmental reconstruction’ may show what it is like to be in the same location as a person, this does not always give a deeper sense of what the sensory experience of that person may be like, as this varies between individuals. As we have seen in this chapter, this becomes especially important when considering autistic people, as they have a markedly different sensory experience than people who are not autistic. We propose that simulations of ‘subjective attention’ and ‘deep subjectivity’ may be more

effective as a means to convey such subjective experiences, and tentatively suggest that these may also lead to a greater sense of empathy for the audience, since they communicate more essential aspects of conscious experience.

## 11.6 Conclusion

This work continues to develop a research trajectory explored by the ‘Affective Audio’ research team, which investigates issues of design and evaluation, of computational systems that represent subjective emotions and altered states of consciousness (see also Weinel et al. 2014a, b; Cunningham and Weinel 2016). In this context we find unique interest in the Autism TMI VR Experience, since the video utilises immersive VR technology and audio-visual design in order to communicate the subjective sensory experience and emotions of an autistic person. While the application is highly novel, it can also be seen as part of a broader emerging trend in the use of VR and 360-videos to convey subjective consciousness and provide empathic tools.

Through our ‘expert analysis’ of the Autism TMI VR Experience we sought to identify the specific design approaches that convey the subjective sensory experience of an autistic person. Following this, the results from the pilot user study began to evaluate the efficacy of these approaches. Whilst indicative, because of the sample size employed, the outcomes of the study suggest that the Autism TMI VR Experience is successful in conveying aspects of attention, vision, hearing and emotion, of an autistic person. Conversely the gustation, olfaction, tactility, proprioception and vestibular senses did not seem to be communicated effectively through the app, even though our analysis found attempts to incorporate some of these features through the design of certain elements in the video. With regards to the latter senses, it is notable that these features are less commonly explored in mass media, such as television, films, video games, radio; hence both the techniques for their representation and audience exposure to such materials may be less well developed.

Despite the difficulties in representing some types of sensory experience, the Autism TMI VR Experience nonetheless seemed to give the viewers’ an overall impression of not only being somewhere else, but also being someone else; and participants reported a sense of empathy from this experience. In this chapter we have tentatively suggested that this sense of empathy may arise due to the representation of non-aural and non-visual features of sensory perception such as attention and emotion. Developing these concepts, we have proposed three varieties of first-person POV including that of ‘deep subjectivity’. In further research it would be useful to compare all three varieties of simulation proposed using bespoke software, in order to explore the effectiveness of otherwise of each for generating empathy.

Among the outcomes of the study, it is also notable that our comparison of the VR and 360 YouTube versions of the Autism TMI VR Experience suggested no clear benefit through the use of VR-cardboard. This may be significant for the developers of such tools, as the use of VR technology may be an unnecessary additional

expense. However, we acknowledge that VR-cardboard offers inferior quality in comparison with more expensive VR hardware solutions, which may offer higher resolutions and greater levels of immersion, leading to a more distinct advantage in terms of empathy levels. A future study should compare the responses of participants using the VR-cardboard, 360 YouTube, Oculus Rift, and HTC Vive devices, to determine if the platform and imaging resolution have an impact upon participants' responses. Furthermore, there may be other benefits for using VR-cardboard that are beyond the scope of our study here, such as the novelty of the VR experience, which may attract a larger audience to the video; and the branding opportunities that the cardboard provides.

To conclude, applications such as the Autism TMI VR Experience present a variety of problems for research, some of which we have begun to explore through the case study presented here, while others have been identified which require further work. Such applications may provide powerful new communicative tools that are useful for generating awareness of autism, but equally the concept of 'deep subjectivity' that we have explored here may also be transferred to other first-person POV forms of media, such as those encountered in computer games, serious games, simulations and immersive entertainment media.

**Acknowledgements** The authors would like to acknowledge: the NAS and Happy Finish for creating the Autism TMI Virtual Reality Experience; the participants who took part in the pilot study; and the reviewers, who provided valuable insights and recommendations.

## References

- Amanatides J (1987) Realism in computer graphics: a survey. *IEEE Comput Graph Appl* 7(1):44–56. <https://doi.org/10.1109/MCG.1987.277026>
- American Psychiatric Association (2013) *Diagnostic and statistical manual of mental disorders*, 5th edn. American Psychiatric Association, Washington, DC
- Arora G, Pousman B (2015) *Clouds Over Sidra* [video]. Within: extraordinary stories in virtual reality. <http://with.in/watch/clouds-over-sidra/>. Accessed 7 Nov 2017
- Asperger H (1944) Die 'autistischen psychopathen' im kindersalter. *Archive für psychiatrie und Nervenkrankheiten* 117:76–136. Reprinted in: Frith U (ed) (1991) *Autism and Asperger Syndrome*. Cambridge University Press, Cambridge
- Atkinson RC, Shiffrin RM (1971) The control of short-term memory. *Sci Am* 225(2):82–90
- Baranek GT (2002) Efficacy of sensory and motor interventions for children with autism. *J Autism Dev Disord* 5(32):397–422. <https://doi.org/10.1023/A:1020541906063>
- BeAnotherLab (2016) The machine to be another. <http://www.themachinetobeanother.org/>. Accessed 7 Nov 2017
- Berkovits L, Eisenhower A, Blacher J (2017) Emotion regulation in young children with autism spectrum disorders. *J Autism Dev Disord* 47(1):68–79. <https://doi.org/10.1007/s10803-016-2922-2>
- Bogdashina O (2010) *Autism and the edges of the known world: sensitivities, language and constructed reality*. Jessica Kingsley Publishers, London
- Bogdashina O (2016) *Sensory perceptual issues in autism and Asperger syndrome*, 2nd edn. Jessica Kingsley Publishers, London



- Borg M, Paprocki M, Madsen C (2014) Perceptual evaluation of photo-realism in real-time 3D augmented reality. In: Proceedings of the 9th international conference on computer graphics theory and applications (GRAPP-2014), pp 377–386. doi: <https://doi.org/10.5220/0004683203770386>
- Broadbent D (1958) Perception and communication. Pergamon Press, New York
- Brugha T, Spiers N (2012) Estimating the prevalence of autism spectrum conditions in adults: extending the 2007 adult psychiatric morbidity survey. UK Data Archive, Colchester, Essex. doi: <https://doi.org/10.5255/UKDA-SN-7082-1>
- Callan DE, Jones JA, Munhall K et al (2003) Neural processes underlying perceptual enhancement by visual speech gestures. *Neuroreport* 14(17):2213–2218
- Calvert GA, Bullmore ET, Brammer MJ et al (1997) Activation of auditory cortex during silent lipreading. *Science* 276(5312):593–596
- Capellen B, Andersson AP (2016) Health improving multi-sensorial and musical environments. In: ACM proceedings of the 10th Audio Mostly: a conference on interaction with sound, Norrköping, Sweden, pp 61–68. doi: <https://doi.org/10.1145/2986416.2986427>
- Craig TKJ, Rus-Calafell M, Ward T et al (2015) The effects of an Audio Visual Assisted Therapy Aid for Refractory auditory hallucinations (AVATAR therapy): study protocol for a randomised controlled trial. *Trials* 16:349. <https://doi.org/10.1186/s13063-015-0888-6>
- Creak M (1961) Schizophrenic syndrome in childhood: Progress report of a working party. *Cereb Palsy Bull* 3:501–504. <https://doi.org/10.1111/j.1469-8749.1961.tb10406.x>
- Cunningham S, Weinel J (2016) The sound of the smell (and taste) of my shoes too: mapping the senses using emotion as a medium. In: ACM proceedings of the 10th Audio Mostly: a conference on interaction with sound, Norrköping, Sweden, pp 28–33. doi: <https://doi.org/10.1145/2986416.2986456>
- Cunningham S, Weinel J, Picking R (2016) In-game intoxication: demonstrating the evaluation of the audio experience of games with a focus on altered states of consciousness. In: Garcia-Ruiz M (ed) Games user research: a case study approach. CRC Press/Taylor & Francis, Boca Raton, pp 97–118
- Gabrielsson A, Lindstrom E (2012) The role of structure in the musical expression of emotions. In: Juslin PN, Sloboda JA (eds) Oxford handbook of music and emotion: theory, research, applications. Oxford University Press, Oxford
- Gauntlett J (2016) In my shoes project. <http://janegauntlett.com/inmyshoesproject/>. Accessed 7 Nov 2017
- Grandin T (2009) Visual abilities and sensory differences in a person with autism. *Biol Psychiatry* 65(1):15–16. <https://doi.org/10.1016/j.biopsych.2008.11.005>
- Guldborg K (2010) Educating children on the autism Spectrum: preconditions for inclusion and notions of “best autism practice” in the early years. *Br J Spec Educ* 37(4):168–174
- Hermann T, Yang J, Nagai Y (2016) EmoSonics – interactive sound interfaces for the externalization of emotions. In: ACM proceedings of the 10th audio mostly: a conference on interaction with sound, Norrköping, Sweden, pp 61–68. doi: <https://doi.org/10.1145/2986416.2986437>
- Kanner L (1943) Autistic disturbances of affective contact. *Nerv Child* 2:217–250
- Kenny L, Hattersley C, Molins B et al (2015) Which terms should be used to describe autism? Perspectives from the UK autism community. *Autism* 20(4):442–462
- Kinnealey M, Pfeiffer B, Miller J, Roan C, Shoener R, Ellner ML (2012) Effect of classroom modification on attention and engagement of students with autism or dyspraxia. *Am J Occup Ther* 66(5):511–519
- Mason LR, Cheren C, Gauntlett J et al (2016) Virtual futures salon: electronic empathy. 26th July 2016, Lights of Soho, London, UK. <http://www.virtualfutures.co.uk/event/vfsalon-electronicempathy/>. Accessed 7 Nov 2017
- Mazefsky CA, Herrington J, Siegel M et al (2013) The role of emotion regulation in autism spectrum disorder. *J Am Acad Child Adolesc Psychiatry* 52(7):679–688. <https://doi.org/10.1016/j.jaac.2013.05.006>
- McGurk H, MacDonald J (1976) Hearing lips and seeing voices. *Nature* 263:747–748

- Murray N, Lee B, Qiao Y, Muntean GM (2016) Olfaction-enhanced multimedia: a survey of application domains, displays, and research challenges. *ACM Comput Surveys (CSUR)* 48(4):56
- NAS (2015) TMI campaign report. The National Autistic Society. <http://www.autism.org.uk/get-involved/tmi/about/report.aspx>. Accessed 7 Nov 2017
- NAS (2016a) Autism TMI virtual reality experience. YouTube. [https://www.youtube.com/watch?v=DgDR\\_gYk\\_a8](https://www.youtube.com/watch?v=DgDR_gYk_a8). Accessed 7 Nov 2017
- NAS (2016b) Can you make it to the end? Behind the scenes [video]. YouTube. <https://www.youtube.com/watch?v=RnQR71AJVH4>. Accessed 7 Nov 2017
- NAS (2016c) Experience too much information in virtual reality. The National Autistic Society. <http://www.autism.org.uk/get-involved/tmi/virtual.aspx>. Accessed 7 Nov 2017
- NAS (2016d) SPELL. The National Autistic Society. <http://www.autism.org.uk/about/strategies/spell.aspx>. Accessed 7 Nov 2017
- NAS (2016e) Too much information. The National Autistic Society. <http://www.autism.org.uk/get-involved/tmi.aspx>. Accessed 7 Nov 2017
- NAS (2016f) What is autism? The National Autistic Society. <http://www.autism.org.uk/about/what-is.aspx>. Accessed 7 Nov 2017
- Noé G (2009) Enter the Void [film]. Fidélité Films (co-production), Wild Bunch (co-production), BUF (co-production)
- Pellicano E (2013) Sensory symptoms in autism: a blooming buzzing confusion? *Child Dev Perspect* 7(3):143–148. <https://doi.org/10.1111/cdep.12031>
- Pfeiffer B, Kinnealey M, Reed C, Herzberg G (2005) Sensory modulation and affective disorders in children and adolescents with Asperger's disorder. *Am J Occup Ther* 59(3):335–345
- Ranasinghe N, Nakatsu R, Hideaki N, Gopalakrishnakone P (2012) Tongue mounted Interface for digitally actuating the sense of taste. In: Proceedings of the 16th IEEE international symposium on wearable computers (ISWC), pp 80–87. doi: <https://doi.org/10.1109/ISWC.2012.16>.
- Russell J (1980) A circumplex model of affect. *J Pers Soc Psychol* 39(6):1161–1178
- Sanchez-Vives MV, Slater M (2005) From presence to consciousness through virtual reality. *Nat Rev Neurosci* 6:332–339. <https://doi.org/10.1038/nrn1651>
- Slater M (2009) Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Phil Trans R Soc B* 364:3549–3557. <https://doi.org/10.1098/rstb.2009.0138>
- Slater M, Wilbur S (1997) A Framework for Immersive Virtual Environments (FIVE): speculations on the role of presence in virtual environments. *Presence Teleop Virt* 6(6):603–616
- Sonnenschein D (2001) Sound design: the expressive power of music, voice and sound effects in cinema. Michael Wiese Productions, Studio City, pp 152–153
- Spencer BS (2006) Incorporating the sense of smell into patient and haptic surgical simulators. *IEEE Trans Inf Technol Biomed* 10(1):168–173
- Treisman AM (1960) Contextual cues in selective listening. *Q J Exp Psychol* 12:242–248. <https://doi.org/10.1080/17470216008416732>
- Weinel J, Cunningham S, Griffiths D (2014a) Sound through the rabbit hole: sound design based on reports of auditory hallucination. In: ACM proceedings of the 9th Audio Mostly: a conference on interaction with sound, Aalborg, Denmark. doi:<https://doi.org/10.1145/2636879.2636883>
- Weinel J, Cunningham S, Griffiths D, Roberts S, Picking R (2014b) Affective audio. *Leonardo Music J* 24:17–20
- Wing L (1969) The handicaps of autistic children: a comparative study. *J Child Psychol Psychiatry* 10(1):1–40
- Williams D (1994) *Somebody somewhere*. Doubleday, New York

# Chapter 12

## Sonification and HCI



Visda Goudarzi

**Abstract** This chapter investigates the second and third wave HCI design processes for an auditory display design and proposes a multifaceted framework. The purpose of using such a framework in a participatory context is to provide the possibility to create shared design knowledge in sonification and building new sonification designs on top of the prior work. From the early stages of the project involving the domain scientists in the process seemed to be an obvious choice. The process works in the sense that we gathered a diverse set of data analysis problems, solutions, and methods that work for data scientists within a sonification framework.

### 12.1 Auditory Versus Visual Representation of Data

Auditory and visual representation of data both have their benefits and drawbacks depending on the application for which they are used. Additionally, information perceived from one of these two modalities can influence the performance on perception in the other one. In order to get a better understanding of sonification, it is necessary to explain where auditory representation of data is advantageous and where visual representation works better.

The science that concerns itself with low-level auditory perception is known as psychoacoustics, and an understanding of the human auditory system is crucial in the use and optimisation of auditory displays. According to Kramer et al. (1999), there are some potential applications where auditory display is more advantageous:

- Monitoring tasks where eyes are busy and an eye free interface is useful to have; e.g. cockpit operations, network monitoring, and factory floors.
- Monitoring in high stress environments. Hennemann et al. presented that response time to an auditory signal can be shorter than a visual one (Henneman

---

V. Goudarzi (✉)

Audio Arts and Acoustics, Columbia College Chicago, Chicago, IL, USA

e-mail: [vgoudarzi@colum.edu](mailto:vgoudarzi@colum.edu)

and Long 1954) which is very useful in stressful situations where an immediate reaction is essential.

- Orienting tasks where ears tell eyes where to look (Perrott et al. 1990). This type of application is very useful when sound indicates the importance of a variable, and then the details of that variable may be delivered visually.
- Monitoring or analysing large data sets. The auditory system allows the ability of backgrounding to listen to some sounds with a low attentional priority while giving enough awareness to those with higher priority.
- Comparing multiple data sets and monitoring multiple tasks is possible because of the capability of parallel listening (Gaver et al. 1991).
- Exploring time-sequenced data with a wide range using auditory displays is possible because of the acute temporal resolution in the hearing sense, which is between milliseconds to several thousand milliseconds.
- Discovering overall trends in data is possible because of auditory gestalt formation (Bregman 1994). We may discern the sound as a whole without guiding our attention to its components. Auditory gestalt allows us to collect meaningful events in a stream of data.
- Our auditory sense is sensitive to temporal changes and this is very useful in analysis of periodic/aperiodic events and temporal processes (Welch and Warren 1980).
- Remembering highly salient sonic patterns could be helpful in pattern recognition in data (Kramer et al. 1999).

Some potentials of visual modality over auditory one are:

- Visualization is culturally more pronounced. Ability to read visualizations has become common knowledge in western cultures (Bieusheuvel 1947).
- It can be created and played back without modern technical means. E.g. it's possible to draw a chart on a piece of paper quickly and discuss it whereas sound and auditory representations require at least devices for creation and playback.
- Possibility to save a discrete state of the data e.g. through taking a screenshot of an animated graph at a specific time.
- The possibility to close the eyes or looking away gives the chance to take a break from the data representation whereas the ears are not made to be shut down at any time.

In addition to these individual characteristics of each sensory modality, how one influences the other has been studied by neuroscientists. These studies have explored how information from different sensory modalities are selected and bound together in the brain to represent objects and events at several stages of perceptual processing. Recent studies have revealed that auditory and visual modalities are closely related and mutually interplaying. E.g. in the domain of motion perception, Soto-Faraco et al. (2004) suggest that visual information influences auditory motion perception and there are common neural substrates to motion perception between the visual and auditory modalities.

Kim et al. (2012) reported the effect of auditory information on visual motion perception. They focus on spatial characteristic or motion of auditory stimuli. For example, they studied the auditory effects on visual motion perception by manipulating the temporal relationship between a transient auditory stimulus and visual event.

In contrary, auditory effects on visual motion perception were reported to be absent or of a smaller size. For examples, Welch and Warren (1980) stated that in spatial tasks where visual perception is more dominant, one will always depend on vision over audition to solve spatial problems. Thus, auditory stimuli can not at all influence one's perception of the location of a visual stimulus. However, studies have established results that contradict this hypothesis and concluding that it is the precision of different sensory inputs that determines their influence on the overall perception. Changes in the location of sound can trigger visual motion perception of a static stimulus in far peripheral vision. E.g. a blinking visual stimulus with a fixed location was perceived to be in a lateral motion when it's onset was synchronized to a sound with an alternating left-right source or when it was accompanied with a virtual stereo noise source smoothly shifting in a horizontal plane.

Alais and Burr (2004), reported that the role of auditory spatial signals in cross-modal localization depends on the spatial reliability of the visual signal. Moreover, Perrott et al. (1990) stated that location discrimination performance at angles of 20° or larger are better for the auditory modality than for the visual. Therefore, auditory spatial information can modulate visual motion perception when moving visual stimuli are presented in peripheral visual field.

The auditory and visual modalities have different ecological purposes, and respond in different ways (Lennox et al. 1999). The fundamental difference is physiological though – human eyes are designed to face forward, and although there is a broad angular range of visibility, the most sensitive part of the eye, only focuses on the central part of the visual scene (Ware 2000), while the ear is often omnidirectional and used to monitor parts of the environment that the eye is not looking at currently. Eye movements and head movements are essential to the viewing of any visual scene, and the ears often direct the eyes to the important stimulus, instead of acting as a parallel information gathering system.

## 12.2 Sonification of Data

Kramer et al. defined sonification as the use of non-speech audio to convey information or perceptualize data. or Sonification is the transformation of data into perceived relations in acoustic signals for the purposes of facilitating communication or interpretation (Kramer et al. 1999).

This definition focuses on two specific points; one is that the sound that conveys information is a non-speech acoustic signal, second is that the output is information or perceptualized data and not raw data. The term 'perceptualization' of scientific data is first used by Grinstead and Smith (1990) interchangeably with the modern

definition of ‘visualisation’, but later used by Auditory Display community free of the sensory bias for auditory and visual display of data.

Later Hermann redefined sonification in more specific terms as a system that uses data as input and generates sound signals as output with these constraints:

- The sound has to reflect objective properties or relations of the data used as input.
- The transformation has to be systematic, meaning that there has to be a precise definition of how the sound is influenced by the data.
- The sonification should be able to create sound that is always structurally identical with previous outputs, given the same data and identical iterations.
- The system has to have the possibility to be used with either the same data or with different data (Hermann 2008).

The latter definition pays special attention to the problem of reproducible and pervasive computing in sonification. Furthermore, it emphasizes on establishing standards by creating identical structures where the data to be sonified is similar. This allows a more systematic and formal comparison of sonification systems.

Sonification can be classified depending on:

- Distributing technology (public/private, interactive/non-interactive, etc.)
- Intended audience/users (data scientists, visually impaired, students, etc.)
- Data source (world wide web, sensors, EEG, etc.)
- Data type (analog, digital, spatial, temporal, etc.)

Besides these classifications, sonification can also be categorized into five techniques in terms of how sound is generated from data (Hermann 2002).

- **Audification:** is the direct conversion of data points into sound samples. In order to make the signal audible, it is usually scaled into a hearable frequency range. E.g. Dombois used audification to perceptualize planetary seismic data (Dombois 2001).
- **Earcons:** are abstract synthetic tones that can be used in structured combinations to create auditory messages (Brewster 1994).
- **Auditory Icons:** are everyday non-speech sounds that directly represent the event that is being sonified. E.g. the sound of a paper basket being emptied represents metaphorically emptying trash in operating systems. Auditory icons are not as abstract as earcons. Bill Gaver (1994) introduced them by adding sounds to visual user interfaces in 1980s.
- **Parameter mapping sonification:** is the mapping of the data values to specific attributes of sounds such as volume, pitch, panning, timbre or indirectly a combination of these attributes.
- **Model based sonification:** provides a setup of a dynamic system which is parameterized from the dataset. The model provides the dynamics that determine the elements’ behaviour in time. Furthermore, some interaction modes are specified so that the user of a sonification model is able to interact with the model. The sonification is the reaction of the data-driven model to the actions of the user (Hermann 2002).

Successful applications of sonification in exploratory data analysis must be paired with a systematic procedure of understanding the working environment in which this analysis is conducted, along with the psychoacoustic principles that affect auditory perception. I discuss some examples of such sonification systems to explore some of their characteristics. The following data sonification tools all have a GUI (Graphical User Interface) and require no programming skills by the data analyst. The data is imported over text or Excel/CSV files as database support doesn't exist in these tools.

- Sonification Sandbox is developed by Sonification Lab at Georgia Tech (Walker and Cothran 2003). It creates auditory graphs using parameter mapping sonification and MIDI output for sound generation. Sonification Sandbox is used for experimenting with various sonification techniques, data analysis, science education, auditory display for blind, and musical interpretation of data (Flowers 2005). The latest version is available for various operating systems.
- xSonify is created by NASA and focuses on sonification of space physics data such as Cassini spacecraft crossing the bow shock of Saturn, and on detecting micrometeoroids impacting Voyager 2 when traversing Saturn's rings. (These impacts were obscured in the plotted data but were clearly evident as hailstorm sounds.) The main user group for this tool are visually impaired scientists and students (Candey et al. 2006). xSonify uses the Java sound API and MIDI output.
- SonifYer: is developed by sonification research group at Berne University of the Arts (Schoon and Dombois 2009). It is mainly used for time series data such as EEG data, seismological data, and fMRI. In SonifYer audification and FM-based parameter mapping sonification is used.
- SoniPy: is based on Python programming language and hosted on sourceforge. It is designed to be a framework for data sonification using components of python for data acquisition, storage, and analysis and adding perceptual mappings and sound synthesis modules into it (Worrall et al. 2007).

In recent years with growth of world wide web and other real time applications, the need for real-time monitoring of multiple data dimensions, such as for monitoring multiple sources of data has evolved. Some examples are financial data sonification systems (Janata and Childs 2004) and (Worrall 2009), twitter data sonification (Dahl et al. 2011) and (Hermann et al. 2012), network data sonification (Worrall 2015), EEG (Hermann et al. 2006), and sonification of astrophysics data (Alexander et al. 2014), to name a few.

### 12.3 Sonification Challenges

There are numerous existing sonification tools with reference to theories of auditory perception and psychophysics but few have been adopted by a specific target user base through analysis of the environment, the nature of the data and the goals of the

application. The assumption is that the auditory display methods have been designed and developed without involving the users throughout the design process. The result of such a design process is a tool that doesn't necessarily fulfill the user's needs, could result in poorly designed displays, or remain at an experimental level. In addition to functionality and usability, pleasure is also a central goal in designing applications. Users want something more than just "usable": they want applications that offer something extra that they can relate to; products that bring not only functional benefits but also emotional ones. Designing aesthetically appealing interfaces involves understanding the users and respecting human diversity.

Some of the open questions that are going to be addressed in this chapter concerning Human-Computer-Interaction (HCI) are:

- What are the most appropriate HCI methods suitable for designing auditory interfaces?
- How is a user-centered design process adapted to the design of sonification interfaces?
- Given a sonification framework, how can a pool of sonifications be created? Where do these techniques fail and where are they superior to visualisation?
- How to develop some standard sonification techniques, which assist the data mining workflow for data scientists?

## 12.4 Evolution of HCI

The main challenge in Human Computer Interaction is that people and computers are different. We focus on human-centered view, but many designs derive from a machine-centered (computer-centered) view as their central concept because it is faster and easier for the designer/developer to create such systems but not for the people who are supposed to use the system.

Because of its interdisciplinary nature, HCI has grown into an unrestrained domain. In the first decade of its evolution, HCI started with studies in the interactions and the relationships between humans and computers, focused on human factors/engineering, and interfaces (especially on the design criteria for graphical user interfaces or GUIs) to create more usable systems. As interface problems were defined more specifically and understood by a more diverse group of users, the main HCI focus started to move beyond the interface (Fischer 1993).<sup>1</sup>

Another motivation behind evolution and change in direction of HCI is the rapid pace of technological advancements in the last decades (e.g. the internet, wireless technologies, mobile devices, wearables, pervasive technologies, tracking devices) which has led to new ways of user experiences, interactions, and communications. Therefore, design methods from different fields such as ethnography, and cognitive

---

<sup>1</sup>As stated by HCI pioneer Douglas Engelbart: If ease of use was the only valid criterion, people would stick to tricycles and never try bicycles.



science have been imported to computer science to study the interactions of humans and computers. Additionally these methods have been transformed to adapt to the new interactions between humans and machines. For example, usability is expanded from traditional goals such as efficiency to user experience goals such as aesthetically pleasing, motivating, and fun (Rogers et al. 2002).

## 12.5 Situated Perspectives: 3rd Wave HCI

Third wave HCI started in mid 2000 and expanded the reach of HCI to homes and everyday culture. In this paradigm, the use context and application types are broadened, users are called actors and participants, cultural differences are notable, and thinking out of the box is supported. Anderson (1994) studied the value of ethnography for design and he pointed out that the core of the paradox is that “to understand what any individual item or part means, you have to see it against the backdrop of the whole.” This idea ties to the observations by Harrison et al. (2007) that the artifact and its context are mutually defining within the third paradigm.

Context has been a central point in the second wave and it has been talked about in many ways but it has still been very vague. Engeström’s triangle was very often used to describe the dimensions of context. However, there seem to be another new dimension in context all the time and an attempt to make a complete analysis of context in the second wave has failed. Other approaches tried to define the notion of boundary object. The attempts have been to create objects and artefacts that are self-contained to travel across contexts of use. This way of thinking can only be applied by design of mobile technologies as the individual user has access to all her personal documents and can work independently of the place. Bødker described the main challenge that led to the third wave in involvement of the users. People need to be involved in design, not only as workers but as someone who brings the whole life experience with. Fischer explored the domain of how system design could be modifiable to enable the users design in use. His early works (Fischer 1994) focused on designing knowledge-based environments in which users can create, reflect, and shape the systems. Furthermore, he introduced a participatory design (co-design) process between environment developers and domain designers. Later on (Fischer and Giaccardi 2006), he addressed the challenges of participatory design methods and introduced the evolution of user-centered development towards meta-design. Improvements in user participation had huge influence on the creation of the third wave. The innovation is user-driven and not task oriented. The cognition is expanded towards exploratory approach where designers seek inspiration from use. To recap, designing according to the third paradigm is a situated and constructive activity of meaning making, rather than problem solving.

The third paradigm has become more apparent because of several reasons including:

- The dynamic contexts of use (the focus moved from workplace to everyday lives);
- Situatedness of interaction;
- Non-task-oriented computing (such as ambient interfaces);
- Emotional human-computer interaction (Harrison et al. 2007)

## 12.6 User Centered Design for Auditory Display

The history of sound in Interface Design is beyond the scope of this chapter but outlining some aspects of Auditory Interface Design and Development is necessary to get a deeper understanding of the core concepts.

Frauenberger describes the history of sound in technology and the use of it in the very first personal computers (Frauenberger and Stockman 2009). He argues that the gaming industry is the main incentive behind the development and improvement of sound in computers. Furthermore, he distinguishes between “Auditory Display” and “Auditory (user) Interface”. The former includes any use of auditory means to convey information, which is equivalent to the definition of sonification by Kramer as “the use of non-speech audio to convey information” (Kramer et al. 2000). “Auditory Display” covers the auditory representation of data as well as the use of sound in user interfaces. “Auditory (user) Interface” is described as a sonically analogous to graphical user interface (GUI) and is mainly common to use for speech interfaces. Frauenberger states that the term “interface” implies a bidirectional communication, while “display” focuses on the presentation and feedback of information. Thus, an auditory user interface includes both ways of interaction.

Analysis of requirements and constraints, understanding the users in the context of the system’s functionality and the tasks that they are involved with are the key constituents for a successful design process, especially in an Auditory Display. The concept of task-oriented and data-sensitive auditory information design method (TaDa!) was developed by Barrass (1996) as the first step for auditory information design. TaDa! starts from a description of a use case scenario, and an analysis of the user’s task, and the characteristics of the data. This analysis informs the specification of the information requirements of the sonification. TaDa! includes some key aspects of requirement analysis in general, and yields a structured way of finding the link between information requirements of a task and the information representation used to achieve the designer’s goals. But the resulting sonification does not support user interaction with the information. Additionally Barrass (2003) explored an extension of design by introducing design patterns into the sonification field.

Frauenberger examined the concept of design patterns for auditory display design by analyzing 23 proceedings of ICAD (International Community for Auditory Display) conferences on the basis of four themes: design process, guid-

ance, rationale, and evaluation. He describes that all papers introduce the application domain, but contextual information did not play a role in the design process. After the in-depth study on design issues, he looks at the field of design in sonification from HCI community's point of view using an online survey. The results of this research (Frauenberger and Stockman 2009) show that the design process for auditory display is mostly unstructured and it provides limited support for the reuse the design knowledge created. Another issue is that methodologies and existing guidance in audio domain are often tied to a specific context and reusing them is only possible within the restricted context (Flowers et al. 2005).

Besides design patterns, another design approach that has been explored in auditory interface design is Ecological Interface Design (EID). EID presents guidelines for the development of displays where a key component is the mapping of real world properties to the interface. Furthermore, it is a design technique that originates from cognitive work analysis (CWA). CWA is a procedure to identify requirements for the interfaces of complex real time systems. EID uses some of the phases of CWA such as work domain analysis, control task analysis, and semantic mapping:

1. Work Domain Analysis provides information about why the system exists, the flow of information through it and its functions. It helps to identify work domain characteristics and relations that are needed to be displayed in any interface. For example, physical properties of work domain may specify if edification or parameter mapping sonification is suitable. At this point the information is not sufficient for interface design.
2. Control Task Analysis provides information about what needs to be done, by whom, when, and how information about activity might be transmitted. It also gives information about temporal relations between tasks. For example, it gives information about which tasks are better suited to be displayed visually and which tasks are more appropriate for auditory display.
3. Semantic Mapping provides information about criteria for choosing interface elements so that goal related task invariants are mapped into perceptual properties of the interface. For example, it gives designers a framework to decide on dimensions of an auditory stimulus, based on knowledge of auditory perception.

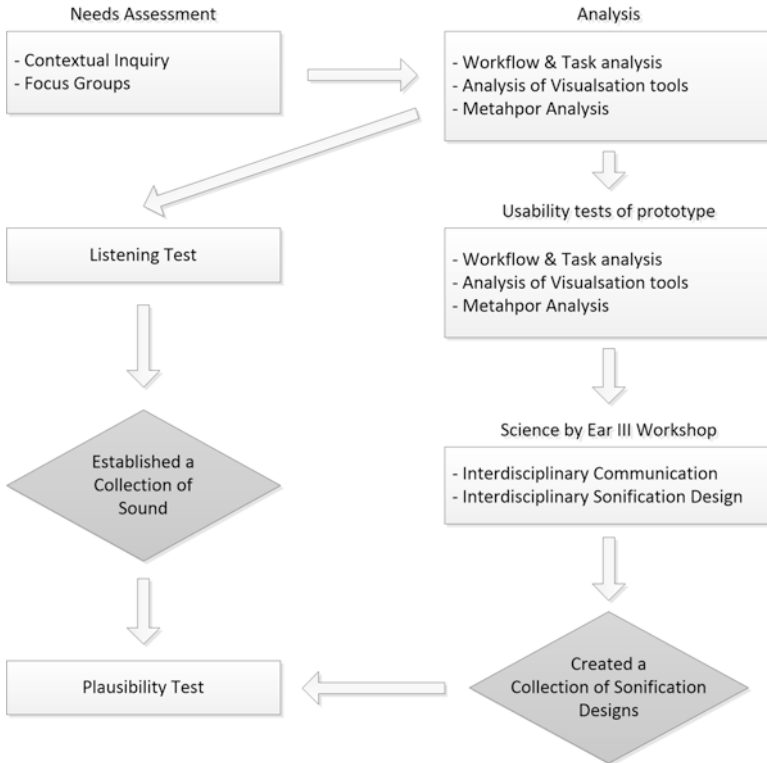
EID differs from user-centered design in that the focus of the analysis is on the work domain rather than the end user or specific task. EID fundamentals are not limited to visual displays, however they have been commonly utilized in visual display design. Gaver (1993) used ecological concepts in his work on auditory icons and earcons and his technique has also been used to the sonification of real time data. Gaver et al. (1991) used ecological approach in the Arkola simulation of a bottling plant and Mynatt (1997) used it in a marine-steam power plant but they did not use a full EID analysis. Instead they emphasized on how to represent physical functions acoustically. A full EID approach with higher order properties in auditory display design was first introduced by Sanderson et al. (2000) argued that if EID is to be used for designing auditory interfaces, in addition to the semantic mapping, an

attentional mapping phase, is needed. Based on knowledge of auditory attention, this phase provides requirements on how an auditory display should control attention alongside other interface elements.

## 12.7 Interdisciplinary Collaborative Work

In sonification of scientific data, designers know very little about the domain science and domain scientists are not familiar with the sonification methodology. The knowledge about the domain science is not given, but evolved during the problem-solving process. Some design challenges in auditory display regarding user-centered approaches are explored, and involvement of domain scientists throughout sonification designs is suggested. We (me and my research group) explored this within a workshop in which sonification experts, domain experts, and programmers worked together to better understand and solve problems collaboratively. The sonification framework that is used during the workshops is described in the previous chapter and the workshop process and how each group worked together during the workshop sessions are examined. Participants worked on pre-defined and exploratory tasks to sonify climate data. Furthermore, they grasped each other's domains; climate scientists especially became more open to use auditory display and sonification as a tool in their data analysis tasks. Resulting sonification prototypes and workshop sessions are documented on a wiki to be used by the sonification community. To get started, we used some of the sonification designs created during the workshop for an online study where participants from science, engineering, and humanities were asked questions about the data behavior by listening to sonifications of bivariate time series.

Frauenberger and Stockman showed that the design process for auditory display is mostly unstructured and it provides limited support to reuse the design knowledge created (Frauenberger and Stockman 2009). Another issue is that methodologies and existing guidance in the auditory domain are often affiliated with a specific context and reusing them is only possible within the specific context (Flowers et al. 2005). A sonification tool as a general software package to develop quick sonification designs for a wide range of scientific domains has been explored by de Campo (2007). Other tools, such as Sonification Sandbox (Walker and Cothran 2003) or SONART (Ben-Tal et al. 2002) have investigated a smaller range of applications. In our approach, we wanted to focus on a specific domain (climate science) and context (as Flowers et al. suggested) but giving a broad range of sonification design possibilities to the users and the power of designing sonifications. We tried to include the users actively in every stage of the design process. Some of their active roles were: choosing and making data sets available that are more challenging for them to analyse, early stage usability tests of the tool, suggesting the missing features, thinking and solving sonification problems with us throughout the interdisciplinary workshop (Fig. 12.1).



**Fig. 12.1** Overview of a sonification design process

Sonification of scientific data requires understanding and expertise in the domain science, sonification design, and software engineering. In order to create useful sonifications, experts design and develop sonification systems iteratively working with the domain scientists. In our approach the aim was to create an interdisciplinary sonification platform which enables climate scientists and sonification researchers to generate sonifications systematically. Climate scientists provided a huge variety of measured and simulated climate data for this goal. The starting point for our approach were previous interdisciplinary sonification workshops which had a broader user group than our project. In *Science by Ear I and II* (de Campo et al. 2006) workshops, domain scientists from different scientific domains with a variety of data (e.g. medical data, sociological data, physics data) participated. Our focal point was one specific domain with a variety and complexity of data sets and problems within this domain. Figure 12.1 provides an overview of various stages of the design process used in our interdisciplinary approach.

## 12.8 Participatory Design Approach

Creating a sonification platform to analyse scientific data that is user-friendly, efficient, and effective requires a broad knowledge of the domain science. The knowledge to understand, frame, and solve problems in the domain science is not given, but is established and advanced during the design. In such a continual process, users become co-designers not only at design time, but also throughout the whole existence of the sonification system. Rather than presenting users with closed systems or predefined sonifications, we intended an iterative system design that evolves by user's engagement to explore and design a variety of sonification possibilities for their problem domain. This allows the users to extend the system to fit to their specific tasks and needs while being assisted by sonification experts in this process. We partially used user-centred (Norman and Draper 1986) and participatory design, but we tried to extend our approach in parts to meta-design (Giaccardi and Fischer 2008) to shift some control from designers to the domain scientists (this approach didn't fully succeed due to complexity and lack of time) by empowering them to create and contribute their own objectives in the sonification design method. The participation of the users in the design decisions go beyond the processes at the design time. We particularly included participatory design (Schuler and Namioka 1993) to involve users in the co-design process with the sonification designers. Participatory Design (PD) is a process that uses early and continual inclusion of the users to produce a technology, such that they actively involve in setting design decisions and planning prototype (Carroll et al. 2000).

The basic rationale behind PD is users' participation in design by involving them in decisions that affect them and the practical issue of accessing their extensive domain knowledge. Moreover, users need to have a sense of ownership in order to make better use of the technology. This concept is first introduced by Scandinavian researchers who tried to rely on union-sponsored workshops and games involving direct interaction between designers and workers (Gregory 2003). They showed that workers who can employ certain control over their work tend to be more motivated and develop more efficient and effective work practices. Technologies are worthless on their own; it is users' expertise that makes them valuable. Some users may simply use a limited set of functionality and ignore the potential value of the technology while others can use it creatively (Yamauchi and Swanson 2010). User participation is an effective technique to empower users.

PD is still developing and hence its research design tends to be adaptable. The later work in PD has tended to add targeted interaction with less intrusive methods such as observation and artifact analysis. According to Spinuzzi (2005) three basic stages in most PD research are:

1. Initial exploration of work; designers meet the users and familiarize themselves with the ways in which the users work together. This exploration includes the technologies used, but also includes workflow and work procedures, routines, teamwork, and other aspects of the work.

2. Discovery processes; designers and users employ various techniques to understand and prioritize work organization to clarify the users' goals and values and to agree on the desired outcome of the project. This stage is often conducted on site.
3. Prototyping; designers and users iteratively shape technological artifacts to fit into the workplace envisioned in Stage 2. Prototyping can be conducted on site or in a lab; involves one or more users; and can be conducted on-the-job if the prototype is a working prototype.

Some techniques of contemporary PD include but are not limited to: video brainstorming (Mackay 2002), structured brainstorming or Bootlegging (Holmquist 2008), and the more physical technique: Bodystorming (Oulasvirta et al. 2003). Bootlegging is a structured brainstorming method that is suitable to generate several ideas within a very multidisciplinary workshop. It's not suited for our workshop because our tasks were not as random and the project was not in such an early stage to benefit from a vast number of ideas. We rather needed a more focused technique.

Participatory workshops adopt the principles of PD by shrinking the user involvement into workshop sessions. The workshops are brisk and powerful tools that help designers to find out specific activities and situations without a thorough investigation of the context. Common PD techniques such as paper prototyping and storyboarding are more suitable for GUI design rather than auditory display design. The adaptation of some of these techniques in the design of an auditory interface in the context of a multidisciplinary workshop has been previously investigated e.g. in Droumeva and Wakkary (2006), Svanaes and Seland (2004), and Taxén (2004). More specifically in the context of data sonification deCampo et al. investigated participatory workshops within the scientific context. Participatory workshop settings have been useful in utilizing participants' creativity (Numa et al. 2008) by simulating environments using role playing (Droumeva and Wakkary 2006; Svanaes and Seland 2004), or in establishing a hands on experience for the users in an iterative process (Taxén 2004). The last approach is the closest to our use of participatory workshop. Given the systemic and experiential nature of our sonification environment, we decided that a participatory workshop as an alternative to controlled experiments could be more useful in the early stages of sonification design. Although we used a sequential process throughout the project, this workshop was a standalone component of the design to involve other stakeholders and sonification experts from other institutions. The users worked with a bigger group of sonification experts than they were usually working with during the project. This led to an almost equal number of sound experts and climate scientists in the process of co-design which created a designer-designer interactive atmosphere rather than designer-user atmosphere. Additionally, we suggest a long-term inclusion of users or user representatives within the design team to be continued after the workshop. Despite the advantages of PD during the design time, sonification systems need to be evolvable to fit new needs and tasks created by users after the completion of the system. Therefore, we needed the domain scientists to be fully involved to contribute and modify the system themselves when new needs arise. Nevertheless, the

sonification design space (de Campo 2007) is huge and impossible to be explored by novice sonification designers. Thus, during the workshops we focused on specific use cases that represent a variety of domain scientists' workflows to explore the relevant design space. Our approach is an open framework for sonification researchers and climate scientists to develop a variety of sonifications but also having the option of using default mappings of climate parameters to sound parameters, suggested by experts. Both groups could work on the same tasks and context. The same script and its configurations could be edited by sonification designers through coding or by climate scientists using the graphical user interface. This form of collaboration is designed and developed in within the platform and makes it possible for both sonification designers and climate scientists to contribute. However, during the workshop the climate scientists were not equipped with their own computers and therefore did not contribute directly into the final sonification designs, but indirectly they were the decision makers regarding several aspects.

## 12.9 Discussion and Future Directions

We investigated the second and third wave HCI design processes for an auditory display and proposed a multifaceted framework. The purpose of using such a framework in a participatory context was to provide the possibility to create shared design knowledge in sonification and building new sonification designs on top of the prior work. From the early stages of the project involving the domain scientists in the process seemed to be an obvious choice. The process worked in the sense that we gathered a diverse set of data analysis problems, solutions, and methods that work for data scientists within our sonification framework.

One of the main challenges that we experienced throughout the process is the domain scientists' skepticism toward sonification and auditory display as a useful tool. The multidisciplinary workshop helped to reduce this skepticism due to the hands-on nature of the hack sessions. The participation of the domain scientists in form of brainstorming sessions was essential but not sufficient. The critical point is that the scientists would not keep on using a framework unless they get hands on experience programming with their own computer at their own workplace. Regular use of the software can improve the involvement of the domain scientists.

Furthermore, the concept of collaborative and participatory design approach fully depends on the support of the community around it. We need to integrate the domain scientists and other sonification experts to use and contribute to the open source software environment in order to transfer and exchange sonification design knowledge. The hope is that this approach and approaches similar to this enable sonification designers and domain scientists to build on each other's works and encourage them to contribute long term in creation of sonification designs efficiently.

The work presented in this chapter looked into the design process of sonifications from human-computer interaction perspective and proposed a framework of



methods to aid the transfer of design knowledge throughout the problem solving lifecycle.

We set out to investigate the evolution and current practices in HCI including the three paradigms. The concept of the third wave as a key element in this work, was reviewed in more detail. Additionally, the focus was on user centered design and user studies where requirements for a methodological design framework were derived from. These led to the development of a sonification framework. Then the design choices and components in the framework are introduced to users and user studies were conducted. The two groups participated in these user studies were domain scientists and sound experts. Finally, participatory design is introduced and an interdisciplinary workshop is explored following sonification designs to investigate the quality of such interdisciplinary work. Sonification designers and domain scientists participated in this workshop and the results provided valuable insights into various aspects of knowledge transfer and sonification design practice.

The intended readers for this chapter are the sonification and auditory display community and the broader HCI research practitioners. Climate scientists and data science community in general may also benefit from reading this work. The contributions to each of these disciplines vary. The hope is that this work will impact the use of sonification and auditory displays in data analysis tools and makes it more common in everyday technology in the long term. The contribution to the display design and HCI community originates from the insights we gained from adapting the HCI design practices and concepts into sonification field using climate science as an exemplar.

The user-centred, methodological approach to sonification and the application of interdisciplinary work with scientists is proposed, however, is not tied to the specific domain of climate and is potentially applicable to any discipline.

This chapter has provided some answers to the questions it intended to investigate, but also left others open and produced some new ones. The following is an effort to emphasize some of the main issues that arose from this work and reflect on them as well as propose future lines of research. The overall research question is:

How to design a systematic approach to develop sonification frameworks that facilitate the efficient use of sonification by domain scientists?

The key design methods suggested in this chapter are extracted from user centered design disciplines of HCI. Needs assessment methods such as contextual inquiry and focus groups helped to gather insights into the domain scientists needs. Furthermore, analytical methods such as workflow and task analysis outlined the way domain scientists work, think, and analyse data. Methods from the third paradigm of HCI were useful to tackle the users needs in their own context. PD methods such as interdisciplinary communication and workshop created a platform for an interdisciplinary knowledge creation and exchange. We definitely did not aim to create new hypothesis in the domain science, rather looked into alternative methods to explore and present complex data. Creating a general tool for development of sonification methods and sound creation was not the central goal, but also a side product of the work. In the process of this work by the author and the experts involved in related projects several sonification design examples are created. The

raw material has been reworked and attached in the projects' wiki. Creating reproducible sonification methods for data mining tasks and comparative research methods in sonification were also by-products that could be shared with Auditory Display communities in specific and HCI communities in general.

To carry on this research field, I suggest to continue the collaboration with domain scientists through training and workshops to create a larger pool of available sonification designs and prototypes. Then, comparing and using these prototypes in combination of visualization techniques to see/hear where they work together or where they are more useful. Additionally, integration of sonification into the scientists' workflow as a regular hack session every few months would be useful to familiarize them and to make it part of their data analysis routine. The author hopes that systematic sonification inspires researchers to develop new methods of sonification and the application of a variety of HCI methods in sonification lead to some standard which assist the data exploration workflow.

## References

- Alais D, Burr D (2004) The ventriloquist effect results from near optimal bimodal integration. *Curr Biol* 14(3):257–262
- Alexander RL, O'Modhrain S, Roberts DA et al (2014) The bird's ear view of space physics: audification as a tool for the spectral analysis of time series data. *J Geophys Res Space Phys* 119(7):5259–5271
- Anderson RJ (1994) Representations and requirements: the value of ethnography in system design. *Human Comput Interact* 9(3):151–182
- Barrass S (1996) Tada! Demonstrations of auditory information design. Unpublished PhD thesis
- Barrass S (2003) Sonification design patterns. In: *Proceedings of International conference on auditory display*, Boston
- Ben-Tal O, Berger J, Cook B et al (2002) The sonification application research toolbox. In: *Proceedings of International conference on auditory display*, Kyoto
- Bieusheuvel S (1947) Psychological tests and their application to non-European peoples. In: *Yearbook of education*. Evans, London, pp 185–207
- Bregman AS (1994) *Auditory scene analysis: the perceptual organization of sound*. MIT Press, Cambridge
- Brewster S (1994) Providing a structured method for integrating non-speech audio into human-computer interfaces. Unpublished PhD thesis
- Candey RM, Schertenleib AM, Merced WL (2006) Xsonify sonification tool for space physics. In: *Proceedings of International conference on auditory display*, London
- Carroll JM, Chin G, Rosson MB, et al (2000) The development of cooperation: five years of participatory design in the virtual school. In: *Proceedings on designing interactive systems: processes, practices, methods, and techniques*. Association for Computing Machinery, New York, pp 239–251
- Dahl L, Herrera J, Wilkerson C (2011) Tweetdreams: making music with the audience and the world using real-time twitter data. In: Jensenius AR, Tveit A, Godøy RI et al (eds) *Proceedings of the international conference on new interfaces for musical expression (NIME)* Oslo, pp 272–275
- de Campo A (2007). A data sonification design space map. In: *Proceedings of the 2nd International workshop on interactive sonification*, York
- de Campo A, Daye C, Frauenberger C et al (2006) Sonification as an interdisciplinary working process. In: *Proceedings of International conference on auditory display*, London

- Dombois F (2001) Using audification in planetary seismology. In: Hiipakka J, Zacharov N, Takala T (eds) Proceedings of the 7th International conference on auditory display, Espoo, pp 227–230
- Droumeva M, Wakkary R (2006). The role of participatory workshops in investigating narrative and sound ecologies in the design of an ambient intelligence audio display. In: Stockman T, Nickerson LV, Frauenberger C et al (eds) Proceedings of the international conference on auditory display, London, pp 36–43
- Fischer G (1993) Beyond human computer interaction: designing useful and usable computational environments. In: People and computers VIII: proceedings of the HCI'93 conference (Loughborough, England), Cambridge University Press, Cambridge, UK, pp 17–31
- Fischer G (1994) Putting the owners of problems in charge with domain-oriented design environments. In: Gilmore D, Winder R, Detienne F (eds) User-centered requirements for software engineering environments. Springer, Heidelberg, pp 297–306
- Fischer G, Giaccardi E (2006) Meta-design: a framework for the future of end-user development. In: Lieberman H, Paternò F, Wulf V (eds) End user development. Springer, Dordrecht, pp 427–457
- Flowers JH (2005) Thirteen years of reflection on auditory graphing: promises, pitfalls, and potential new directions. In: Brazil E (ed) Proceedings of International conference on auditory display, Limerick
- Flowers JH, Turnage KD, Buhman DC (2005) Desktop data sonification: comments on flowers et al ICAD 1996. *ACM Trans Appl Percept* 2(4):473–476
- Frauenberger C, Stockman T (2009) Auditory display design—an investigation of a design pattern approach. *Int J Hum Comput Stud* 67(11):907–922
- Gaver W (1993) How in the world do we hear? Explorations in ecological acoustics. *Ecol Psychol* 5:283–313
- Gaver W (1994) Using and creating auditory icons. In: Kramer G (ed) Auditory display: sonification, audification and auditory interfaces. Addison-Wesley Publishing, Reading, pp 417–446
- Gaver W, Smith RB, O'Shea T (1991) Effective sounds in complex systems: the ARKola simulation. In: Robertson S, Olson G, Olson JS (eds) Proceedings of the SIGCHI conference on human factors in computing systems, New Orleans, pp 85–90
- Giaccardi E, Fischer G (2008) Creativity and evolution: a metadesign perspective. *Digital Creativity* 19(1):19–32
- Gregory J (2003) Scandinavian approaches to participatory design. *Int J Eng Educ* 19(1):62–74
- Grinstein G, Smith S (1990) The perceptualization of scientific data. In: Proceedings of the SPIE/SPSE conference on electronic imaging, pp 190–199
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: Alt. Chi. Session at the SIGCHI Conference on human factors in computing systems, San Jose, California, USA, pp 1–18
- Henneman RH, Long ER (1954) A comparison of the visual and auditory senses as channels for data presentation. In: Technical report 54(363). Wright Air Development Center, Dayton
- Hermann T (2002) Sonification for exploratory data analysis. PhD thesis, Universität Bielefeld
- Hermann T (2008) Taxonomy and definitions for sonification and auditory display. In: Susini OW (ed) Proceedings of the 14th International conference on auditory display, Paris
- Hermann T, Baier G, Stephani U, et al (2006) Vocal Sonification of pathologic EEG features. In: Stockman T, Nickerson LV, Frauenberger C et al (eds) Proceedings of the 12th International conference on auditory display, London
- Hermann T, Nehls AV, Eitel F et al (2012) Tweetscaples: Real-time sonification of twitter data streams for radio broadcasting. In: Nees MA, Walker BN, Freeman J (eds) Proceedings of ICAD, Atlanta
- Holmquist LE (2008) Bootlegging: multidisciplinary brainstorming with cut-ups. In: Simonsen J, Roberston T (eds) Proceedings of the tenth anniversary conference on participatory design. Indiana University, pp 158–161
- Janata P, Childs E (2004) MarketBuzz: a Sonification of real-time financial data. In: Barrass S, Vickers P (eds) Proceedings of the tenth international conference on auditory display, Sydney

- Kim R, Peters MA, Shams L (2012) 0+ 1> 1 how adding non-informative sound improves performance on a visual task. *Psychol Sci* 23(1):6–12
- Kramer G, Walker B, Bonebright T et al (1999) Sonification report: status of the field and research agenda, report prepared for the National Science Foundation by members of the International Community for Auditory Display, available via <http://www.icad.org/websiteV2.0/References/nsf.html>. Accessed 1 Jan 2018
- Kramer G, Walker B, Bonebright T et al (2000) Sonification report: Status of the field and research agenda, available via <http://digitalcommons.unl.edu/psychfacpub/444/>. Accessed 01 Jan 2018
- Lennox PP, Myatt T, Vaughan JM (1999) From surround to true 3-d. In: 16th International Audio Engineering Society conference on spatial sound reproduction, Rovaniemi, Finland
- Mackay WE (2002) Video to support interaction design, DVD. ACM, New York. ISBN: 1-58113-516-5
- Mynatt ED (1997) Auditory icons. In: Hearst MA (ed) *Dissonance on audio interfaces*. IEEE Intelligent Systems, vol 12, pp 10–16
- Norman DA, Draper SW (1986) *User centered system design*. L. Erlbaum Associates, Hillsdale
- Numa K, Toriumi K, Tanaka K et al (2008) Knowledge-based intelligent information and engineering systems. Springer, Berlin, pp 823–830
- Oulasvirta A, Kurvinen E, Kankainen T (2003) Understanding contexts by being there: case studies in bodystorming. *Pers Ubiquit Comput* 7(2):125–134
- Perrott DR, Saberi K, Brown K et al (1990) Auditory psychomotor coordination and visual search performance. *Percept Psychophys* 48(3):214–226
- Rogers Y, Preece J, Sharp H (2002) *Interaction design: beyond human-computer interaction*. Wiley, New York
- Sanderson P, Anderson J, Watson M (2000) Extending ecological interface design to auditory displays. In: *Proceedings of the 10th Australasian conference on computer-human interaction*, Sydney
- Schoon A, Dombois F (2009) Sonification in music. In: *International conference on auditory display*, Copenhagen
- Schuler D, Namioka A (1993) *Participatory design: principles and practices*. L. Erlbaum Associates, Hillsdale
- Soto-Faraco S, Spence C, Kingstone A (2004) Cross-modal dynamic capture: congruency effects in the perception of motion across sensory modalities. *J Exp Psychol Hum Percept Perform* 30:330–345
- Spinuzzi C (2005) The methodology of participatory design. *Tech Commun* 52(2):163–174
- Svanaes, D, Seland G (2004) Putting the users center stage: role playing and low-fi prototyping enable end users to design mobile systems. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, Vienna
- Taxén G (2004) Introducing participatory design in museums. In: *Proceedings of the eighth conference on participatory design: artful integration: interweaving media, materials and practices*, Toronto
- Walker BN, Cothran JT (2003) Sonification sandbox: a graphical toolkit for auditory graphs. In: *Proceedings of International conference on auditory display*, Boston
- Ware C (2000) *Information visualization: perception for design*. Morgan Kaufman, San Francisco
- Welch RB, Warren DH (1980) Immediate perceptual response to intersensory discrepancy. *Psychol Bull* 88(3):638
- Worrall D (2009) The use of sonic articulation in identifying correlation in capital market trading data. In: *Proceedings of the 15th International conference on auditory display*, Copenhagen
- Worrall D (2015) Realtime sonification and visualisation of network metadata. In *Proceedings of the 21st International conference on auditory display*, Graz
- Worrall D, Bylstra M, Barrass S et al (2007) The design of an extendable software framework for sonification research and auditory display. In: *Proceedings of International conference on auditory display*, Montreal
- Yamauchi Y, Swanson EB (2010) Local assimilation of an enterprise system: situated learning by means of familiarity pockets. *Inf Organ* 20(3–4):187–206

# Chapter 13

## Media Poetics and Cognition in Colocative Audiovisual Displays



Michael Filimowicz

**Abstract** Colocative display is the technique of sonically articulating the screen area of audiovisual media by dynamically placing and animating the associated sounds in spatial localization to their visual cues. Two such systems have been designed to date, the author’s Pixelphonics system and the Allosphere facility at University of California at Santa Barbara. While both technologies are prototypes, and thus lacking in a rich historical tradition that might inform what film theorist David Bordwell has called historical poetics, Bordwell’s concept of analytical and theoretical poetics can be fruitfully brought to bear to elucidate the general principles for making colocative audiovisual media and applications. The poetics, or ‘principles of making’ colocative media are situated within a discussion of the empirical dimensions of auditory localization, cognition and attentional resources, general audiovisual practices, acoustics and phenomenology. A new concept, that of the soundscape, is introduced to hone in on the particular design affordances of colocative displays. This inquiry blends second and third wave HCI approaches in its hybridization of humanist media theories with cognitivist-attentional usability perspectives.

### 13.1 Audiovisual Colocation

*[W]hat cinema most wants is to come to life.*

*Sound is often the warrant and enactment of this coming to life (Connor 2013: loc 2529 & 2538)*

Colocative sound design refers to the technique of sonically articulating the screen area of audiovisual media by dynamically placing and animating the associated sounds in spatial localization to their visual cues. To date there are two such

---

M. Filimowicz (✉)

School of Interactive Arts & Technology, Simon Fraser University, Surrey, BC, Canada

e-mail: [michael@filimowi.cz](mailto:michael@filimowi.cz)

© Springer International Publishing AG, part of Springer Nature 2018

M. Filimowicz, V. Tzankova (eds.), *New Directions in Third Wave Human-Computer*

*Interaction: Volume 1 - Technologies*, Human-Computer Interaction Series,

[https://doi.org/10.1007/978-3-319-73356-2\\_13](https://doi.org/10.1007/978-3-319-73356-2_13)



**Fig. 13.1** The Pixelphonic prototype. 8' × 4' 3 mm Alupanel screen for projection-based display, utilizing an 8 × 4 2D array of audio exciters to vibrate the screen from its backside (author's image)

systems which match this definition, the Pixelphonics<sup>1</sup> technology designed by the author of the present study, and the Allosphere facility at the University of California at Santa Barbara. Pixelphonics operates by utilizing a 2D array of audio exciters attached to the backside of the screen surface, which can be either a projection-based or monitor-based display (Fig. 13.1). The array of exciters is connected to signal distribution and software, as well as post-production processes and workflows, which applies channel or object-based audio approaches to the visuals in time-based media. The sounds associated with the moving images are mapped to screen areas according to a logic of XY axes superimposed onto the screen. The Allosphere, by contrast, occupies an entire building at the UCSB campus (Allosphere Research Facility). The screen is spherical in its shape, enclosed within a 3-story anechoic chamber, and the audience may consist of up to ~30 individuals situated on a walkway which runs through the center of the sphere. The screen is made of perforated aluminum that is transparent to sound and reflective to light. Twenty-six high definition projectors and fifty-four speakers (Allosphere Research Facility) plus subwoofers are arranged to map audiovisual media in an immersive environment that is enveloping in all directions to support “ensemble-style interactions” (Kuchera-Morin et al. 2014: 17).

<sup>1</sup>Formally described in *Apparatus, Method and System for Co-locating Visual Images and Associated Sound* (U.S. Provisional Patent No. 62/482725, 2017).

Since both systems are essentially new representational devices, offering novel platforms for the display of interactive audiovisual media, the discussion which follows will be largely, though not entirely, content neutral. For Pixelphonics, 14 application areas have been identified, including, for example: games, video, animation, simulation-based training, large scale immersive exhibition displays, process control, command and control, and video conferencing. Pixelphonics is explicitly designed to be a general purpose display technology, with a flexible and scalable architecture to potentially accommodate any home, workplace and public setting. The authors of the Allosphere literature have focused primarily on exploratory and immersive interactive scientific visualization (the system has received funding from the National Science Foundation, for instance) and the use by artists for experimental creative works which may not necessarily rely on data-driven source material. Allosphere researchers, in their publications, have imagined the possibility of exporting the technology in the form of specialized installations at contexts such as science museums and planetariums (Höllner et al. 2007). Both technologies build on “long histories of decisions that have created recognizably distinct configurations of audiovisual space” (Stillwell 2013: loc 2595).

This essay describes a poetics for colocative sound design that is intended to be applicable to both kinds of systems, whether relating to a 2D screen or immersive 3D sphere. There is, however, some possibility for design hybrids of the two systems. Pixelphonic systems need not be flat but can be built out of curvilinear screen surfaces and walls at scales approaching the Allosphere, so even this difference of a flat 2D screen and 3D sphere is not a stark one, and there is plenty of room for a spatial continuum between them.

## 13.2 Poetics

This chapter will build on and extend Bordwell’s (2007) conception of cinema poetics. Bordwell distinguishes what he calls poetics from two other dominant approaches to cinema studies, interpretivist and reflectionist. The interpretivist he also refers to as dogma-driven, which describes film scholars who see their task as the interpretation of film works through referring them to other major discourses, such as phenomenology, Marxism, psychoanalysis and feminism. The interpretivist masters a conceptual framework, then seeks parallels in film works for explicating their meaning. The reflectionist, on the other hand, insists that media always be referred back to culture, and argues that media is a reflection of society and, through a process of circular reasoning, claims in parallel that society reflects itself in its media.

Poetics is distinct from both approaches, and “has no privileged semantic field, no core of procedures for interpreting textual features, and no unique rhetorical tactics” (12) and unlike interpretivist and reflectionist methods, is not primarily concerned with producing a hermeneutics.

*Poetics* derives from the Greek word *poiesis*, or *active making*. The poetics of any artistic medium studies the finished work as the result of a process of construction – a process that includes a craft component (such as rules of thumb), the more general principles according to which the work is composed, and its functions, effects and uses. Any inquiry into fundamental principles by which artifacts in any representational medium are constructed, and the effects that flow from those principles, can fall within the domain of poetics [emphasis in original].

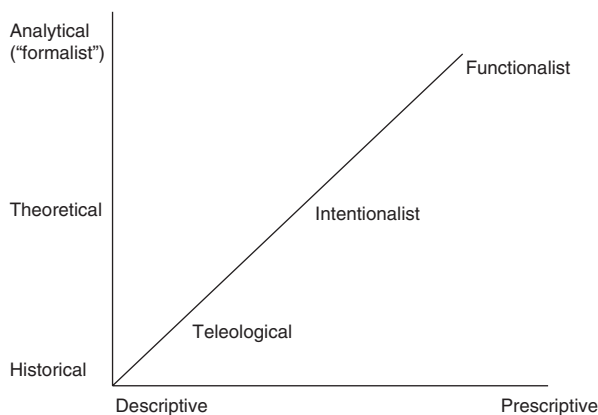
It is worth noting that in this passage, while Bordwell places poetics in relation to specifically artworks, in the reference to “any artistic medium,” he also suggests the possibility of non-artistic works, by using the broader term “artifacts” which, when combined with the semantic openness of the words “functions, effects and uses” could in principle expand poetics to the use of representational technologies beyond artworks. Poetics is inquiry into making representations. Colocative audio-visual systems like the Allosphere and Pixelphonics can easily include the display of artistic media but are not limited to it.

Bordwell proceeds to distinguish several forms of poetics. Analytical poetics, which is sometimes characterized as formalist, seeks out similar and particular devices within an artwork or a selection of works. Theoretical poetics has the aim of defining conditions or classification schemes, as exemplified by Aristotle’s *Poetics*. Historical poetics tries to understand the forms of artworks across time or within defined periods. These three varieties of poetics will often be woven together in any particular inquiry into poetics, though one or another may predominate. Kinds of poetics may also be defined in another manner as being either descriptive or prescriptive, which defines principles of making in terms that are either neutral as to their relative importance and merit, or biased toward emphasizing a particular poetic sensibility. In yet another vector explicating varieties of poetics, Bordwell identifies teleological (e.g. the evolution of cinema), intentionalist (the motivations of filmmakers) and functionalist (instituted dynamics that fulfill systemic norms) modes of poetic inquiry. Bordwell implicitly proposes what one might describe as a ‘three dimensional’ model of poetics, defined along three vectors of distinctions which we can model as follows (Fig. 13.2):

Thus, while Bordwell rejects the idea of a “privileged semantic field,” a “core of procedures” and “unique rhetorical tactics,” he replaces these with discursive modes articulated as a conceptual possibility space for a certain kind of inquiry to unfold within. This discursive possibility space articulates and supports Bordwell’s conception of poetics as a form of “rational and empirical inquiry” (15) which “appeals to intersubjectively available data that are in principle amenable to alternative explanation.” This empirical character does not make poetics scientific or ‘fact-based’ but describes how principles of making are to be clearly grounded in the making of the artifact itself. However, as depicted above (Fig. 13.2), poetics complements the artifact with a very well structured set of writing styles which can concatenate a complex diversity of approaches toward explicating its principles of making.

What...is the status of the “principles” studied by poetics? I’d argue that the principles should be conceived as underlying concepts, constitutive or regulative, governing the sorts of material that can be used in a film and the possible ways in which it can be formed.





**Fig. 13.2** Bordwell's 'three dimensional' model of cinema poetics, showing the three vectors of distinctions posed as possible approaches (author's image)

[T]he poetics I propose looks at artistic form as an organizing principle that works not on "content" but rather on *materials* (23) [emphasis in original]

While this three dimensional model of poetics' discursive tendencies is already rich with intersecting possibilities, Bordwell elucidates what he calls "a research framework" (24) that proposes six varieties of questions to be pursued in any poetics inquiry. These research/question trajectories are defined as a set of 'six p's': particulars, patterns, purposes, principles, practices and processing. The six research foci are presented as a linear subsumption hierarchy, with the six p's nested into each other as they become articulated at higher orders of organization. Particulars refer to details, such as "a line of dialogue, or a certain cut, or a moment in a performance, or an unusual sound" (24). Details are isolates and at most can be organized according to a list, while poetics proper begins when particulars come to be organized as patterns. Patterns are recurring items of interest for inquiry. Bordwell's examples are a "hero's single wisecrack" and the associations of low-key lighting characters and settings. The patterns that are most salient for poetics are the purposes that can be assigned to it. The purpose of the pattern is functional, fulfilling some need of the story, or solving a problem to produce a desired effect. Principles emerge when particulars, patterns and purposes are identified across multiple films and begin to function as norms or conventions. These norms can be relatively local, as with the placement of narrative climactic scenes, or global, as in constructing scenes out of shots of individual characters. Practices situate the previous categories as being rooted in activities, materials and tools, subject to the means-end reasoning of filmmakers and to the socio-economic constraints and opportunities of cinema's institutions (28).

The final 'p,' processing, receives the most extended treatment, and refers to the cognitive science model of top-down (concept and goal driven) and bottom-up

(perceptually organized) psychological processes. “On the whole, bottom-up processes are fast, involuntary, cheap in cognitive resources, and fairly consistent across observers....Top-down processes are slower, more voluntary, more expensive in cognitive resources, and more variable across observers” (45). These two vectors of psychological processes interact through recursive loops or “complicated feedback and input-output among many mental systems.” Bordwell articulates the layers between bottom-up “data-driven” and top-down “concept-driven” processes according to a gradient, moving from bottom to top, of Perception, Comprehension and Appropriation. Perception refers to the organization of sensory information presented by the medium. The organization of perception, e.g. through the use of style and film craft, develops comprehension, or “organizing the stimulus for uptake” (50). Appropriation works at a higher level than comprehension, and is not just another cognitive level but also indicates a shift toward audience control, or filmmakers’ lack of control. Appropriation, while describing the cognitive mode of audience-produced mashups for instance, also applies to filmmakers, as it is perfectly possible to make a film, using professional craft and careful techniques, while also being personally dismissive and disliking of the work being done.

The ability of the cognitive layers to influence each other is limited and diminishes the further along the continuum they are from each other. This is because the “feedback systems can’t go all the way down or all the way up” (46). For instance, no amount of mental willpower will allow one to perceive what is literally not presented for perception (e.g. one cannot will a red car shown in a film to be perceived as a blue car). Likewise, no amount of directorial acumen or expensive production values, working at the levels of perception and comprehension, can force audiences to react to films in ways desired by filmmakers at the appropriation level. A final note on processing is given to aspects of emotion, which Bordwell finds operating at all three cognitive levels. He identifies emotion as not being part of cognition *per se* (as much of the cognitive science literature also proposes) but rather sees emotion as working in tandem with cognitive capacities.

While Bordwell’s poetics are applied to works of film, the present discussion is focused on a new perceptual experience and affordance that has emerged in two prototype display technologies, and the way these shape the form of audiovisual media, specifically through collocating audio with the associated visuals in the screen area. Of the two technologies used for collocative sound design, the Allosphere has been used as the medium for a number of ‘finished’ works that have been designed for it and published as such in the research-creation literature. The Pixelphonic system instead has a suite of ‘content scenarios’ in which rapid prototyping techniques are brought to bear on short demonstrations of possible use cases. Both systems are forms of audiovisual representation, and like film qualify as general media with specific effects on content, form and experience which come about due to the associated practices of making.


The poetics of collocative sound design presented here, using Bordwell’s schema, is articulated according to Analytical/Theoretic-Prescriptive-Intentionalist dimensions, applied to questions most pertinent to particulars, practices, principles and processing, and traces the recursive relations through perceptual and comprehensive

cognitive levels. As will be discussed in the next section, the empirical sources for grounding a poetics for colocative audiovisual media can be found in psychophysics (localization, capture, multimodal fusion), cognition (attentional focal/ambient structure), the technology itself (its modifications and design), the application context (content types and use scenarios), phenomenological investigation (experiential gestalts), acoustic considerations (planar dispersion and point source effects), and audiovisual practices and forms on which colocative sound design can be brought to bear (e.g. games, animation, film, simulation, scientific visualization, video conference, situational awareness in media rich work environments, etc.).

### 13.3 The Soundscape

To hone in on the specific poetics of colocative sound design, I will introduce the term *soundscape* as a new parameter (a seventh ‘p’ perhaps!) in the design of audiovisual works in which sounds are dynamically placed within the specific screen areas defined by their associated visual cues. The rationale for introducing this term is grounded in several connected motivations, taking for its main background the established notion of the ‘soundtrack’ (which is also written as a single word) and the empirical-theoretical background of auditory scene analysis. The soundscape differs from an auditory scene in that the latter term designates any environment in which auditory events occur, while the former term restricts itself to screen-based media. The ‘scene’ of the soundscape coincides with the spatially framed presentation of audiovisual media, which has as its background theatrical, operatic and other stage-based live performance practices. When transposed into filmic language, the term indicates the hierarchical relationship in which the fundamental ‘building block’ of edited visual languages is taken to be the shot, which is cut together with other shots to compose a scene. “There is near-universal agreement about the minimal unit of the visual language of film: it is the shot, that which exists between one cut and another. Shots have visible edges” (Connor 2013: loc 2442).

The shot/scene distinction is particularly relevant for the kind of bottom-up and top-down (perception/comprehension) distinctions drawn by Bordwell. Shots are the most direct causal connection to the actual proto-filmic event – the reality rendered by the camera apparatus through indexical processes, whether analog or digital, which in semiotic terms yield an iconic imagery of resemblances. The scene, in contrast, denotes a logic of construction. The shot/scene distinction is strongly recursive if construed through cognitive feedback processes, since shots do not in any way naturally build themselves up into various stylistic constructions, such as shot/reverse shot, close-up/medium shot, A-roll/B-roll, and so on. Similarly, the conception of a scene in film requires sequencing a vocabulary of shot types and camera movements, often pre-visualized in other media like storyboards or animatics.



**Fig. 13.3** The badger-protagonist falls from his underground burrow into the missile control room installed beneath it, triggering an alarm which flashes in the upper left screen area




**Fig. 13.4** A change of perspective implies that now we are above rather than behind the badger, which if realistically rendered would put the sound of the alarm in the lower right quadrant of the image

The necessity for a concept like the soundscape becomes apparent when media produced for non-colocated audiovisual presentation is run through a colocative system, which will tend to reveal new problematics that were not present in the original production. The frames below, from a scene in the animated film *Badgered* (Sharon Colman 2005), illustrate how spatial-visual information presents new problems when presented in a colocative system (Figs. 13.3, 13.4, and 13.5).

In a mono or even stereo presentation of this scene, these three different spatial locations of the sounding alarm does not produce any perceptual problematic – whether our perspective is behind or above the badger, the alarm can continue to sound without producing effects of “spatial disorientation” (Donnelly 2013: loc 6998) and discontinuity. What the soundscape describes is the new form of spatial continuity produced by colocative audiovisuals. Outside of montage and experimental approaches, producers of audiovisual media typically seek to avoid overtly disruptive effects that diminish immersion or involvement with the mediated experience. Examples of unwanted disruptions abound in the history of cinema sound development. To name but a few well-known examples: audience confusion in the

**Fig. 13.5** A wider shot of the same perspective now moves the spatial location of the sounding alarm to the upper right area of the screen



earliest implementations of stereo imaging with widescreen film formats, in which actors' voices were decentered from the screen center and were felt to travel in odd ways about the theater space; the 'exit door' effect whereby sounds placed in the surround array direct the audience's attention to the walls of the theater, where the sounds seem to be emanating from; the discomfort with extremely loud soundtracks afforded by digital soundtracks with greater dynamic range, leading theater owners to lower the volume of the screening space below equipment manufacturers' recommendations (Kerins 2010: loc 675); sound that was too quiet and out of sync in the early days of cinema.

Even in avant-garde approaches, the kind of spatial disruption that can be produced by a literal remapping of sounds in a colocative remix of extant films would likely be unwelcome by an audience, as there is no general aesthetic approach of mismatching sounds to their 2D spatial placement for the sake of annoying audiences. Traditionally the logic of 'asynchronist' approaches references a logic of counterpoint whereby sounds and images don't always repeat each other but instead offer differing or complementary kinds of information, as for example expressed in Bresson's 'Notes for Sound' (Bresson 1985: 149):

To know what business that sound (or that image) has there.

What is for the eye must not duplicate what is for the ear.

If the eye is entirely won, give nothing or almost nothing to the ear\*. One can not be at the same time all eye and all ear.

When a sound can replace an image, cut the image or neutralize it. The ear goes more toward the within, the eye toward the outer.

A sound must never come to the help of an image, nor an image to help the of sound...

Image and sound must not support each other, but must work each in turn through a sort of relay.

The eye solicited alone makes the ear impatient, the ear solicited alone ,makes the eye impatient. Use these impatiences....

\* And vice versa, if the ear is entirely won, give nothing to the eye.

Spatial disruption and disorientation are not a form of counterpoint, but rather read as mistakes and poor cinematic craft. The soundscape concept thus has a prescriptive dimension to it, calling for a new modelling of the audiovisual space so that such spatial disruptions do not occur. This adds a new conceptual and aesthetic layer to previsualization and classic techniques of spatial continuity, such as the practice of establishing a 180° arc for camera placement around a 'line of action' in

a scene so that characters always remain in a consistent relative position to each other within the frame.

Potentially, colocative sound design might work as a kind of ‘corrective’ or counter-impetus to what Bordwell elsewhere has called “intensified continuity” (2002) which describes the newer style of ‘amped up’ continuity used in contemporary, and particularly, Hollywood-produced popular films.

[T]here have been some significant stylistic changes over the last 40 years. The crucial technical devices aren’t brand new – many go back to the silent cinema – but recently they’ve become very salient, and they’ve been blended into a fairly distinct style. Far from rejecting traditional continuity in the name of fragmentation and incoherence, the new style amounts to an intensification of established techniques. Intensified continuity is traditional continuity amped up, raised to a higher pitch of emphasis. It is the dominant style of American mass-audience films today. (16)

Bordwell identifies four specific techniques which have led to this intensification of continuity: faster editing, bipolar extremes of lens lengths, more close framings in dialogue scenes, and a free-ranging camera. All of these forms of intensifying continuity have spatial implications for the soundscape, and given the strong possibility of easily causing spatial-auditory disorientation by constantly shifting visuals through editing, camera movements and lens lengths, one can conjecture that colocation would foster a counter-tendency towards less intensified continuity, whereby the spatial integrity of audiovisual scenes takes on a new importance in structuring the presentation of visual information.

The soundscape does not replace the concept of the soundtrack. Rather, a soundscape can be regarded as either a new kind of soundtrack, or a new requirement or parameter for synchronized audio information in screen-based media. The term ‘soundtrack’ is often disparaged in filmsound criticism, as it is sometimes regarded as establishing a “subordination of sound to sight” (Connor 2013: loc 2360). My use of the term ‘soundtrack’ dispenses with this hypersensitivity towards sound’s status, and implies no such hierarchical privileging between image and sound. Rather, it simply names a foundational material condition, namely that audiovisual media clearly is composed of image tracks and audio tracks, and typically multiple tracks of each, which when finalized in post-production results in a locked cut for both sound and visuals. Audiovisual media also includes many other tracks which we might call ‘informatic’ tracks: timecode, sprocket holes, or control tracks. The ways that the spatiotemporal structure of an equal interval is embodied in media (e.g. through repeating spatial frames and synchronization mechanisms) could be referred to as yet additional ‘interval tracks’ which enforces the spatial and temporal regularities on which all the other tracks depend on for smooth playback, general use and manipulation.

Another critical complaint that has been voiced regarding the term ‘soundtrack’ is that it suggests to some critics that it somehow denies the reality of a multisensory experience, conveying a notion that sound in cinema can be regarded as an “acoustic isolate” (Chion 2013: loc 6388), following scientific traditions which have erroneously posited a separation of the senses for purposes of experimentation that does not in fact match our embodied and situated experiences with audiovisual media.

Chion, for instance, has claimed that there is no soundtrack, and has coined alternate terms like the “audio-visiogenic” (loc 6713) to describe a multi-sensorial cinema. Such alternative terms are deemed needed since the soundtrack is supposed to indicate a mentality of simply ‘summing’ sound and image, whereas Chion’s neologism expresses a holistic multimodal experience that is “greater than the sum of the parts” (loc 6318). While one can of course purchase “the soundtrack” to a film and listen to it as what Chion also calls a “sensory isolate”(loc 6378), this stand-alone soundtrack is not the whole soundtrack of a film, which would contain all of the dialogue and effects sounds, but just the score or licensed musical material.

It seems to be a critical overstatement to claim that the existence or use of the word “soundtrack” somehow fissions off one’s senses into separate realms. Again, the soundtrack has an empirical basis and reference in actual materials and practices. For instance, the final visual cut is typically developed by one group of professionals in one set of networked editing facilities, while the final audio mix is completed by another group of experts with their specialized technologies in mixing suites. Sound and image are sometimes assembled together – as when an editor makes a cut that includes some temp tracks and production sound – and separately – as when a field recordist wanders around a locale carrying only sound recording equipment. The multisensory features of audiovisual media can be functionally separated during production and post-production stages, just as they can be when a score is listened to as music. But they of course are also integrated, during presentation to an audience and the production of a final master. The idea that somehow the rich multisensory fusion of audiovisual mediation is compromised by pointing out that sometimes sounds exist on tracks, and can propagate through discrete channels, and be sold as standalone albums seems to impart a kind of preciousness to sound, as though something in cinema breaks if one of its components is analyzed in distinction from others. These forms of what one can call ‘soundtrack denial’ are in my view a rather hyperbolic critical concern. A poetics of colocative sound design, as a form of rational and empirical inquiry, can readily enough acknowledge that soundtracks of various kinds do in fact actually exist, in one form or another.

### 13.4 The Perceptual Background

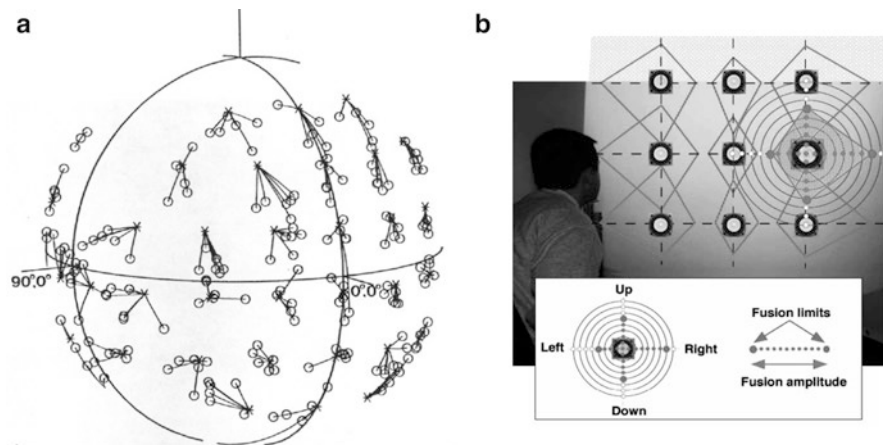
The empirical perceptual basis for colocative audiovisual perception can be parsed out across several domains that fit the cognitive model of bottom-up, top-down and recursive (feedback based) processes. As the experimental literature that pertains to the relevant perceptual matters is vast, I will have to limit the discussion to a summary of the main ideas and approaches. The most salient scientific literature in connection to the design of colocative systems can be defined under the general headings of: auditory localization, the auditory scene, multisensory enhancement, perceptual capture, and attentional resources. These will be summarized below to give a sense of the rich empirical perceptual capacities that are taken up in colocative audiovisual practices.

### 13.4.1 Psychophysics

According to principles established in the fields of environmental psychology (Gibson 1966) and auditory scene analysis (Bregman 1990), vision is superior for conveying information about objects, while audition is superior for providing information about events. A colocative system therefore has better ‘evental resolution’ relative to other forms of audiovisual display, enhancing attentional notice of new events tied to their dynamic visual representation. In the field of audio postproduction for moving image media, principles of sound design have been developed to improve overall clarity of information by assigning sound sources to different emitters in a spatialized array (Holman 2007). Spatialized audio information is clearer for perception and understanding, as masking and interference effects are minimized when many audio channels are not overloaded into a few loudspeakers. A wide body of empirical evidence has shown that distinct neurocognitive systems encode semantic and spatial positioning information in different pathways, via an action-based positional ‘how system’ and a semantic ‘what system’ (Richardson and Spivey 2000). Screen-based auditory colocalization of visual cues can be understood as a new ‘attentional resource’ (Wickens 2008) harnessed by colocative systems (discussed in more detail shortly).

The capacity of human listeners to localize sound in a free field has been the subject of extensive psychophysical investigation. Sound localization is performed by the auditory system in the horizontal plane by interaural differences of the phases and levels of wavefronts arriving differently at each ear, and also by differences in the dynamic envelopes of frequency components. In the vertical dimension, spatial discrimination occurs primarily through differences in harmonic spectra produced by the pinna (the outer ear). The most common scheme for measuring auditory localization is a bipolar coordinate system defined by azimuth and elevation, for horizontal and vertical degrees, respectively. Localization is generally better along a horizontal dimension, however vertical localization can be superior in the areas most peripheral to a listener (Makous and Middlebrooks 1990). With simple multimodal stimuli, localization is organized within ‘fusion areas’ of co-occurring auditory and visual signals (Fig. 13.6b). The MAA (Minimum Audible Angle) measures the precision, or the just noticeable difference, of sounds positioned in different locations. The MAA is frequency dependent and varies with laterality, and for anterior (frontal) sounds ranges from  $1^\circ$  to  $3^\circ$  for sounds close to the midline, to  $7^\circ$  or more for the most peripheral sounds. The perception of the movement of sound across a 2D plane is defined through the Minimum Audible Movement Angle (MAMA). MAMAs are affected by azimuth, duration of the sound, velocity (distance covered in time), bandwidth (wide band sounds feature smaller angles relative to pure tones), and training. Ignoring training effects, the optimal velocity for accurate auditory movement perception ranges between  $1^\circ/s$  and  $20^\circ/s$  (Carlile and Leung 2016). Listeners are also able to accurately identify the movement paths for trajectories of auditory cues. The figures below illustrate some of the relevant auditory spatial resolution findings from the empirical literature, which can inform





**Fig. 13.6** (a) Stimulus and response locations drawn on an imaginary sphere, as if looking from a point  $30^\circ$  to a subject's right and elevated  $10^\circ$ , with stimulus rendered as asterisks and open circles indicating five responses to localization cues. (Middlebrooks and Green 1991) (b) Visual-auditory fusion areas of perceived coincidence of simple multimodal stimuli. (Godfroy et al. 2003)

colocative system design as well as the poetics of works developed for such systems.

Our senses have evolved to work in a coordinated fashion, not in isolation as separate processing systems. The daily environment consists of multisensory data flows which provide concurrent and complementary information about the same external event or object, to maximize accurate detection and minimize error probabilities. Perceptual discrimination works to either integrate or segregate multisensory cues, depending on spatiotemporal coincidence, to decide whether these cues originate from the same or different sources and events. Spatial dislocation or incongruence in visuals and audio can result in conflict perceptual situations that need to be resolved, usually favoring one or another sensory modality, through sensory merging or separation (Ursino et al. 2011; McGovern et al. 2016). Incongruent stimuli received across sensory modalities are typically 'captured' by a 'dominant' modality through cross-modal modulation and dominance (e.g. the so-called 'ventriloquism effect'). Such sensory capture phenomena can be modified through training effects, which have been shown to reduce the magnitude of audio-visual ventriloquism and to narrow the temporal window for the perception of simultaneity.

Multisensory integration of auditory and visual modalities has an amplification effect on information processing, resulting in faster reaction times and lower error rates compared to unimodal stimuli. Multisensory integration involves neuronal activity in the Superior Colliculus, though enhancements in unimodal neurons also results from cross-modal modulation. These enhancements can be significant (called "super-additive") or moderate ("simple-additive"). Sensory degradation can also result through information overload of too many competing signals ("sub-additive")

enhancement) (Ursino et al. 2011). Crossmodal attention shifting, which can be either object or location-based, occurs when a location cue in one modality guides attention to stimuli in another modality in goal-driven (top down) or stimulus-driven (bottom up) recursive cognitive processes (Wright and Ward 2008).

### 13.4.2 *The Auditory Scene*

The soundscape, as a new design parameter for screen-based audio information which reconfigures spatial continuity, takes its other (than the soundtrack) conceptual motivation from the field of Auditory Scene Analysis (ASA). Early auditory research was to a large extent based on the “medical study of deafness” (Bregman 1990: 1) with focal areas on sensation of weak sounds, perception of loudness, and noise exposure. Auditory scene analysis (ASA) addresses “perceptual or ecological questions about audition” which are concerned with “higher level principles of organization.” ASA, which both describes a field of research and the actual perceptual analysis of sensory inputs performed by auditory processes, parses the complex mixture of environmental sounds which produces the perception of individual sounds. ASA studies how it is that our neurocognitive capacities are able to extract information about different acoustic events occurring simultaneously as part of the same auditory field of co-occurring sounds.

The term ‘stream’ in ASA is used in preference to the commonplace word ‘sound.’ Streams in an auditory scene communicate information about events, and an event may have multiple sounds. For instance, the word “shoe” has a “shhh” and an “oo” phoneme and could be said to comprise multiple sounds. Or an auditory stream such as footsteps can be taken as a single acoustic event or described as multiple sounds. An auditory scene is the field from which specific kinds of information are transduced from the environment by audition. ASA researchers frequently analyze sound recordings through spectrographic representations, as “there is some reason to believe that the human auditory system provides the brain with a pattern of neural excitation that is very much like a spectrogram” (7).

The central phenomena of ASA is the ‘grouping’ of streams, the principles of the ‘belongingness’ of sounds so that we attach them to streams. Scene analysis refers to the neurological processes which parse the signal inputs to the auditory system and synthesizes information from the environment through perception.

The goal of scene analysis is the recovery of separate descriptions of each separate thing in the environment. What are these things? In vision, we are focused on objects. Light is reflected off objects, bounces back and forth between them, and eventually some of it reaches our eyes. Our visual sense uses this light to form separate descriptions of the individual objects. These descriptions include the object’s shape, size, distance, coloring, and so on.

Then what sort of information is conveyed by sound? Sound is created when things of various types happen. The wind blows, an animal scurries through a clearing, the fire burns, a person calls. Acoustic information, therefore, tells us about physical “happenings.” Many happenings go on at the same time in the world, each one a distinct event. If we are to react

to them as distinct, there has to be a level of mental description in which there are separate representations of individual ones. (9–10)

In ASA, *stream* refers to the “perceptual representation” and *sound* or *acoustic event* refer to “the physical cause” (10). A stream is analogous to an object in visual experience, and auditory processing assigns properties to specific events since many things frequently occur at any given time, so that we do not perceive a “mush of... properties” (11) Audition provides complementary information to vision, conveying the acoustic energy that belongs to visual objects as they undergo events. While with visual objects things can occlude each other, in the sonic world it is as though all objects are transparent to each other. We can hear happenings through or behind objects or around corners.

This belongingness of sounds to streams, which convey the happenings of objects, appears to depend on familiar gestalt principles like similarity, proximity, continuity, closure and exclusive allocation. The auditory system constructs multiple representations or interpretations of stimuli in parallel, setting up a competition amongst the built up descriptions of environmental information. These alternate perceptions may enter awareness periodically or even be switched to with intentional effort. As in Gestalt, within-group and between-group “forces of attraction” (20) set up a competition for attentional representation. “In general, all the Gestalt principles of grouping can be interpreted as rules for scene analysis” (24) Sounds with strong similarities to each other, whether in spatial location, frequency, harmonic spectra, or intensities will tend to be grouped as belonging to the same event.

Scene analysis adds new conceptual resources that are based on principles well-known in visual perception, corroborating “the Gestalt theorists [who] saw the principles of organization as following from the general properties of neural tissue [and] they focused on similarities between the senses rather than on the differences” (36) ASA distinguishes Gestalt effects specific to audition, e.g. the communication of events in a transparent acoustic world. “For humans, sound serves to supplement vision by supplying information about the nature of events, defining the “energetics” of a situation” (37)

Having evolved in a world of mixtures, humans have developed heuristic mechanisms capable of decomposing them. Because the conditions under which decomposition must be done are extremely variable, no single method is guaranteed to succeed. Therefore a number of heuristic criteria must be used to decide how to group the acoustic evidence. These criteria are allowed to combine their effects in a process very much like voting. No one factor will necessarily vote correctly, but if there are many of them, competing with or reinforcing one another, the right description of the input should generally emerge. (33)

Auditory streams render information about “a coherent physical event” (39). To accurately model an event, scene analysis sets up multiple competing heuristics which combine to give a final shape to auditory experience. Despite the evident importance of spatial localization of sound sources, frequency (as in perceived pitch, not rate of occurrence) heuristics seem to be the most dominant in the overall perceptual ‘voting’ which parses the streams. Moreover, frequency information in the real world, as opposed to the laboratory, is complex, and so the harmonic

frequencies which comprise a unified timbre is perhaps the most salient cue against which other heuristics are compared. Timbre poses an interesting nuance to the idea that streams communicate happenings, since timbral information says much about the ‘body’ of a sound-producing object, being tied to many physical dimensions (e.g. material, density, weight, shape, surrounding environment) of a physical object. While ASA authors tend not pursue this thought of the ‘object-ness’ of timbre, if it is true that harmonic-frequency components of sounds play the strongest role shaping the perception of a stream, as appears to be suggested by the ASA literature, then this does point to an object-prioritization of the auditory system, albeit of course a focus that is concerned with something that is *happening* to that object.

Bregman indicates that spatial position is a critical heuristic for parsing an auditory scene:

We will now shift our attention to one of the strongest scene-analysis principles. This is the one that says that acoustic components that can be localized as coming from the same position in space should be assigned to the same stream. We know that the cue of spatial location is a very important one in allowing a listener to follow the words of one speaker in a noisy environment. If we cover one of our ears in such a situation our ability to select the desired voice gets much worse. (292)

This is a reference to the ‘cocktail party effect’ which allows us to single out a speaker’s voice in a crowded and loud social environment. One difficulty presented in understanding the role that spatial location plays in articulating streams is the tendency of too-similar sounds to ‘fuse’ into the same tone, an effect greatly exacerbated by the practice of using single frequency sinusoidal tones in ASA laboratory experiments. Bregman notes some psychoacoustic evidence that a sine wave tone emanating from different speaker positions can be distinguished by their locations. However, many of these experiments are designed to measure sensory “fusion, but not its opposite – independent localization” (295). This has been under-studied, and Bregman comments,

There will have to be extensive research on separate judgments of locations for simultaneous pure tones before we can know exactly how finely tuned the auditory system is for making separate frequency-specific localizations. The research will be hard to do because the mere simultaneity of the tones will tend to fuse them. We may need special indirect methods for the study of such localizations.

What is even more striking is that the other cues get a voice in deciding not only how many sounds are present but even where they are coming from in space. The auditory system seems to want to hear all the parts of one sound as coming from the same location, and so when other cues favor the fusion of components, discrepant localizations for these components may be ignored. It is as if the auditory system wanted to tell a nice, consistent story about the sound. (305)

Despite this intriguing insight into the semi-fictional character of percepts – the way the auditory system may sometimes privilege coherence over correspondence, to put it one way – an epistemological claim of ASA is that the non-pathological functioning of the auditory system produces accurate and corresponding representations of external stimuli. A ‘successful’ percept in this context is one where the

report by the subject matches the known actual stimulus presented by the experimenters.

### 13.4.3 *Attentional Resources*

Multiple Resource Theory (MRT) is a framework for addressing and predicting “multitask workload overload” (Wickens 2008: 451) through design considerations based on its ‘4D model’ of modality-based cognitive resources. Its robustness over time has been based on two premises: that there should be plausible neurological structures to support these four dimensions, and that the framework should be relatively easy to use in making design considerations. It has been applied often towards predicting breakdowns in dual-task workload, e.g. in situations similar to when use of a cellphone can degrade the performance of driving, due to competition for mental resources. The four dimensions are modelled as (Fig. 13.7):

1. Stages of Processing defined as the spectrum Perception-Cognition-Responding
2. Codes of Processing which distinguishes verbal and linguistic versus spatial (e.g. kinesthetic) activities
3. Modalities which distinguish auditory from visual perceptions.
4. Visual Channels which distinguish between focal and ambient vision.

The premise for design is that time-sharing performance in multi-tasking will be better when cognitive demands are parsed across multiple levels (e.g. visual/auditory, spatial/linguistic) rather than competing for the same attentional resources. This model has been translated into mathematical formulae which aim for predictive value in system performance evaluation. The specific neurological bases for the model are as follows (451):

1. “Perceptual-cognitive activity” is associated with the brain’s central sulcus area, while “motor and action oriented activity is anterior.”
2. Spatial and verbal activities are distributed across left and right hemispheres of the brain.
3. “Auditory and visual processes are distinctly associated with auditory and visual cortices.”
4. “Focal and ambient vision are supported by ventral and dorsal visual pathways, respectively.”

With this “neurophysiological plausibility” in place, Wickens discusses the design applicability as follows:

I felt it important that the dimensions of the model coincide with relatively straightforward decisions that a designer could make in configuring a task or work space to support multi-task activities: Should one use a keyboard or voice? Spoken words, tones, or text? Graphs or digits? Can one ask people to control while engaged in visual search or memory rehearsal?

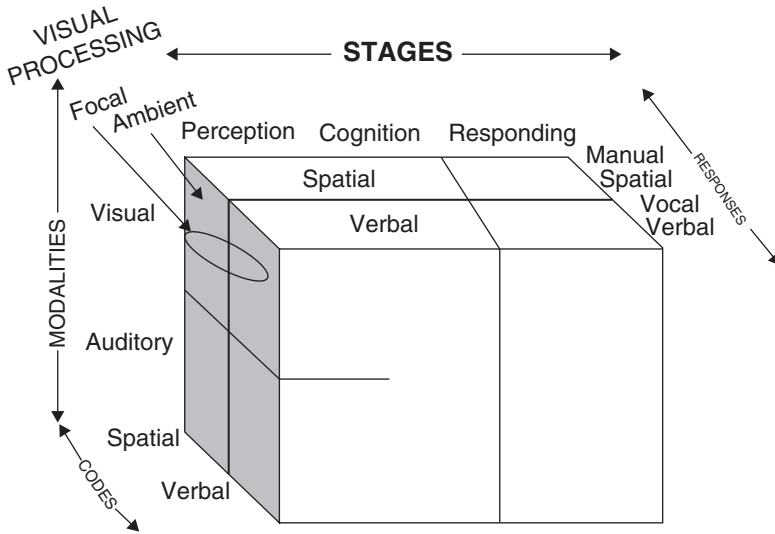


Fig. 13.7 Wickens' 4D model of attentional resources (450)

These two aspects of the model – its neurological underpinnings and its ease for use in design – has enabled it to “stand the test of time in its ability to account for three decades of dual-task research and to support design decisions.” Other research has added subprocesses to the model, differentiating within auditory resources other aspects like “spatial positioning, spatial quantitative, auditory linguistic, and auditory emotional resources” (451). The 4D model is practically implemented as a matrix of values, organizing the MRT cognitive resources in tables and matrices so as to obtain and sum values within cells. There are two types of quantification produced by such grids, a “demand component” and “resource conflict component” or a “total interference component” (452). The values are based on a gradient of “demand component” whereby tasks are rated as being either “automated (= 0), easy (= 1) or difficult (= 2), “independently of which resource may be demanded” which means that the resources are not weighted. In this way a dual task demand score can range “from 0 (two automated tasks) to 4 (two difficult tasks).” The resource conflict values can range from 0 to 8, as they are taken as the sum of the two activities competing for the same resource. The scores produced by the matrices serve as predictive indication of breakdown in competing tasks:

Only in the region where overload is imposed by multiple tasks does the multiple resource theory make an important contribution to mental workload by predicting how much performance will fail. (453)

Wickens suggests that the tactile should be considered as an additional modality in future research. However the tactile itself is a subdivision of the category of the “haptic” and is paired theoretically with the “kinesthetic,” (Tappeiner 2013) which

suggests a new modality pairing in a similar manner to the other modal pairs, and producing the new contrast sets of verbal/spatial, visual/auditory, focal/ambient + tactile/kinesthetic. Wickens' dimension of "response" includes the kinesthetic as an output (completing a task) rather than as an input (informing a task), and thus the very notion of a spatial task is implicitly based on a modality that could also be its own independent dimension as a resource. For example, in driving the kinesthetic act of steering is an output interacting with the input of the kinesthetic sense of the driver's body and the car's roll and lean, which are distinct modal dimensions.

The separation of spatial and verbal resources seemingly accounts for the relatively high degree of efficiency with which manual and vocal responses can be timeshared, assuming that manual responses are usually spatial in nature (tracking, steering, joystick or mouse movement) and vocal ones are usually verbal (speaking). (Wickens 2002: 166)

With respect to colocative audiovisual systems, this relationship of attentional resources to task performance more directly connects to interactive media and applications areas where viewer-listeners are also users. Games, simulations, telepresence and other interactive and immersive environments engage users with cinematic experiences in which task performance takes on a clearer role in the construction of a soundscape than would be the case with linear media like narrative film and video.

### 13.5 Phenomenology

Since its founding by Husserl, phenomenology has developed a vast literature encompassing a family of related approaches and methodologies, with variations in its schools of thought characterized by antecedent descriptors – transcendental, cognitive, structural, hermeneutic, existential, social, psychological, biological, empirical and experimental phenomenology, to name the more prominent kinds. What all these approaches have in common is a grounding in first person subjective perspective validated by intersubjective corroboration. Despite the proliferation in styles of phenomenological inquiry, one can broadly distinguish all forms of phenomenological thought with respect to how the natural attitude is regarded, and the specific status of the *epoché* – the suspension or 'bracketing' of the natural attitude, regarding all phenomena strictly as constitutive acts of consciousness, and deprived of external and habituated reality.

The phenomenological explication I will describe below is somewhat mid-way between a 'full embrace' of the natural attitude through lifeworld analysis – an alternative approach developed by thinkers such as Alfred Schutz and Harold Garfinkel – and structural, also sometimes called cognitive and transcendental phenomenology, as developed originally by Husserl. 'Transcendental' in this sense means inquiry into the conditions for the possibility of consciousness, which is worth noting as sometimes in English the term is taken to mean 'transcendent' which would imply something beyond consciousness. The Husserlian style inquires into permanent ideal structures of mind, or 'eidetic essences' that are in principle

intersubjectively available to anyone's consciousness, and which are revealed through the process of epoche. The middle approach I will illustrate – which can be described as an inquiry into the cognitive structures of experience through lifeworld analysis – has affinities with Merleau-Ponty's focus on embodied experience and Horst's (2005) argument that most psychophysical data is in fact phenomenological in character. As Merleau-Ponty's writing and ideas are well-known, I will discuss Horst's conception of the phenomenological character of psychophysical data in more detail.

Horst's conception of phenomenology can be characterized as an approach that dispenses with the procedure of epoche, since in his view what is most critical for an approach to be considered phenomenological is that it rely on a first person report that has intersubjective validity. In the context of experiment design, the 'contamination' of perceptual data by attention, decision making and memory – which are understood as higher order cognitive processes – is an ongoing debate within the psychophysical fields. These fields may include other areas and terms used, such as "psychophysiology," "neurocognition" and so forth. It is not the goal here to elucidate all the varieties of empirical research into perception, and so I will keep to the term "psychophysics" to describe experiment-based approaches to perception. Perceptual experiments tend to share certain common features, for example: the use of very simple stimuli (e.g. sine waves or white noise); a lab-based setting where the researcher can manipulate variables to generate effects; a high number of randomized trials; and some manner of eliciting the experience of the human subject, for instance by pressing a button, moving a lever, or merely looking in some direction while an eye or head tracking apparatus documents saccadic movements or head position, and so forth. Because the bulk of psychophysical research depends on obtaining a human response to stimuli in order to yield a data point, Horst argues that in a fundamental sense, such experiments have at their core phenomenological data.

There are forms of perceptual experiments which study phenomena at a more neurological level, in which perceptual responses can be detected in the neurons but which never register in conscious experience. Research into blindsight, or "the ability of individuals with blindness to detect and respond to visual stimuli despite lacking awareness of having seen anything" (Blindsight) is of this kind, where there is no phenomenological correlate to a neurological event. However, most psychophysical research is dependent on obtaining first person reports of subjective experiences, which Horst argues makes phenomenology central to the psychophysical enterprise.

[S]ince psychophysics is the major supplier of data that constrain theories of perception, phenomenological properties make up an important portion of the data that theories of perception try to explain. (Horst 2005: 1)



Psychophysics is widely regarded as the portion of psychology that really has become scientific, and it depends very heavily upon phenomenology. On the one hand, its domain includes phenomenologically described mental states (percepts). On the other hand, its methodology requires subjective access to the first-person, experiential, phenomenological character of these percepts. And without such a phenomenologically-based psychophysics we lose many of the data that it is the business of theoretical psychology of perception to explain. (14)

[T]he best established part of scientific psychology is essentially committed to phenomenological properties of mental states, both as its domain and as a necessary part of its methodology. (16)

So long as one allows that subjective reports may be elicited in ways other than through verbal accounts – as specified by traditional phenomenological approaches – and if one does not require of phenomenology that it always embody the methodology of epoche, Horst’s argument is compelling and can be considered as a lifeworld-grounded conception of phenomenological research. Instead of the complete suspension of the natural attitude, what reveals the eidetic-cognitive structures of perception – the form of the constituting acts of consciousness – in psychophysical research is experimental method: multiple subjects, many randomized trials, and methods for obtaining the first person report of experience. Experimental psychophysical phenomenology can therefore be understood both as being grounded in the natural attitude and lifeworld, and also aiming for Husserlian ideal types of cognitive structures.

For our purposes of focusing on the poetics of colocative sound design, every home-based entertainment system provides the intersubjective context which in principle is available to anyone to experience for themselves the perceptual condition I will now describe. Namely, I will claim that much of the phenomenon of visual capture in media – the sensory dominance of auditory location by visuals, also called “the ventriloquism effect” – is of an attentional, not a perceptual, character, and can be overdriven through the development of new perceptual and attentional thresholds as one gains embodied experience with colocative audiovisual systems.

With regards to the ventriloquism effect – the visual capture of sound so that the audio emitters’ actual location is not disruptive to visually mediated experience – we can distinguish a perceptual and an attentional component. Since the norm with audiovisual media, especially in the context of home-based systems, involves spatially dislocated sound sources relative to the image area, much of our acceptance of sound emanating from dislocated speakers has formed in our habitus with media, rather than in hard-wired neurological or cognitive capacities.

Habitus is a term mobilized by Bourdieu to refer to those embedded and embodied dispositions produced and reproduced in the mediation between objective – or material – conditions and subjective experience; or, to put it more simply, the ways in which – without any conscious effort on our part – our habits are intertwined with our social and material habitat. (Lacey 2017: loc 5709)

During a home viewing of audiovisual media, little attentional effort is required to perceive the sound emanating from spatially dislocated point sources. Recalling

Bordwell's point that bottom-up and top-down feedback processes cannot proceed all the way up or down through the processing levels, we can describe this phenomenologically proposed attentional shift as a cognitive operation that takes place between his posited levels of Appropriation and Comprehension. The ease with which it is possible to discern that dialogue playing in one's living room setup is actually emanating from a spot somewhere beneath the bottom of the screen, or that sound effects and music are coming from sources well beyond the screen edges, points to the non-perceptual component at work in the daily ventriloquist effect of our lifeworld media habits. When sound and image emanate from very close spatial positions, the perceptual component of the visual capture of sound takes over at a level in between perception and comprehension, and no amount of willful appropriation could change what we experience (or at least, much cognitive effort must be expended to try to override the perceptual capture). The attentional component becomes more apparent as the spatial incongruence of a sound and its image increases. It is phenomenologically possible to discern three levels of sound-image localization with respect to the ventriloquism effect: (1) 'true' multimodal fusion, wherein the audio emanates from within the spatial fusion areas described by Godfroy et al. (2003) because of their close spatial proximity; (2) Visual capture, whereby incongruence of audio and visual spatial positions is resolved by cognitively favoring the image as the locus of the sound when sound and image sources are further apart; and (3) Attentional acceptance of dislocated point sources, established by habituated experience with audiovisual systems. One can experience these three phenomenological structures simply by moving one's speakers around in the living room with respect to the video screen and its image, a straightforward and informal experiment which offers intersubjective corroboration for this set of distinctions.

As perceptual and attentional thresholds shift through habituation to colocative sound design, during media production one can learn to notice when sounds are 'out of place' with their image in a colocative display, because the emplaced combination of sound and image produces new expectations and sensitivities toward audiovisual media. During the audio post-production phases of colocative sound design, the sense of colocation between sound and image is just another aspect of the overall multisensory experience to pay attention to. This reinforces experiential dimensions with respect to traditional dislocative media, i.e. once one has become acclimated to colocative media, it can sometimes seem odd that, in a typical home setup, the sounds associated with moving images are in fact often located several feet away, to the left of the right of the screen. Put another way, the attentional component of the ventriloquism effect can be understood as an "audile technique" (Lacey 2017: loc 5713) building everyday "auditory capital" which is

accumulated via the deliberate cultivation of particular listening skills, but also via a process of gradual "incorporation" of historically contingent listening techniques. In its objectified form, auditory capital is accumulated in specific material, acoustic and media objects, for example both the sonic arts and the various technologies that mediate sound. (loc 5693)

Colocative audiovisual systems reveal that our usual manner of perceiving synchronized media through spatially dislocated sensory information channels actually involves an attentional habit, whereas we may have hitherto regarded it as a form of natural perception.

### 13.6 Acoustic Considerations

A brief discussion of the acoustic (electrical-mechanical) aspects of colocative sound design is warranted. Both the Allosphere and Pixelphonic systems rely on a spatialized array of audio emitters which are essentially point sources, which can in principle be utilized with either discrete channels (i.e. one channel of sound assigned to each emitter), or make use of a diffusion approach whereby the same signal is routed to multiple emitters. Another approach, matrixed audio information, was used in older Dolby Stereo systems but has been surpassed as memory storage and processing speeds have improved, removing past limits on the number of channels that can be accommodated by today's digital media (Kerins 2010: loc 634). There are today considered to be three kinds of audio representations: channel-based, object-based, and transform domain (He 2017). In typical film, video and musical media, audio is channel based, in the sense that discrete audio tracks are assigned to a more limited number of channels from which they are eventually encoded onto the final medium. For instance, a film's mix may have hundreds of audio tracks where clips are sorted according to a practical information architecture – the most common categories being Voice, Effects and Music – but during mixdown these hundreds of discrete tracks are bussed to a final output format, such as stereo or 5.1 surround (five channels and one low frequency effects channel). With object-based audio, which is more typical of gaming platforms, sounds are attached to virtual objects within virtual spaces, though a channel-based approach to the sonic construction might still be used for cut scenes of linear animation that are placed within the gameplay. In games, the assignment to an audio emitter in a spatialized array – for instance, a speaker in one's living room – is performed in real-time and based on a user's interactions. A transform-domain representation of audio aims to reproduce an originally propagating wavefront. For example, one can arrange dozens of microphones at a locale, and play them back through a similar array of speakers, with the idea of reconstituting the original acoustic wavefronts. As hardware platforms, both the Allosphere and Pixelphonics could accommodate, in theory at least, any of these forms of sonic representation. It is a matter for the software and signal distribution to determine what approach to audio representation is being used in a given instance. Moreover, one possibility with Pixelphonics is the assignment of audio channels beyond the screen area, as nothing in principle confines each discrete channel to the screen. This is also possible in a sense with the Allosphere, but since it is designed as an immersive sphere, any non-screen audio channels would have to be located outside the acoustic envelope of the display.

With a vibrating screen concept, the acoustic phenomenon of planar dispersion – whereby the surface beyond the exciter also vibrates in a gradually attenuating fashion – fills up the screen space in between each audio exciter with sound, which could not occur with a speaker array. Pixelphonics produces ‘smeared’ point-sources relative to a speaker-based system like the Allosphere, because the screen vibrates beyond the attached exciter in an attenuated gradient across the entire surface. In both systems, the spatial distance between emitters would ultimately be determined through factors by which the scale effects of the display size are incorporated into the hardware configuration and system design.

### 13.7 Synthesizing a Poetics of Colocative Sound Design

Having situated audiovisual colocation in its most important theoretic and practical contexts, we can now offer some provisional statements as to what constitutes its form of “screen-centric” (Kerins 2010: loc 11,349) poetics. We can say: *Audiovisual colocation emplaces sounds with their visual sources, and in doing so increases the resolution of the perceived and comprehended rendering of events on screen. It requires the construction of a soundscape which is a new design requirement for screen-based audio information, so as to avoid disorienting and dislocative effects when integrated with other continuity techniques. A soundscape reconfigures spatial continuity with reference to the rendering of visualized sonic events, which has repercussions for shot planning and scene editing, information visualizations and overlays, telepresent operators, virtual objects, and other content types and applications for colocative displays. Audiences and users of colocative works may develop a refined perceptual and attentional sense of when such sounds are or are not adequately emplaced with their visual cues, via a reduction in the magnitude of the ventriloquism effect. There are three primary colocative techniques: 1) to emplace sounds by locating them within visual-auditory fusion areas that are formed at the level of perceptual organization; 2) to create new movement paths and trajectories for sounds within the screen area; and 3) to use sound to draw attention to images in the ambient visual field which may not be at the center of focal attention, through crossmodal attention shifting. Like many developments in new audiovisual technologies, there is an implicit gesture toward realism and fidelity to natural perception, since in everyday experience sounds do in fact emanate from object-sources undergoing events in a spatiotemporal context.*

These poetic principles apply to any application of colocative audiovisual media: games, film, animation, teleconferencing, simulation-based training, scientific visualization, command and control installations, process control displays, linear or interactive media, etc. The principles would not apply, of course, in instances where a contrarian aesthetics is used that deliberately refuses colocation, in an analogous manner that some avant-garde works eschew synchronized sound. However, even this eschewing can itself be a reference to and take as its baseline mainstream principles and practices.

This poetics associated with colocative sound design will in some ways run counter to certain tendencies in criticism, aesthetics and commentary. There is a trope one frequently comes across, to the effect of sound's ephemeral and ungraspable character, a sense of its being bodiless and invisible, exceeding conceptual and perceptual capacities to pin it down, and since it is time-based, lacking edges or boundaries. Sound is said to have a "radical heterogeneity" (Connor 2013: loc 2303). The body of sound is "diffuse and intermittent. It is intense, but evanescent. It has no place to reside or come to rest....The sound of the screen is not primarily 'on' the screen, but in the listener" (loc 2473). "Sound, as Michel Chion has emphasized, has no frame. There is nothing to hold it in, so there is nothing to hold it in" (loc 2466). "Sound, like water, has no border, no clear outline to distinguish it from not it" (Kim-Cohn 2017: loc 1630. "What I want to suggest is that sonic bodies don't exist" (loc 1636). Sound "suggests a spatiality less concrete and more ephemeral"(LaBelle 2017: loc 7146)

Can we not understand the spaces generated by sound as always already unfixed, vibrant, and coproduced? Built from a multiplicity of events and actants that comingle and conflict moment by moment, and whose invisibility may throw into question what constitutes identity? (loc 7144)

While much about this trope makes a kind of exaggerated intuitive sense by playing up sound's temporal dimensions, we saw above that in fact psychophysical experiments have identified spatial 'fusion areas' for coincident auditory-visual stimuli, that trajectory-based spatial shapes of sound are accurately identifiable, that the auditory system has considerable localization capacities, and that hearing in general renders object-based events. By dynamically associating sounds in time-based media with their visual cues, colocative sound design breaks with this critical-rhetorical sensibility of sound lacking a body, spatial definition and boundary. Rather, it constructs an emplacement effect which produces an alternative experience to what has been characterized as "schizophonic sound (i.e., sound that is split from its source and recontextualized)" (Herbert 2017: loc 5976). Emplacement is the colocative equivalent to an experiential concept like "in sync" that emerged with synchronization practices. Just as synchronized media can sometimes be out of sync (usually read as a disruptive mistake but sometimes exploited for creative purposes), emplacement likewise gives rise to the new possibility of sound being 'non-emplaced.'

Audiovisual colocative media can be integrated with and complementary to established techniques and experiences involving surround sound. The Allosphere, for instance, is also a surround array, albeit one that completely envelopes one within a spherical environment. The Pixelphonic system articulates the screen area itself, which can be adjoined to a surround system, and its screen is scalable and flexible to the point of offering some degree of envelopment as well, depending on its spatial construction and scale. Screen-centric emplacement, in other words, can coexist with the envelopment approach that has been the dominant focus of spatialized cinematic audio since Fantasound.

Colocated audiovisuals have been imagined before, but now there are at least two working systems which pose solutions to this possibility.

But if sound could be localized either on screen or in the orchestra pit, then it could also, at least theoretically, be reproduced at different spots on the screen, depending on where the person speaking is portrayed. Never to my knowledge realized, this project, reported by J.C. Kroesen in July, 1928, demonstrates the extent to which early technicians assumed the necessity of tying sound to the image. “The screen,” said Kroesen, “should be divided and so arranged that sound will be reproduced only at or as near the point of action as possible. (Altman 1992: 48)

Ideally, as stereo expert Harvey Fletcher has pointed out, every square inch of the screen should have a separate speaker for sound, while an infinite number of speakers and tracks would be needed to duplicate sounds emanating from offscreen space. (Belton 1992: 165)

Whether or not colocative sound design brings media “closer to life,” one can suggest that it may increase the presence aspect of immersive mediated experiences, which is a discussion for another place and time.

## References

- Allosphere Research Facility (2017) University of California Santa Barbara. <http://www.allosphere.ucsb.edu/index.html>
- Altman R (1992) Sound space. In: Altman R (ed) *Sound theory sound practice*. Routledge, New York/London
- Belton J (1992) 1950s magnetic sound: the frozen revolution. In: Altman R (ed) *Sound theory sound practice*. Routledge, New York/London
- Blindsight. Merriam-Webster Medical Dictionary Online: <https://www.merriam-webster.com/medical/blindsight>
- Bordwell D (2002) Intensified continuity visual style in contemporary American film. *Film Q* 55(3):16–28
- Bordwell D (2007) *The poetics of cinema*. Routledge, New York
- Bregman A (1990) *Auditory scene analysis*. MIT Press, Cambridge, MA
- Bresson R (1985) Notes on sound. In: Weis E, Belton J (eds) *Film sound: theory and practice*. Columbia University Press, New York
- Carlile S, Leung J (2016) The perception of auditory motion. *Trends Hearing* 20:1–19
- Chion M (2013) Sensory aspects of contemporary cinema. In: Richardson J, Gorbman C, Vernallis C (eds) *The Oxford handbook of new audiovisual aesthetics*. Oxford University Press, Oxford. Amazon Kindle edition
- Coleman S (2005) *Badgered. On: a collection of 2005 academy award nominated short films* (DVD). Magnolia Home Entertainment
- Connor S (2013) Sounding out film. In: Richardson J, Gorbman C, Vernallis C (eds) *The Oxford handbook of new audiovisual aesthetics*. Oxford University Press, Oxford. Amazon Kindle edition
- Donnelly KJ (2013) Extending film aesthetics: audio beyond visuals. In: Richardson J, Gorbman C, Vernallis C (eds) *The Oxford handbook of new audiovisual aesthetics*. Oxford University Press, Oxford. Amazon Kindle edition
- Gibson JJ (1966) *The senses considered as perceptual systems*. Houghton Mifflin Co., Boston
- Godfroy M, Roumes C, Dauchy P (2003) Spatial variations of visual-auditory fusion areas. *Perception* 32:1233–1245

- He JJ (2017) Literature review on spatial audio. In: Gan WS, Kuo CCI, Zheng TF, Barni M (eds) Spatial audio reproduction with primary ambient extraction, Springer briefs in signal processing. Springer, Basel
- Herbert R (2017) Sonic subjectivities. In: Cobussen M, Meelberg V, Truax B (eds) The Routledge companion to sounding art. Routledge, New York. Amazon Kindle edition
- Höllerer T, Kuchera-Morin J, Amatriain X (2007) The allosphere: a large-scale immersive surround-view instrument. In: EDT'07: proceedings of the 2007 workshop on emerging displays technologies: images and beyond: the future of displays and interaction, San Diego, CA, USA, Aug, 2007, Article No. 3. doi 10.1145/1278240.1278243
- Holman T (2007) Surround sound up and running. Focal Press, New York
- Horst S (2005) Phenomenology and psychophysics. *Phenomenol Cogn Sci* 4:1–21
- Kerins M (2010) Beyond Dolby (stereo): cinema in the digital sound age. Indiana University Press, Bloomington
- Kim-Cohn S (2017) Dams, weirs, and damn weird ears— post-ergonal sound. Auditory capital, media publics and the sounding arts. In: Cobussen M, Meelberg V, Truax B (eds) The Routledge companion to sounding art. Routledge, New York. Amazon Kindle edition
- Kuchera-Morin J, Wright M, Wakefield G, Roberts C, Adderton D, Sajadi B, Höllerer T, Majumder A (2014) Immersive full-surround multi-user system design. *Comput Graph* 40:10–21
- LaBelle B (2017) Restless acoustics, emergent publics. In: Cobussen M, Meelberg V, Truax B (eds) The Routledge companion to sounding art. Routledge, New York. Amazon Kindle edition
- Lacey K (2017) Auditory capital, media publics and the sounding arts. In: Cobussen M, Meelberg V, Truax B (eds) The Routledge companion to sounding art. Routledge, New York. Amazon Kindle edition
- Makous JC, Middlebrooks JC (1990) Two-dimensional sound localization by human listeners. *J Acoust Soc Am* 87:2188–2200
- McGovern DP, Roudaia E, Newell FN, Roach NW (2016) Perceptual learning shapes multisensory causal inference via two distinct mechanisms. *Sci Rep* 6 Article number 24673. <https://www.nature.com/articles/srep24673>
- Middlebrooks JC, Green DM (1991) Sound localization by human listeners. *Annu Rev Psychol* 42:135–159
- Richardson DC, Spivey MJ (2000) Representation, space and Hollywood Squares: looking at things that aren't there anymore. *Cognition* 76(3):269–295
- Stillwell RJ (2013) Audiovisual space in an era of technological convergence. In: Richardson J, Gorbman C, Vernallis C (eds) The Oxford handbook of new audiovisual aesthetics. Oxford University Press, Oxford. Amazon Kindle edition
- Tappeiner H (2013) Robotics: what is the difference between tactile feedback and haptic feedback? Quora.com. <https://www.quora.com/Robotics-What-is-the-difference-between-tactile-feedback-and-haptic-feedback>
- Ursino M, Magosso E, Cuppini C (2011) Sensory fusion. In: Cutsuridis V, Hussain A, Taylor JG (eds) Perception-action cycle: models, architectures, and hardware. Springer series in cognitive and neural systems. Springer International Publishing AG, Basel, pp 23–62
- Wickens CD (2002) Multiple resources and performance prediction. *Theoretical issues in ergonomic. Science* 3(2):166
- Wickens CD (2008) Multiple resources and mental workload. *Hum Factors* 50(3):449–455
- Wright RD, Ward LM (2008) Orienting of attention. Oxford University Press, Oxford

# Chapter 14

## Language Technology and 3rd Wave HCI: Towards Phatic Communication and Situated Interaction



Lars Borin and Jens Edlund

**Abstract** In the field of language technology, researchers are starting to pay more attention to various interactional aspects of language – a development prompted by a confluence of factors, and one which applies equally to the processing of written and spoken language. Notably, the so-called ‘phatic’ aspects of linguistic communication are coming into focus in this work, where linguistic interaction is increasingly recognized as being fundamentally situated. This development resonates well with the concerns of third wave HCI, which involves a shift in focus from stating the requirements on HCI design primarily in terms of “context-free” information flow, to a view where it is recognized that HCI – just like interaction among humans – is indissolubly embedded in complex, shifting contexts. These – together with the different backgrounds and intentions of interaction participants – shape the interaction in ways which are not readily understandable in terms of rational information exchange, but which are nevertheless central aspects of the interaction, and which therefore must be taken into account in HCI design, including its linguistic aspects, forming the focus of this chapter.

### 14.1 Introduction

An interesting – and appropriate – point of departure for a chapter on the intersection of language technologies and 3rd wave HCI could be some relevant and spectacular communicative faux pas reported in media in recent years. For example:

- The inability of IBM’s Watson to get to grips with the proper usage and nuances of colloquial expression (Smith 2013)

---

L. Borin (✉)  
University of Gothenburg, Gothenburg, Sweden  
e-mail: [lars.borin@svenska.gu.se](mailto:lars.borin@svenska.gu.se)

J. Edlund  
KTH Royal Institute of Technology, Stockholm, Sweden



- Microsoft’s Twitter bot Tay turning into “a Hitler-loving sex robot” in 24 hours (Horton 2016; Wikipedia 2017)
- Facebook bots making up their own – non-human-understandable – language (LaFrance 2017)

Importantly in this context, these all involve attempts to use language technology (LT; strictly speaking text technology; see below) in order to design human-like human–computer interaction (HCI). These examples also highlight some major stumbling blocks that may materialize in this kind of endeavor.

According to practitioners, the HCI field is currently going through a paradigm shift, into its “third paradigm” (Harrison et al. 2007, 2011) or “third wave” (Bødker 2006, 2015).

This involves a shift in focus from (1) framing the requirements on HCI design primarily in terms of (mainly context-independent) information flow, to (2) a view where it is recognized that HCI – just like interaction among humans – is indissolubly embedded in complex, shifting contexts. These – together with the different backgrounds and intentions of interaction participants – shape the interaction in ways which are not readily understandable in terms of rational information exchange, but which are nevertheless central aspects of the interaction, and which therefore must be taken into account in HCI design.

Forming an increasingly important component of HCI, LT and the related discipline of linguistics<sup>1</sup> have seen similar developments in recent decades. In the next section, we briefly outline relevant developments in linguistics, in particular those often summed up under the heading phatic commun(ica)tion.

The term ‘language technology’ is of fairly recent origin, coined as an umbrella term for two historically distinct (and largely separate) fields: one dealing with automatic linguistic processing of written texts – called ‘text technology’ here, but more commonly labelled as computational linguistics, natural language processing, or language engineering in the literature, and the other – speech technology which addresses the nature of spoken language, ranging from signal processing and engineering issues of automatic speech analysis and synthesis to issues of cognitive science and psychology relating to the pragmatic and social goals and effects of variations in spoken interaction.<sup>2</sup>

Even though there is an increasing interaction between these two strands of LT, they are still distinct enough that it makes sense to treat them separately in the

---

<sup>1</sup>The relationship between LT and linguistics is a complex one (Reiter 2007; Jones 2007; Wintner 2009), but since the subject matter of both is human language and linguistic behavior, and since they share a common heritage, including a nontrivial set of ontological and methodological assumptions, they nevertheless tend to come up against the same problems.

<sup>2</sup>There is now also an emerging field of sign language technology, dealing with the problems of processing the various signed languages of the Deaf (the Ethnologue lists no less than 177 sign languages worldwide; Simons and Fennig 2017). This field is still very much in its infancy, and we will not be able to discuss it further here.

present context. Text technology is discussed in sect. 14.3 and speech technology in sect. 14.4 below.<sup>3</sup>

In this chapter, we will not concern ourselves with existing spoken dialog system technology (e.g., Jokinen and McTear 2009), but rather look at trends in basic LT research with the potential to result in better tools for designing all kinds of linguistic interaction between people and computers. As issues of computational grammatical or semantic analysis of language have been intensively researched, we focus on emerging aspects of LT which so far have not been extensively addressed, but have started to attract the attention of many practitioners in the field, especially within the context of 3rd wave HCI.

## 14.2 Phatic Communion in Linguistic Interaction

Phatic communion is a term introduced by the British anthropologist Malinowski (1923), designating “a type of speech in which ties of union are created by a mere exchange of words”, e.g., “pure sociabilities and gossip” (Malinowski 1923: 315). The term was later adopted and given a partly new content by the linguist Jakobson (1960), and has since generated a respectable number of linguistic publications (e.g., Laver 1975; Coupland et al. 1992; Žegarac and Clark 1999; Stenström and Jørgensen 2008; Senft 2009) on phatic communication and phatic linguistic devices. The concept of phatic linguistic devices builds on the realization that most human interactions that utilize language are complex, generally fulfilling more than one need simultaneously, and thus may include elements of phaticity even in linguistic interactions serving mainly some other purpose, where said elements are typically signalled by specific linguistic devices.

Similarly to the “second paradigm” of HCI, linguistics and LT too are developed around the basic premise that a language’s primary function is communication about facts, and that it has emerged as a consequence of increasing need for goal-directed activity coordination (hunting, farming, etc.). Contrary to this, Dunbar (1996, 2004) has suggested that the primary driving force behind the evolution of languages in humans has been the need for a social bonding mechanism which would allow the inclusion of larger social group sizes (as opposed to primate grooming), where gossiping has been established as the most prototypical linguistic activity.

Humans have utilized language for at least 100,000 years, and possibly for as long a half a million years, while writing appeared on the scene only less than 10,000 years ago. All human groups have a spoken language (and/or a signed

---

<sup>3</sup>To confuse matters further, LT – especially the written-language kind – is often treated as belonging to the field of artificial intelligence (AI), and in fact the three cases of communication breakdown mentioned at the beginning of this chapter have generally been reported as failures of AI. There is no doubt that these involve specifically LT, and also no doubt that LT is a thriving scientific discipline in its own right, with numerous conferences and journals.

language in special circumstances), while writing is a regular feature of perhaps at most one fifth of the world's languages today. Universal or near-universal literacy – or even a recourse to a widely used written standard language – is a very recent and quite limited phenomenon characterizing perhaps as little as 5% of all languages. Thus, writing is something of a mayfly in the larger picture of languages. However, for various practical reasons, it has informed linguistic inquiry and LT research at the expense of everyday spoken language, to a point where it has been contended that linguistics is fundamentally written-language biased (Linell 1982). Studies of the spoken origins of language have not always been widely accepted as an entry point for linguistic study. In the 1800s, the issue was considered so infested by opinions and speculation that the influential Société de Linguistique de Paris explicitly prohibited any discussion of it in a bylaw: “The Society will accept no communication dealing with either the origin of language or the creation of a universal language” (Quote from Stam 1976). To date, the genesis of communication in general and language in particular is conspicuously underrepresented in linguistics, as Bickerton (1990: 105) notes: “While this literature [on the origins of language] includes contributions from anthropologists, ethologists, palaeontologists, philosophers, and members of other assorted disciplines, hardly any of it is written by professional linguists”.

So while the primary mode of human language use is spoken interaction, the form of language most studied by linguists and computational linguists is written monological sequences of statements of fact. Arguably, these fields have been shaped by their emergence at a time, place, and a societal stratum characterized by high literacy, a standardizing prescriptive attitude to language, and a monopoly on the means of mass communication often exemplifying what Spolsky (2004) has termed “ideological monolingualism”. Notably, this has created an ideological climate which has provided fertile ground for ideas about linguistic deficit, such as Bernstein's restricted and elaborate codes (Bernstein 1973; Jones 2013), or Hansegård's semilingualism (Hansegård 1968, 1990; Martin-Jones and Romaine 1986).

After standing tall for over 200 years (Anderson 2006; Rutten 2016), this edifice started crumbling at the end of the twentieth century, with the consequences that we see a reemergence onto the public scene of such phenomena as language variation and phatic communication, all ignored by linguists and language technologists alike. Social media are arguably more like (some forms of) everyday spoken interaction than written Medline abstracts, newspaper articles, or even novels. Natural spoken interaction turns out to be dominated by social topics, accounting for two thirds of speaking time (Dunbar 2004: 105f).

Dunbar (1996, 2004) does not refer specifically to phatic communication and its treatment in linguistic research, but it is clear that the linguists' notion of phatic communication and phatic linguistic devices correspond almost perfectly to Dunbar's “conversation on social topics/gossip” (Dunbar 2004: 105).

The notion of phatic communication provides us with a possible explanation on both what went wrong with Watson, Tay and the Facebook bots (see above) and on where to look for ways of improving LT for enhancing HCI. Both language variation

and cursing can be described, at heart, as phatic linguistic devices. The emergence and maintenance of ‘lects’ (dialects, sociolects, registers, etc.) serve a social function, and their production and recognition in interaction serves to establish “common ground” in the social sense (Monk 2003), signalling: “This is the kind of person I am. Are you one of us?” The use of certain emotionally charged linguistic items and constructions similarly can function as socially cohesive devices. Charles Darwin himself seems to have been on the same track:

Languages owe their origins to the imitation and modification of various natural sounds – such as the voices of other animals, and man’s own instinctive cries – aided by signs and gestures. When we look at sexual selection, we shall see that primeval man – or rather some early progenitor of man – probably used his voice largely, as does one of the gibbon-apes at the present day in producing true musical cadences; we may conclude from a widely-spread analogy that this power would have been especially exerted during the courtship of the sexes, serving to express various emotions, as love, jealousy, triumph, and serving as a challenge to their rivals. (Darwin 1871: 56).

It is worth noting that anything intended as a cohesive device can also cause estrangement, if the interlocutor understands it as an attempt to pigeonhole her as something which she does not wish to be a part of.

Within this context, it has been suggested that new media – in particular social media – correlate with and promote “phatic culture” (Miller 2008) and sometimes even constitute “phatic technology” (e.g., Vetere et al. 2005; Wang et al. 2011). This raises some fundamental questions about HCI design. An important part of the third-paradigm program seems to be that HCI should not and cannot be neatly segmented into content/activity on the one hand, and (independent) context, on the other (Dourish 2004). Context here is construed to encompass not only the spatial and temporal setting of an activity, but crucially also its social setting- including the social roles and status of participants, i.e., the domain of phatic communication and phatic linguistic devices. Consequently, and to the extent that the HCI design involves linguistic interaction, language technology ought to occupy a central role to play here, provided that it is prepared to deal with these aspects of language. Fortunately, the field is progressing in this direction further developing both text- and speech technologies as we will see in the following sections.

### 14.3 Text Technology Closes in on the World

Although these have not been seen as the most central concerns in linguistics, the social and interactional aspects of language have been explored by a number of linguists from various perspectives, including the ethnomethodologically oriented field of conversation analysis (e.g., Sacks 1992; Hutchby and Wooffitt 1998; Sidnell and Stivers 2013), as well as the work on phatic communication mentioned above.

On the other hand, LT has until very recently adhered to a view on language where phatic communication has been seen as peripheral at best to the concerns of

the field. This view has been shifting, where there has been a degree of coinciding between the developments in HCI and LT, which arguably reflects a more general trend in how the role of computers and computing has changed over history.

In text technology, more specifically, issues have arisen from the practical need to deal with large volumes of online text, with a clear dominance of social media content, and the concomitant observation of the obvious fact that social media represent a kind of text which does not conform to expectations on text and text structure as incorporated in existing tools for automated language processing.

This has forced the field to return to fundamental questions about language and linguistic interaction (even if often not recognized as such by many practitioners, coming as they often do from a computer-science or engineering background). At the same time, such questions can be attacked in new ways, since (1) there is now an unprecedented, truly enormous volume of empirical language data available to researchers; and (2) computational methods for data-driven processing (machine learning) have advanced at a rapid pace, assisted by the continuous exponential growth in processing and storage capacity of computing equipment that we have seen over the last half century.

Specifically, text technology researchers are now taking on the following problems related to (1) situated linguistic interaction,<sup>4</sup> which addresses how language is used to convey sentiment, emotion and attitudes, and (2) language variation and multilinguality categories,<sup>5</sup> which deal with the demography (in a wide sense) of the participants of the interaction, and in particular its linguistic expression.

Note that these are not HCI applications per se, but rather basic technologies which can be brought to bear on the task of building systems including (written or spoken) linguistic interaction informed by third-wave HCI design principles.

### **14.3.1 *Sentiment Analysis***

There has been interest in sentiment analysis and related problems among LT researchers for a long time, but the field took off in earnest around the turn of the millenium, in the wake of a new wave of data-driven NLP and access to large volumes of social-media content (Pang and Lee 2008; Tsytsarau and Palpanas 2012; Clavel and Callejas 2016). This field of inquiry goes under several names – sentiment analysis, opinion mining, and subjectivity analysis, among others – and these terms are in practice mostly used as synonymous, even if some authors try to make a distinction among them.

The interest in sentiment analysis has also been driven by increasing commercial interests. Notably, sentiment analysis is as a rule a component of so-called ‘business intelligence’ applications, information access, and processing systems deployed by commercial actors, as companies are interested not only in finding out if their

---

<sup>4</sup>Addressed in Sect. 3.1.

<sup>5</sup>Addressed in Sect. 3.2.

products and services are discussed in social and other online media, but also whether the opinions expressed there are positive or negative.

Current work on in this area focuses on four kinds of problems – viz. (1) sentiment analysis; (2) opinion mining (who expresses which opinion about what); (3) emotion analysis; and (4) argument analysis. Research in these areas is conducted using a combination of lexicon-driven and machine-learning approaches. Most work involves at least some element of prior lexical knowledge, but as online language is highly variable (see the next section), corpus-driven methods are almost always also used for dealing with both out-of-vocabulary items and context-dependent sentiment values.

Since recognizing and classifying stance and emotion are important to successful communication, these technologies are integral to 3rd wave HCI design principles in systems with elements of linguistic interaction.

### ***14.3.2 Language Variation and Multilinguality***

Given that language boundaries are arbitrary, it follows that from a linguistic point of view, language variation and multilinguality represent the same kind of phenomenon: namely that humans use different forms of language depending on who they are (or who they would like to portray themselves as being) and on the context in which they interact. Sometimes we refer to these forms as belonging to one and the same language, as dialects, sociolects, registers, etc., where at other times, we may classify them as different languages. This classification is only partly made on objective grounds, such as mutual intelligibility or (dis) similarity of language systems. In the case of related language varieties, this is often treated as a political decision.<sup>6</sup>

The notion of separate language is often tied to the existence of a written standard. Language variation in the narrower sense studied in text technology research focuses on what is sometimes called “non-canonical language” in the context of non-standard written language. Language variation has been held artificially at bay by the way written media functioned before the advent of the WWW, with “language watchdogs” – editors, proofreaders, secretaries, etc. – guarding against deviation from the written standard. In natural spoken interaction, one important source of information about our interlocutors is the language variety of their expression. We are observing more of these dynamics in online written language, where there has been intensive ongoing research on building LT tools for processing such situations (e.g., Bamman et al. 2014; Eisenstein 2015; Pavalanathan and Eisenstein 2015; Doyle et al. 2017; Pavalanathan et al. 2017; Jurgens et al. 2017).

---

<sup>6</sup>Some well-known examples are the continental Nordic (or Scandinavian) languages Danish, Norwegian and Swedish; Bosnian, Croatian and Serbian; Bulgarian and Macedonian; Hindi and Urdu; and others.

Multilinguality in a narrower sense enters the picture in two ways. Firstly, through the fact that there are many languages in the world. The widely used standard catalogue Ethnologue (Simons and Fennig 2017) lists about 6700 non-signed languages in the world. Most of these are spoken by small communities; the median number of speakers is 8000, while the mean is about 1.4 million. This indicates that the distribution of languages over speakers is quite skewed, with about 1400 languages having more than 1 million speakers. If we consider only written languages this figure drops to slightly below 400 languages.

Secondly, social media are characterized by frequent code switching, where other-language material is inserted into the text. In modern Western European contexts, this ‘other language’ is often English. In a minority-language text, the ‘other language’ may instead be insertions of the surrounding majority language (and also in English).

Dealing with the first kind of multilinguality obviously requires language technology for all the languages in question. In the great majority of cases, this actually means speech technology (see below). But even in those cases where the language is conventionally written, the present methods for building text analysis tools are almost exclusively based on machine learning and require very large amounts of text as training data – a precondition which is rarely fulfilled for all but a few languages.

For the second kind of multilinguality, text processing tools need to recognize instances of code switching and their extent in order to invoke the appropriate processing module. The goal for a text understanding system should obviously be to ‘understand’ the intercalated other-language material in the same way that would be expected of the intended readership of the text.

Lui et al. (2014) describe a system for determining which languages are present in a text and their proportions, but which does not identify individual passages as the system treats texts as bags of words, while other researchers (e.g. Solorio and Liu 2008; Yamaguchi and Tanaka-Ishii 2012; King and Abney 2013) tackle the more ambitious problem of segmenting texts into different-language passages.

Language variation and multilinguality are obviously also central elements of situated interaction, whose recognition has the potential to convey information about participants, essential to third-wave HCI design desiderata.

## 14.4 Speech Technology and 3rd Wave HCI

In addition to being different from text in that it is a more primal older form of language, several other aspects of speech differ fundamentally from text: (1) where speech is transient and exists in the moment, text is static and can be reviewed over and over; (2) where speech is highly variable and no two realizations of the same word are the same, different instances of a written word are largely the same; (3) where speech from more than one source can coexist in time, a consumer of text is

bound to one instance at a time. As the list continues, it may be argued that the differences between speech and text are greater than the similarities.

Nevertheless, speech technology has undergone a journey similar to that of text technology. The field established itself in the 1950s at which time it had a strong focus on speech as a means of transmitting information. This was evidenced by the fact that many of the early laboratories were telephony companies (e.g. Bell Labs) and the academic labs had names like “The Speech Transmission Laboratory”. In 1951, while at MIT, Gunnar Fant famously stated that “we speak to be heard in order to be understood” (in Jakobson et al. 1951). It is likely that they meant that the textual content of the speech is what was to be understood, rather than any additional phatic content in the speech.

As late as the turn of the last century, work in the subarea of automatic speech recognition (ASR) would nearly always remove any filled pauses, interjections, hesitations, and other ‘non-well-formed’ material from speech – such material was given the perfunctory label ‘garbage’. Over the last decade, change has emerged, and increasing effort is spent recognising and understanding the full content of spoken interaction (see e.g. the subarea of social signal processing).

#### ***14.4.1 Human-Computer Interaction Through Speech***

As noted, linguistics and its adjacent fields tend to treat speech as a (poorly formed) special case of text, in spite of the primary and original nature of speech and the derivative nature of text. In a similarly reductive manner, HCI traditionally treats speech interfaces as (1) task-driven, with a simple well-defined task; (2) strictly turn-based, so that user and system take turns speaking in an ‘orderly’ manner; and (3) substitutes or alternatives for some other, existing manner of performing the same task (e.g. command-line commands, DTMF signals).

A typical example of this (early) discourse is the debate on anthropomorphic versus tool-like interfaces, which was concerned with whether speech interfaces should or should not strive to mimic humans. The strongest voices argued that computers should not mimic humans (Jönsson and Dahlbäck 1988; Shneiderman and Maes 1997; see Qvarfordt 2004 for a review of the discussion in general). These discussions sought to find a general solution that would apply regardless of the purpose and context of the systems in question, and many discussions of speech interfaces, in particular in HCI, still linger in a space where content/action is disjoint from context. Today, however, the range of applications of speech in human-computer and human-robot interaction is widening so rapidly that it is becoming obvious that one solution does not fit all, and with that, we can finally put the lengthy and eloquent, but entirely moot, discussions to rest.

Before delving into particular aspects of speech interfaces with a special relation to 3rd wave HCI, let us note that speech interfaces are not spared the type of spectacular HCI bloopers that opened this chapter. In the wake of a range of successful so called ‘assistants’ from telephony giants such as Apple and Google, Samsung



recently decided to provide telephone customers with their own assistant, Bixby. Albeit on a platform with a user base that often takes pride in modifying their telephones, Samsung pushed Bixby by devoting one out of three hardware buttons entirely to the assistant, while attempting to block the option to reconfigure that button. At the same time, Bixby was endowed with the capacity to learn from its users what they most commonly wanted to know. In 2017, this combination made Bixby present the web search “how remove bixby” as a first suggestion to its users.

Next, we shall look at a few areas to exemplify how current speech technology aims at interaction that is of a less direct nature than manipulating buttons in a GUI or sending commands for a computer to execute. In these areas, we consider features of human interaction that lack obvious analogues in direct manipulation, such as collaborative emergent interaction and phatic communication.

### ***14.4.2 Physical Collaboration***

Not long ago, the interaction between robots and humans consisted almost exclusively of professionals supervising industrial robots in production lines. Robots were industrial, and they occupied spaces that either were out of bounds to humans or allowed access to trained professionals only. For all practical purposes, robots and humans dwelled in separate realities, and any human-robot encounters were surrounded by strict safety regulations (often, these regulations stated that the robot be powered down before the occurrence of any encounter). Robots that actually interacted with people in any meaningful way belonged to the domain of (science) fiction. In this domain, authors spent considerable effort pondering the nature of such interactions and how they would affect our safety. Isaac Asimov’s Three Laws of Robotics (1943) serves as a well-known example. Asimov’s laws were intended to be ubiquitous and to apply to each robot, programmatically. They were designed to eliminate the risk of robots acting in a way that is harmful to people.

To date, no ubiquitous (international) real-world robot laws exist. Still, the ways in which robots interact with people have changed dramatically. No longer are robots and humans confined to separate realities. Ranging from the relatively harmless, such as embodied virtual assistants and social robots (e.g. Nao, Paro, FurHat), through the slightly more concerning (e.g. robotic vacuum cleaners, lawn mowers, which could conceivably cause physical harm to a person), through the obviously frightening (e.g. Big Dog, Atlas), to the deliberately and blatantly harmful, such as weaponized drones, robots now coexist with people. This change has been rapid, and the window of opportunity to agree on guidelines and rules before the state of affairs gets dangerous has passed. We are left with one option only: making the best of the situation. Here, speech has a significant role to play. In order to illustrate the role of speech in current human-robot interaction, we take collaborative manufacturing as an example, as it is currently a promising field with a lively research community in which collaborative speech is particularly useful. The area targets

manufacturing where robots and humans work together physically, each contributing their strengths toward a common goal.

First, speech provides the necessary coordination in collaborative work. When a group of people lift a heavy object together, they manage their respective inputs using a blend of gestures, actions, and speech (e.g. “Heave!”). This behaviour comes naturally and is well-known to humans. Implementing such interaction in human-machine interactions (HMI) requires a level of linguistic and contextual coordination that fundamentally deviates from what is usually considered in speech HMI. This interaction is fast and situated with actions and reactions emerging simultaneously in what is similar to dancing, and quite the opposite of the slow, context independent and pre-structured ‘table tennis’ game envisioned in HCI descriptions of speech (and dialogue).

Second, physical collaboration requires trust, which provides an opportunity for phatic elements to help people and machines acknowledge that they know each other (i.e. that they have successfully attempted to collaborate and that they understand each other’s means of communication). Thirdly and finally, we tie back to safety. The immediate reaction for a person who is in pain is to signal this by calling out. If a human and a robot lift a heavy object together, and the human calls out in pain, it is crucial that the robot does not respond with “What do you mean by ‘Ow’”, “I do not know how to do ‘Ow’”, or “I believe you said ‘ow’, is that correct?”. Instead, it should cease its action instantly, then figure out whether it was correct in assuming pain and figure out how to relieve the pain. In addition to increased requirements on speed and incremental, continuous processing, this calls for an understanding of both the context and the manner in which the message is spoken – the system must be able to tell “Oh my God” as an expression of astonishment apart from “Oh my God” as an expression of pain.

In these types of applications, speech HMI is a collaborative, emergent, and trust-based activity without simple task-based goals – a specification that fits well in 3rd generation HCI.

## 14.5 Summary and Conclusions

As we have tried to show above, current developments in language technology – both text technology and speech technology – resonate well with the ideas underpinning 3rd wave HCI. This work is still at an initial stage, aiming at developing robust language analysis components capable of recognizing and analyzing phatic communication and situated interaction. In the near future we expect to see the results of this work incorporated in actual HCI applications, such as chatbots and spoken dialog systems, hopefully making them more pragmatically mature than the examples given at beginning of this chapter.

**Acknowledgements** The work on this chapter has been made possible by financial support from several sources: the Swedish Research Council through its funding of the Towards a Knowledge-based Culturomics Research Program and Swe-Clarín, the Swedish node of the European CLARIN ERIC research infrastructure, as well as the University of Gothenburg through its funding of the Språkbanken research infrastructure.

## References

- Anderson B (2006) *Imagined communities*, 2nd edn. Verso, London
- Bamman D, Eisenstein J, Schnoebelen T (2014) Gender identity and lexical variation in social media. *J Sociolinguistics* 18(2):135–160
- Bernstein B (1973) *Class, codes and control. Volume 1: Theoretical studies towards a sociology of language*. Paladin, St Albans
- Bickerton D (1990) *Language & species*. University of Chicago Press, Chicago
- Bødker S (2006) When second wave HCI meets third wave challenges. In: *Proceedings of the 4th Nordic conference on human-computer interaction: changing roles*. ACM, Oslo, pp 1–8
- Bødker S (2015) Third-wave HCI, 10 years later—participation and sharing. *Interactions* 22(5):24–31
- Clavel C, Callejas Z (2016) Sentiment analysis: from opinion mining to human-agent interaction. *IEEE Trans Affect Comput* 7(1):74–93
- Coupland J, Coupland N, Robinson JD (1992) “How are you?”: negotiating phatic communion. *Lang Soc* 21:207–230
- Darwin C (1871) *The descent of man, and selection in relation to sex*. John Murray, London
- Dourish P (2004) What we talk about when we talk about context. *Pers Ubiquit Comput* 8(1):19–30
- Doyle G, Goldberg A, Srivastava S, Frank M (2017) Alignment at work: using language to distinguish the internalization and self-regulation components of cultural fit in organizations. In: *Proceedings of ACL 2017*. ACL, Vancouver, pp 603–612
- Dunbar RIM (1996) *Grooming, gossip and the evolution of language*. Harvard University Press, Cambridge, MA
- Dunbar RIM (2004) Gossip in evolutionary perspective. *Rev Gen Psychol* 8(2):100–110
- Eisenstein J (2015) Systematic patterning in phonologically-motivated orthographic variation. *J Sociolinguistics* 19(2):161–188
- Hansegård NE (1968) *Tvåspråkighet eller halvspråkighet*. Bonniers, Stockholm
- Hansegård NE (1990) *Den norrbottensfinska språkfrågan: En återblick på halvspråkighetsdebatten*. Uppsala universitet, Centrum för multietnisk forskning, Uppsala
- Harrison S, Sengers P, Tatar D (2007) The three paradigms of HCI, presentation at the Alt. Chi session at the SIGCHI conference on human factors in computing systems San Jose, CA
- Harrison S, Sengers P, Tatar D (2011) Making epistemological trouble: third-paradigm HCI as successor science. *Interact Comput* 23:385–392
- Horton H (2016) Microsoft deletes ‘teen girl’ AI after it became a Hitler-loving sex robot within 24 hours. *The Telegraph* Published online 2016–03–24. <http://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit/>. Accessed on 10 Aug 2017
- Hutchby I, Wooffitt R (1998) *Conversation analysis*. Polity Press, Cambridge
- Jönsson A, Dahlbäck N (1988) Talking to a computer is not like talking to your best friend. *Scandinavian conference on artificial intelligence (SCAI’88)*. IOS, Tromsø
- Jakobson R (1960) Closing statement: linguistics and poetics. In: Sebeok TA (ed) *Style in language*. MIT Press, Cambridge, MA, pp 350–377
- Jakobson R, Fant G, Halle M (1951) *Preliminaries to speech analysis. The distinctive features and their correlates*, Technical report. MIT, Cambridge, MA

- Jokinen K, McTear M (2009) Spoken dialogue systems. *Synth Lect Hum Lang Technol* 2(1):1–151
- Jones KS (2007) Last words: computational linguistics: what about the linguistics? *Comput Linguist* 33(3):437
- Jones PE (2013) Bernstein's 'codes' and the linguistics of 'deficit'. *Lang Educ* 27(2):161–179
- Jurgens D, Tsvetkov Y, Jurafsky D (2017) Incorporating dialectal variability for socially equitable language identification. In: *Proceedings of ACL 2017*. ACL, Vancouver, pp 51–57
- King B, Abney S (2013) Labeling the languages of words in mixed-language documents using weakly supervised methods. In: *Proceedings of NAACL-HLT 2013*. ACL, Atlanta, pp 1110–1119
- LaFrance A (2017) An artificial intelligence developed its own non-human language. *The Atlantic* Published online 2017-06-15; <https://www.theatlantic.com/technology/archive/2017/06/artificial-intelligence-develops-its-own-non-human-language/530436/>. Accessed on 10 Aug 2017
- Laver J (1975) Communicative functions of phatic communion. In: Kendon A, Har-ris RM, Key MR (eds) *Organization of behavior in face-to-face interaction*. Mouton, The Hague, pp 215–238
- Linell P (1982) *The written language bias in linguistics*. University of Linköping, Department of Communication Studies, Linköping
- Lui M, Lau JH, Baldwin T (2014) Automatic detection and language identification of multilingual documents. *Trans Assoc Comput Linguist* 2:27–40
- Malinowski B (1923) The problem of meaning in primitive languages. In: Ogden CK, Richards IA (eds) *The meaning of meaning*. K. Paul, Trend, Trubner, London, pp 296–336
- Martin-Jones M, Romaine S (1986) Semilingualism: a half-baked theory of communicative competence. *Appl Ling* 7(1):26–38
- Miller V (2008) New media, networking and phatic culture. *Convergence* 14(4):387–400
- Monk A (2003) Common ground in electronically mediated communication: Clark's theory of language use. In: Carroll JM (ed) *HCI models, theories, and frameworks: toward a multidisciplinary science*. Morgan Kaufmann, San Fran-cisco, pp 265–289
- Pang B, Lee L (2008) Opinion mining and sentiment analysis. *Found Trends Inf Retr* 2(1–2):1–135
- Pavalanathan U, Eisenstein J (2015) Audience-modulated variation in online social media. *Am Speech* 90(2):187–213
- Pavalanathan U, Fitzpatrick J, Kiesling S, Eisenstein J (2017) A multidimensional lexicon for interpersonal stancetaking. In: *Proceedings of ACL 2017*. ACL, Vancouver, pp 884–895
- Qvarfordt P (2004) *Eyes on multimodal interaction*. PhD thesis, Linköping University, Department of Computer Science
- Reiter E (2007) Last words: the shrinking horizons of computational linguistics. *Comput Linguist* 33(2):283
- Rutten G (2016) Historicizing diagglossia. *J Socioling* 20(1):6–30
- Sacks H (1992) *Lectures on conversation*. Vols. I & II, Jefferson G (ed). Basil Blackwell, Oxford
- Senft G (2009) Phatic communion. In: Östman JO, Verschueren J (eds) *Culture and language use*. John Benjamins, Amsterdam, pp 215–238
- Shneiderman B, Maes P (1997) Direct manipulation vs. interface agents. *Interactions* 4(6):42–61
- Sidnell J, Stivers T (eds) (2013) *The handbook of conversation analysis*. Wiley Blackwell, Chichester
- Smith, D (2013) IBM's Watson gets a 'swear filter' after learning the urban dictionary. *International Business Times*. Published online 2013-01-10; <http://www.ibtimes.com/ibms-watson-gets-swear-filter-after-learning-urban-dictionary-1007734>. Accessed on 10 Aug 2017
- Solorio T, Liu Y (2008) Learning to predict code-switching points. In: *Proceedings of EMNLP 2008*. ACL, Honolulu, pp 973–981
- Spolsky B (2004) *Language policy*. Cambridge University Press, Cambridge. .
- Stam JH (1976) *Inquiries into the origin of languages*. Harper & Row, New York
- Stenström AB, Jörgensen AM (2008) A matter of politeness? A contrastive study of phatic talk in teenage conversation. *Pragmatics* 18(4):635–657
- Tsytssarau M, Palpanas T (2012) Survey on mining subjective data on the web. *Data Min Knowl Disc* 24:478–514

- Vetere F, Howard S, Gibbs MR (2005) Phatic technologies: sustaining sociability through ubiquitous computing. In: CHI'05 workshop – first international workshop on social implications of ubiquitous computing at the 23th international conference on human factors in computing systems. ACM, Portland
- Wang V, Tucker JV, Rihll TE (2011) On phatic technologies for creating and main-taining human relationships. *Technol Soc* 33:44–51
- Wikipedia (2017) Tay (bot). [https://en.wikipedia.org/wiki/Tay\\_\(bot\)](https://en.wikipedia.org/wiki/Tay_(bot)). Accessed on 10 Aug 2017
- Wintner S (2009) Last words: what science underlies natural language engineering? *Comput Linguist* 35(4), doi:<https://doi.org/10.1162/coli.2009.35.4.35409>, <http://aclanthology.coli.uni-saarland.de/pdf/J/J09/J09-4012.pdf>, 641
- Yamaguchi H, Tanaka-Ishii K (2012) Text segmentation by language using minimum description length. In: *Proceedings of ACL 2012*. ACL, Jeju Island, pp 969–978
- Žegarac V, Clark B (1999) Phatic interpretations and phatic communication. *J Linguist* 35(2):321–346
- Simons GF, Fennig CD (2017) *Ethnologue: languages of the world*, 20th edn. SIL International, Dallas. Online version: <http://www.ethnologue.com>

# Chapter 15

## Sensorial Computing



Puja Varsani, Ralph Moseley, Simon Jones, Carl James-Reynolds,  
Eris Chinellato, and Juan Carlos Augusto

**Abstract** Humanity is at a special time in its relationship with technology where there is an increasing likelihood of artificially replicating characteristics which we thought were in the realm of the distinctly human. Artificial Intelligence and Robotics are making the news increasingly often and replication of body parts is also making progress. This chapter looks at senses which although not exclusively human have a powerful potential to support other higher functions and aspects of human life. Technology has been developing nature-inspired artifacts which resemble somehow their human counterparts with specific practical applications. So far these explorations have been mostly isolated. It is only a matter of time, however, until these become physically and logically connected into a cooperative fashion – in fact Robots typically use them although not always in their full capacity. Such developments will provide machines with interface capacities of a higher order, bringing new powerful tools to solve new problems, whilst raising unexpected scenarios and challenges for our societies. Regardless of whether we want it or not, it seems impossible to stop technological progress in this direction. We assume this development is here to stay. Thus we look at the artificial and human synergies, how interaction with machines is influenced by sense-like interface capabilities – namely “sensorial computing”.

### 15.1 Background

We are witnessing a historical transition where technology has become an integral part of almost every aspect of human life experience. Until recently, this dynamic was forecasted and explored with the advent of what was then termed Ubiquitous Computing (Weiser 1999). Sensing technology is now helping drive the transport we use, gives us feedback on our physical performance, and directs us toward our

---

P. Varsani (✉) · R. Moseley · S. Jones · C. James-Reynolds · E. Chinellato · J. C. Augusto  
Middlesex University London, London, UK  
e-mail: [p.varsani@mdx.ac.uk](mailto:p.varsani@mdx.ac.uk)

destinations. These technologies can be assembled to form technological bubbles which in turn can create sophisticated services in a variety of daily life situations (Augusto et al. 2013).

However, there is still a divide between humans and technology where that symbiosis is definitely embryonic. We experience numerous difficulties on a daily basis to access the services we need, and although progress has been made (Aghajan et al. 2009), there is certainly still much scope for progress. Moving forward from the current bottleneck of seamless integration of technology and humans, this chapter looks one step further into how artificial systems can complement human perception. There is clearly a gap when it comes to comparing how humans perceive their surroundings and how current systems do. This chapter considers what the strengths of humans and systems are, and how future developments in this area of technology can benefit human-system interactions.

Harrison et al. (2007) define the 3rd wave of HCI as a “phenomenologically situated paradigm” focusing on meaning and meaning creation, human experience represented through multiple perspectives, and their interrelationships. Whilst we are still understanding and exploring this paradigm, interaction design is also developing new strands of enquiry as we equip devices with a range of sensors and tools to mimic and extend human perception. Our understanding of these interactions is very limited as we move from using etic to emic markers. Ryan and Cacioppo (1981) explain that from an audio perspective, speech markers derived from non-content speech cues with analysable characteristics are etic. Emic markers, in contrast, include “social meaningfulness in terms of interpersonal judgements.”

What makes this area particularly exciting is that the interpretation of multimodal stimuli is not yet fully understood in humans, but also that we have yet to fully apprehend optimal ways of processing this electronically in order to develop new applications that support us in our endeavours. A computer analysing raw data and making judgements of gender or emotions, for example, is using etic markers. As we move to developing contextual awareness and multimodal input, computers will be able to utilise emic markers.

## 15.2 Human Sensing

### 15.2.1 Hearing

The human auditory system provides information about sounds (air vibration) and their location. This is useful for communication and safety, for example, as we can hear a car moving towards us. Human hearing is affected by age. We become less sensitive to higher frequencies as we get older, although hearing loss is influenced by a wide range of factors. We rely heavily on sound for communication and there are a variety of phenomena such as the cocktail party effect that allow us to communicate effectively despite background noise. Two sounds presented with a very

short delay are heard as one sound and this is known as the precedence effect. It assists us with localisation of sound (Cremer 1948). Our ears act as multiple band filters and it is the harmonic content of sound and its change over time that allows us to identify a range of different instruments and voices (Bregman 1990). We are also capable of detecting emotions in a voice and this is an important part of communication. Humans have created sounds for entertainment and we use a wide range of rhythms and scale types. Scales and tempo may be linked to emotion, however there are cultural differences in how these might be interpreted (Balkwill and Thompson 1999). Our ability to produce tonal variation using synthesis allows us to create new sounds and to resynthesise existing sound.

There are a number of interesting psychoacoustic phenomena that have been exploited in our processing of sound. An example is that of frequency masking where the presence of a masking frequency can prevent a quieter frequency in the same filter band from being heard. Temporal masking is a phenomenon where the masking effect is caused by the onset of a sudden masking sound where audio preceding and following the masking sound is not registered by the listener. This phenomenon is exploited in lossy compression schemes such as the MPEG standard (ISO/IEC JTC1/SC29/WG11 1999). The human auditory system does not support equal sensitivity to different frequencies at low volume. This was originally described as the Fletcher Munson Curve (Fletcher and Munson 1933), but more recent measurements have given us Equal Loudness Contours (ISO226:2003E). Compensation for this has led to the development of Loudness controls on amplifiers that allow the bass to be boosted when audio is listened to at low amplitude.

### 15.2.2 *Vision*

Among all human senses, vision is by far the one which involves the most neural networks in the brain. This is certainly due to its importance and complexity. Indeed, progress in animal and human vision research has shown that what we perceive as an image is the product of very complex, distributed, and interdependent cognitive processes involving most of the brain. The anatomy and functionality of human visual areas is rather well known, but how all these areas interact among themselves and with others, according to internal and external factors is a very large and active field of research in cognitive psychology and neurosciences (Tootell et al. 2003). Here, we only summarise a few fundamental concepts.

First of all, vision is not a deterministic process. This means that the same scene can elicit substantially different responses depending on the subject internal state and external context. Some of the factors affecting visual perception are environmental context, task, concurrent stimuli, social conditionings, attention allocation, personal experience and knowledge, motor goals, and others.

Second, vision is not a unidirectional process, nor a linear one. Indeed, while the main direct visual pathway is slightly predominant, visual perception derives from the integration of the neural activation of parallel pathways, and is strongly biased



by top-down projections coming from other brain regions – for example, goal-related areas in the prefrontal cortex (Chinellato and del Pobil 2009).

As all other senses, vision is not immune to cross-modal effects. On the contrary, concurrent stimuli from other senses, and/or priming and inhibitory effects are directly affecting visual perception by changing the focus of attention or biasing perceptual goals, or inhibiting/enhancing responses according to the congruence/incongruence of concurrent stimuli.

For what concerns social interaction, vision is certainly a major factor in establishing and successfully carrying out most human interactions, either personal, professional, or casual. At the same time, the social context is one of the main factors in visual perception, strongly affecting attention allocation. A large number of studies have shown how a variety of visual stimuli related to gestures, expressions, postures, and all types of movements, even minimal ones, are taken into account during social interactions. Indeed, subjects perceive the interaction as somehow incomplete if any of those stimuli is missing.

### 15.2.3 *Smell*

Our olfactory sense has only really been integrated with electronic devices and media in the Geist of the 3rd wave. Suzuki et al. (2014) provide a summary from early experiments in the 1950s to more recent use of sniff cards. The use of smell has largely been associated with limited attempts to enrich perceptual experiences. However, we have been equipping devices with a range of chemical sensors to facilitate a wide range of tasks including the tuning of cars (exhaust analysers), breath tests and smoke /gas alarms.

The human olfactory sense of smell is closely linked with that of taste. Much of what we consider to be flavour comes from our olfactory sense and like other senses it changes with age and external factors such as smoking (Doty et al. 1984). Smell serves many functions and can act as a warning of danger. For example, we can detect the difference between a house fire and overbrowned toast. Smell is also of use in less critical situations as it allows us to detect if milk has ‘gone off’. However there are counter examples, such as that of Durian – a fruit that tastes delicious despite smelling of rotting flesh and failure to detect some toxins such as carbon monoxide. The human nose is an extremely sensitive organ (McGann 2017). It contains approximately 50 million cells acting as primary receptors. Despite noses being far more sensitive than our eyes in terms of sensors, we have a very limited shared vocabulary for articulating odours or classifying them (Groen et al. 2017). Our noses interpret smells differently from each nostril (Herz et al. 1999). Odours are linked to memory, other stimuli, and context that give rise to significant changes to autonomic nervous system responses, even when we sleep (He et al. 2014; Herz et al. 2004; Braun and Cheok 2014).

### **15.2.4 Taste**

The tongue is the primary organ of taste in the gustatory system and it is closely linked to the olfactory system. Its upper surface is covered in taste buds equipped with taste receptor cells (Li et al. 2002). Each taste receptor receives multiple chemical substances creating a single taste. There are five basic tastes: sweetness, sourness, saltiness, umami and bitterness. The latter is generally found unpleasant while the others provide a pleasurable sensation. Discovery of the bitterness receptors in taste cells was made in 2000, shortly followed by sweetness and umami receptors (Adler et al. 2000). Further clarification is expected in the future on mechanisms behind sourness and saltiness receptors and their physiological significance.

### **15.2.5 Touch**

Skin is the largest, oldest and most sensitive organ in the human body and is said to be the first medium of communication (Juhani 2014). Covered in 1000's of nerve endings, dermis is a layer of the skin providing the sense of touch. The brain processes touch via two parallel pathways: (1) the first provides facts such as vibration, pressure, texture, and (2) the second social and emotional information. Movements of the hands and fingers allow humans to strike a surface to detect texture, palpate gently or trace edges to judge shape, press to determine hardness, etc. (Prescott et al. 2011). Touch plays an important role in child development, social relationships and can feel different based on the social context of the encounter (Linden 2015).

### **15.2.6 Multisensory Processing**

The availability of multiple sensory modalities and their continuous integration are almost universal features in the animal kingdom. Almost all animals have multiple ways of perceiving. Nevertheless, compared to the historical development of detailed research dedicated to sensory modalities, the mechanisms behind concurrent processing of stimuli coming from different senses are relatively unexplored. In fact, only recently multimodal (or multisensory) integration has been obtaining increasing interest in cognitive and brain science research.

Even though some of the basic defining principles behind multisensory processing are well known, more studies are shedding new light on the interaction mechanisms among various modalities. At the same time, those studies are aiming to clarify the relevance of multimodal integration for sensorimotor interactions and behaviour in general. For example, it is somehow obvious that at least vision and hearing (and touch in many cases) are involved in everyday human social interactions,

but the strict relation among these modalities and the way they modify each other's perception and the whole integrated percept are still matters of study. What is clearly being established is that sensory modalities can hardly be considered independently, since they affect each other not only for the whole integrated feeling, but down to the perception of each single modality alone (i.e., what we actually see depends also on what we are hearing, and the other way round).

## 15.3 Artificial Sensing

### 15.3.1 *Vision*

Somehow mimicking the importance of natural vision, artificial vision (also called computer vision) is a prominent field in modern day computer science and artificial intelligence research. The focus of computer vision has gradually shifted from image acquisition and processing to image understanding and the representation of knowledge deriving from visual inputs. The subfields of computer vision most related to human machine interactions are object and pose recognition, tracking, and face and expression recognition.

Currently, computer vision systems are able to reliably identify objects and faces they have been trained to recognise, but their generalisation capabilities and unsupervised learning skills are still limited, compared to natural systems. Human pose recognition – a fundamental problem in social interactions – has recently seen significant advances thanks to the use of depth sensors. Such sensors do not perform proper visual processing, being more akin to sonar data processing, but they are usually paired with standard cameras to produce richer representations of scenes in 3D, particularly when human subjects are involved.

The ability of artificial visual systems to take into account contextual information is still quite limited, and mostly related to general advancements in artificial intelligence and knowledge engineering. In fact, despite substantial progress in recognition accuracy often obtained with novel techniques based on the deep learning, convolutional neural network framework, artificial visual processing is still a mostly bottom-up, unidirectional process, and its integration with previous knowledge and memory occurs only at the later processing stages for high level representations. It has been mentioned above that this is not the case for natural vision, which is a tightly interconnected multidirectional process.

### 15.3.2 *Taste*

An electronic tongue or taste sensors have been present since 1990 in the food industry as a method of evaluating taste. Some sensors discriminate solution samples and others are used to quantify the intensity of each type of taste using a taste 'scale' (Kobayashi et al. 2010; Toko 2000). Electronic tongue concepts are based on combining the technologies for measuring potentiometry, voltammetry, amperometry, electrochemical impedance spectroscopy and hybrid approaches (potentiometry, voltammetry and conductivity (Kumar et al. 2012). Metallic electrodes are used to measure electrodes in voltammetric measurements to gather different potential responses. Taste sensors are composed of lipid/polymer membranes for transforming information of taste substances into electric signals (Kumar et al. 2012). The output shows different patterns for chemical substances with different taste qualities. Some taste interactions which occur, such as bitterness and sweetness, can be detected and quantified using taste sensors.

### 15.3.3 *Hearing*

Sounds have been incorporated into a wide range of devices and in the first wave were used as notification of events (Brixen 2007). Sonification is used in devices such as ECG and Geiger counters which have audio interfaces. In the second wave, sounds were chosen to represent task completion and user feedback, for example the click of a button or the sound of a bin being emptied. Sound was also used with speech synthesis and the ability to read out text. This has provided support for those with impaired vision. The use of sound by computers has also been developing with speech recognition products originally targeting administrative productivity. More recently, mobile devices have provided much of the impetus for speech recognition where output with Siri is probably the most recognised. More recently, Alexa and similar devices have brought these technologies into the home, where they integrate into the home network (Moskvitch 2017).

Microphones can have very wide frequency response and in arrays can capture the location of a sound accurately. The frequency response is limited by Nyquist theorem (1928) and the upper limit is determined as half the sampling frequency, the increasing processing speed of computers has been reflected in the increased available sampling rate. We can put our sound through Fourier transforms (Fourier 1822) and extract detailed frequency information about tone. This allows us to identify individuals from their voices. Levels of stress and emotion can be detected from the voice with a limited degree of accuracy (Scherer 2003). Audio analysis also allows us to forensically analyse signals for powerline noise and other subtle audio components that the human ear is not sensitive enough to identify (Kajstura et al. 2005). Likewise, we can also make use of ultrasonic noise that might be used for echolocation.

### 15.3.4 *Touch*

Tactile sensing comes with challenges in the simulation of skin characteristics with high resolution, sensitivity and rapid response (Hou et al. 2014). Electronic skin should also be able to sense temperature and humidity simultaneously and differentiate diverse mechanical forces. It should also be able to recognise both medium pressure (for object manipulation) and low pressure (gentle touch) (Kim et al. 2011). This can be achieved with flexible and stretchable pressure sensor arrays which are constructed based on different transduction mechanisms and structural design to imitate tactile sensing (Wang et al. 2015). Transduction methods such as piezoresistivity, capacitance, and piezoelectricity convert external stimuli into an analogue electronic signal. To meet new challenges and opportunities other transduction methods such as optics, wireless antennas, and triboelectricity are undergoing rapid development (Chen et al. 2014). Development in the flexibility and stretchable pressure sensors is significant for skin to maintain pressure-sensing ability under complex mechanical deformation. Other recent developments in this field include the production of self-powered e-skins which combine wireless technology and energy/data transfer and the integration with high-density flexible circuits (Wang et al. 2015). One of the limitations is the development of highly intelligent e-skins that can sense and respond to the changes in the external environment.

### 15.3.5 *Smell*

Olfactory Interfaces, often described as “Electronic Noses” are often created from chemical sensor arrays in conjunction with pattern recognition systems. There is a large range of sensor types with new developments in nano-scale sensors and in Terahertz spectroscopy. Sensors can detect a wide range of particles in the air, but sensors are often limited to specific chemicals and may need higher concentrations of an odour in order to detect it. Uses of this technology range from the simple detection of ethanol in breath to pollution sensing and as a means of detecting hazardous materials such as explosives. Sensors are also developed for detecting certain diseases and to ensure that food is fresh. Some sensors are limited by the need to recalibrate and some types are “poisoned over time” (Villarreal and Gordillo 2016). In some aspects, this is similar to human noses where a strong odour may overstimulate the sensors, which then require time to recover.

**Table 15.1** Comparing media

	Hearing	Vision	Smell	Taste	Touch
Natural	Ears	Eyes	Nose	Tongue	Skin
Artificial	Microphones	Cameras	E-nose	E-tongue	Tactile sensor

**Table 15.2** Comparing current perceived levels of achievement

	Natural	Artificial
Hearing	High	High
Vision	High	High
Smell	High (domestic use)	Medium
Taste	High	Medium
Touch	High	Medium

### 15.3.6 *Multisensory Processing*

In robotics, multisensory processing is usually employed to complement defective or noisy sensory modalities, e.g. by substituting the information from sensors when they become missing or unreliable (e.g. to compensate for occlusions). In other instances, multisensory processing consists of switching among modalities (e.g. from visual planning to tactile feedback when interacting with objects). Indeed, there is a lack of research on the general rules about how various modalities can be integrated, and the possible advantages of such integration on the interaction of robots with objects and people.

For what concerns human robot interactions, multisensory processing is mostly referred to the integration of voice recognition with vision, and sometimes touch. Again, this is done almost exclusively in order to solve specific issues in specific scenarios where different sensory modalities are employed with different goals and tasks, not in a principled integrated way as we know is the case for natural system.

Table 15.1 relates the natural senses and some of their artificial manifestations. Table 15.2 confronts them in terms of perceived qualitative and quantitative all round effectiveness. Even though the comparison is arbitrary and dependent on the specific criteria of interests, general development trends can be drawn. It is fairly uncontroversial that artificial smell and taste sense are the least developed. Touch is also underdeveloped compared to hearing and vision, but is improving quickly. More development is needed in the synergies of these all. In fact, the principles behind multimodal integration in nature are still not fully understood, and the integration of different sensory modalities in artificial systems is typically done ad-hoc on specific applications (e.g. visuo-tactile systems for grasping). A thorough study with computational modelling of the general principles behind multimodal integration is required before this can fully develop artificially.

## 15.4 Potential Application Domains

As the threshold of current applications is pushed back, HCI is not fully realising the many concerns of what is considered ‘3rd wave’; emotion, embodiment, experience. There is the idea within this wave of the context of the system and its relation to the environment or knowledge base in which it resides. It is not just what is created and constitutes a system but also what is not created – the Zen of systems methodology and implementation.

There is a departure from fixating upon the interface, as seemed to be the case in previous paradigms, and a move toward experience and context. The shift is primarily to one of relationships, the validity of the data, and the exploration of a subject or question.

We can break the application domains down according to a framework (Downey 2015) based on varying perspectives on the self; that is, at the Hedonic, Eudaimonic and social/interpersonal levels. These levels correspond to the use of technologies to induce positive subjective experiences (hedonistic), support a sense of growth in the individual (Eudaimonic) or using technology to improve social integration, connectedness or holism. It is possible that some technologies cross boundaries between these viewpoints: for example, Individual or Subjective Well-Being (Hedonistic) or Psychological Well-Being (Eudaimonic).

### 15.4.1 *Social Well-Being*

The web of ubiquitous interconnections that constitutes the internet and which embrace individual connections at many levels allows for the sharing of data for spatial and temporal merger. For example, data can be collected about an individual’s metrics (which could include basic bio information) at any given point in space and time along direct information about that person’s context. When this information is enabled into Virtual Reality or mixed reality systems, it becomes possible to share experiences at varying levels of degree (one-to-one, one-to-many, many-to-many cardinality).

An individual can experience an event taking place remotely. A concert where the agent or managed perspective of experience is handled by a production team or simply by the first-hand experiencer in the audience can serve as an example. The managed perspective of experience may not be limited to a human agent’s viewpoint; this could include input from more abstract thought – art, music or other synthesised realities. This managed space could be directed in a live sense or programmatically. That is, events can be triggered interactively over time or in a sequenced fashion; or even a mixture of various forms of input can trigger and cascade scenes.

### 15.4.2 *Surrogates and Multi-user Avatars*

An interesting aspect of shared realities is the idea of group access to a single avatar. This could be within virtual worlds, such as Second Life. It may be as simple as time sharing of an avatar to experience a reality through that particular guise, identity or persona, or taken to a point where the agent's senses are plugged into the many, with either one or many individual drivers making decisions for the avatar. This sharing of one avatar has already been seen in SL by the Naughty Auties<sup>1</sup> (A virtual resource centre within the virtual world) to practice social interaction.

There is also the idea of the surrogate – the avatar itself becomes instantiated in physical form (as perhaps a robot) and its inputs and outputs tied into a person(s) Virtual Reality system. Here, the individual becomes the driver of the surrogate. The benefits here are that there is no danger to the human operator: the robot can be used in hazardous environments, or situations problematic for a human. Taken to its logical extreme, the movie *Surrogates* shows how humans of the future interact via their robotic counterparts mainly, and have left their human bodies to atrophy. Unfortunately, in the movie a mechanism has been found which enables fail safes to be overcome and direct fatal injury to be inflicted on the owner via the robot. An interesting aspect of the movie is the social backlash by some to the surrogate paradigm which leads to the creation of areas free from such entities. Any technology which supplants what it means to be human is bound to create some disruption at a social level based on philosophical or religious concerns.

### 15.4.3 *Loneliness, the Human Condition*

As AI progresses and gets closer to the goal of mimicking human intelligence, there is the possibility of creating – what appears to the user, a personality to interact with (as in “Bicentennial Man”<sup>2</sup>). This progress is exemplified by products such as Siri, Google and Amazon's Alexa. Many tasks can be automated through the use of such ‘intelligences’ and there is some ability to converse – ask what the weather is like, setting of reminders and suchlike. This is useful not only on mobile operating systems but on mainstream operating systems too. The capacity for a meaningful conversation has not been fully developed yet as most users report that ‘there is something lacking’.

The movie “Her”<sup>3</sup> illustrates how these dynamics can progress. For example, setting up the operating system (OS) with a few basic questions about the user can assist the system in deducing the rest from further conversations. In this movie, there is emotional content to the human – computer dialogue, so much so the human

---

<sup>1</sup> <http://edition.cnn.com/2008/HEALTH/conditions/03/28/sl.autism.irpt/index.html>

<sup>2</sup> <http://www.imdb.com/title/tt0182789/>

<sup>3</sup> <http://www.imdb.com/title/tt1798709/>



falls in love with the OS... only to find himself eventually rejected – after all, all those apparently singular OS's were in fact working together and learning from their human masters as they happily chatted.

It may be that humans could quite easily be tricked into believing and emotionally investing in what seems to feel to have human qualities but lacks the qualities of the deeper aspects of human-ness. There have been cases – for example in Japan, where males have been attaching themselves not to real girlfriends but instead, simulacra in the form of virtual girls accessible on mobile devices and computers. After all, the virtual girl can be made in anyway the user requires. This has become such an important factor that it is actually affecting the birth rate of the country.<sup>4</sup> These 'otaku' as they are known, prefer the company of their related anime and manga. With some consumers of these technologies withdrawing from normal relationships, sex and marriage (Sheutz and Arnold 2016).

Some of the extremes of the human condition could in theory be remedied to a degree by interaction with artificial intelligences operating within the remit of 3rd wave HCI. One of these primary aspects has already been tackled in the many forms it takes: loneliness. Loneliness occurs when an individual feels isolated and their social interactions do not match their expectations. Loneliness occurs predominantly in the following groups of individuals:

- Older age groups at home, care facilities, hospices and hospitals.
- Children with illnesses that isolate them from their peers.
- People with particular diseases and conditions.

A good example of new technology enlisting new paradigms of interaction is the telepresence robot AV1.<sup>5</sup> AV1 acts as the avatar of the ill child, attending classes as their communicating presence. The robot sits at a vacant desk and acts as the eyes and ears of the child - even turning its head for a full view of the classroom. Should the child want to attract the attention of the teacher, the robot will flash blue. Interestingly, if the child wants to communicate with a nearby friend they can do so by an available whisper mode, keeping their comments from being overheard by the teacher. There is evidence of the AV1 providing an emotional as well as logistic link for the user. Children have personalised the robot – giving it gender or customisations. It has been envisaged that schools would purchase these units, or others like them, for their shared use by pupils in need. AV1 is being used in other contexts too – including to stop the isolation of older people.

---

<sup>4</sup> <http://www.bbc.com/news/magazine-24614830>

<sup>5</sup> <https://www.theguardian.com/technology/2017/aug/13/robot-to-help-sick-children-norwegian-start-up>

### ***15.4.4 User Emotional and Creative Investment in Virtual World Spaces***

A large amount of time and creative effort has been previously invested into virtual spaces such as Second Life. The creative endeavours were extremely large in scope at times. For example, space could be controlled to provide the exact environment one required varying from a spacey abstract scape involving seeing an alien landscape with galaxies floating in the background, to stepping stones reaching across open voids and stars colliding in the skies. Yet another zone was the scene of a Japanese area complete with pagoda and cherry blossom trees.

The landscapes in these simulations are constructed either in world or via art packages, the products of which are imported in. There is the structure for the objects, the texture or surface content, and, if a complex object, they may also contain a program or script. This script allows the performance of tasks and interactions with a user. For example, if the object is a simulation of a television, the script will give it the functionality of a TV – streaming an image to its external texture and allowing control by the user.

Being allowed to control every factor in a simulated area allows for enormous creativity. As this technology becomes even more immersive, the impact on humans will be dramatic.

### ***15.4.5 Biofeedback***

It is possible to provide a biofeedback loop to control the environment that a person is immersed in. This could be done to train a person in certain situations or to provide an environment suited to targeting specific body or mental states.

The basic means of achieving this is by introducing an EEG reader (and other body metrics if required such as heartbeat or skin resistance) to the usual VR headset. Software can then follow the EEG pattern for specific signatures in reaction to the environment. Using this reaction to an immersed scene, further alterations can be introduced and again the signature monitored. In this way, a target state is induced and a biofeedback loop created.

An example of this are meditation training systems, some of which are based on Buddhist or yogic techniques which have existed for hundreds of years. Meditation states have specific brainwave signatures which are known. These become the target state to be induced with a specific route of immersion to get there. This may be the concentration on particular objects (known as Kasina). The objects and a surrounding environment are constructed in VR where software checks the focus of the individual and attempts to hold their concentration by manipulating the scene. The response in the EEG is watched and further alterations occur.

**Table 15.3** Extended senses

Sight	Beyond normal ranges of human bringing in infrared, ultraviolet; the microscopic or macroscopic
Hearing	Hearing what cannot be heard – Ultrasonic, the extremely quiet, listening to data itself
Taste	Tasting an image, audio...artificial tastes
Smell	Smelling fragrances produced remotely
Touch	Remote touch – Touching through robots and devices at a distance

A simple version of this may be the raising of a ball which flies higher, the closer the state is to being achieved – thus providing a reward to the individual and feedback that they are inducing the desired state of mind (Moseley 2016).

It is not only training that the introduction of an EEG or other body metrics could be used for. Movies can be constructed to change depending on how the watcher is feeling and therefore relating to the material. Within VR, both games and interactive stories can be made in this way which are malleable or self-creating, reacting perhaps to how the participant feels.

#### 15.4.6 *Human/Machine Merger – Boundaries*

Humans can be enhanced, or have deficiencies reduced, by the integration of machine parts into the body. These enhancements can be mechanical, communicative or information-based. Where the human body has failed or never assumed the generally accepted norm in the first place, mechanical parts – intelligent or otherwise – can be cleverly embedded. This is the generally accepted form of the cyborg or bionic man.

In recent years, and with the exploration of wet wiring and neural interfaces, there has been a development toward the direct connection between the brain and computer. In its most basic form, we see how this could unfold by looking at augmented reality – the meshing of the human, in their environment and the informational overlay that can be presented with spatial and temporal data, as well as network integration. What if, the technology could be directly and discreetly projected at some point into a brain's optical systems? Given the technologies with biofeedback discussed here, it could be seen how eventually humans could well merge at varying degrees with machine networks and intelligences. It can be seen how this 'connected self' easily fits into the 3rd wave model, in the sharing or social mode.

The current state of mobile technologies gives clues to what will be possible when a human integrates the technology further. This trans-human, or post-human state also gives rise to other possibilities – extended knowing through access to high speed networks and databases and also possibilities that give rise to almost super-human senses – extending ranges, see Table 15.3.

**Table 15.4** Extended knowing (Telegnosis)

Maps	Spatial awareness (Where am I in relation to ...?)
Proximity	Where is ...? (a person, an object ...)
Ideas, concepts, messaging	Telepathy (what do you/they think?)
When	Temporal awareness (when does this happen?)
Prescience	The fact of knowing something in advance (how probable is this?)

In short, the sense ranges can be extended beyond their normal human capacities but also, there is the ability to swap information in a synesthetic sense, tasting what is normally seen or hearing what can usually be only seen. Data, subjected to specific algorithms, could be represented as streams of audio, light or even taste. This can be known as sensory substitution (O’Connell 2017), see Table 15.4.

A human therefore, at this level, becomes enmeshed deeper in the world around themselves in both a social and communicative sense.

The issue may no longer be the lack of information, but being able to survive in a veritable sea of information and being able to filter what is contextually appropriate for an individual. This embedded node, the person, will require a filtering system to survive, not unlike how modern malware and antivirus systems work.

### 15.4.7 Identity in the Trans-human Self

There are scholars who proclaim that with the deep embedding of the individual within the social and technological matrix there is a loss of self (Moseley 2016). This is almost the reverse of what we see in the HCI 3rd wave priorities at psychological and social levels.

Instead of self-realisation (the Eudaimonic level), it is argued that the individual can be battered by an onslaught of social media related conformity, or become lost in the “lonely crowd” (Turkle 2011). Alternatively, there is the hope that those very same algorithms that threaten individuality become true personal assistants and saviours of mental well-being, acting as filters and guides through the smog of a polluted data cloud.

Is it possible that we will need AI simply to act as our proxy – to keep up with the social interactions going on around us? A trained AI with our persona at heart, forever watching the statuses of others, aligning our interests and only passing back what is good for our sanity? A social “trading floor” where the currency is “likes” and “dislikes” and “friends” and “unfriends” are passed in microseconds on huge scales.

What about AI run amok? (Lantham 2017) If the above holds true, then what happens when there is no “real” person to speak to behind the proxy and instead there is simply an intelligence or army of cloned intelligences with the interests of a company, organisation or nation state behind it?

## 15.5 Ethics

The developments in 3rd wave HCI outlined above have profound social and ethical implications. The field of ICT ethics reminds us that emerging technologies are very rarely ethically neutral (Jones 2016). The way such technologies are designed, implemented and used in particular domains raises ethical issues. These issues tend to arise when core ethical principles are at stake, principles such as privacy, dignity, autonomy or equality of access. The technologies involved in 3rd wave HCI are no exception in this regard, particularly those that mimic, augment or replace human sensorial perceptions, and especially where these technologies are ubiquitous and embedded in everyday life, or where they interface directly with the human body.

Such technologies undoubtedly introduce many benefits. They have the potential to increase human well-being, security and quality of life in a range of application domains. Electronic sensors have numerous potential environmental and safety benefits. Olfactory sensors, and object and facial recognition, for example, can be used to protect or warn humans of hazardous substances, toxins or other threats in contexts as diverse as air quality monitoring, radiation detection, food safety, counter-terrorism, border control and forensics. The benefits of virtual and mixed reality systems in training and entertainment are well known. These systems also offer the potential to negotiate hazardous environments remotely in ways that minimise risks to humans. Such potential will be further enhanced as these virtual reality systems become increasingly connected to human senses, through surrogate avatars or physical robots. The shift in modes of input away from traditional interfaces which require particular sensory abilities and motor skills also have the potential to empower individuals with impairments and disabilities by supporting a broader range of senses, including auditory, speech or touch. Novel forms of interaction can be more accessible for frail or elderly users of technology, or those with reduced cognitive or mental abilities. Sensor technologies which extract data from the human body, such as EEG readers or heart-rate monitors, can benefit individuals with complex health and social care needs. These can potentially enhance bodily integrity and well-being, and postpone institutionalisation, in the case of the elderly, for example. However, if human impairments can be reduced by these technologies, and human sensory capabilities enhanced, the question immediately arises, who will have access to such technologies and their benefits? Will they be affordable and equally accessible to all, irrespective of age, gender, disability or social class? If certain social groups are excluded from these benefits, there is the real prospect that such technologies might serve to exacerbate existing inequalities.

If 3rd wave HCI represents a new paradigm of interaction between humans and machines, it also involves a step-change in the collection and processing of personal data. This paradigm represents both a quantitative shift in the amount of personalised data that is collected from multiple sources and sensor inputs, and a qualitative shift in the level of detail of this data. Data is extracted directly from the human body. Its physiological and psychological characteristics are monitored through devices that interface with the body or are embedded within it. The fact that this data

can include sensitive, lifestyle, biometric and medical data presents new challenges for data protection frameworks and privacy rights. Questions arise about how and where this data is stored, and crucially, who or what this data is shared with or disclosed to. This could include a range of stakeholders, third parties and services, both commercial and public. The prospect of data being exchanged between these various entities, without the consent or knowledge of individuals, has concerned privacy advocates (Jones 2016). Where more personal data is extracted than people realise, expectations about privacy boundaries can be destabilised and undermined. These concerns are heightened where data is aggregated and mined for intelligence about individual and collective human behaviour. The commercial use of such data raises the prospect of scenarios such as the targeting of users with behavioural advertising. Here, various sensor technologies are deployed to monitor levels of stress or emotional response, as detected in an individual's voice, posture, heart rate or facial expression. This data is then used to target consumers at moments of maximum receptivity or vulnerability. The danger in such scenarios is that people become increasingly transparent to ever more intrusive, ubiquitous and fine-grained modes of surveillance which became ever more difficult to escape. In the hands of State agencies, such as law enforcement and intelligence, or when allied to military agendas, such tools can easily be abused in ways that violate bodily and psychological integrity, and threaten human rights.

These technologies not only provide new ways of knowing about, and reporting on, past and present human activities. They also offer ways of predicting future human behaviour, based on profiling and risk assessment. The use of such profiling to variously predict crime, social unrest or acts of terrorism can be ethically contentious. This is particularly the case when such predictive profiling is used to make decisions which might be based on inaccurate information, or which results in discriminatory treatment, depriving people of rights and opportunities.

The disappearance of the traditional interface in 3rd wave HCI, and the shift to embodied modes of processing and interaction raises issues around transparency and human control. This is especially the case where this processing occurs surreptitiously in ways that are less visible to human inspection. These issues are made more pressing by "black boxed" technologies which are specifically designed to operate invisibly or automatically without human intervention (Jones et al. 2015). In such scenarios, users may, or may not be aware that they are interacting with technologies which are embedded in their immediate environment. Where technologies start to disappear from human consciousness in this way, and where these operations are opaque, they become more vulnerable to abuse, and serious questions arise about the extent of human agency and control. It is the prospect of humans merging, to varying degrees, with intelligent machines through biofeedback and direct connections with the brain, that perhaps raise some of the most profound ethical questions regarding 3rd wave HCI. How might these technologies affect human social relationships? How might they change human identity, and what it means to be "human" itself?

## 15.6 Conclusions and Future Work

As technology develops through faster network connections, complex mobile devices, virtual and augmented mixed reality and autonomous systems, it can be seen how the 3rd wave of HCI advances immersion, embeddedness and socially networked connections.

This chapter utilized 3rd wave Positive Technology Framework to analyze how social networks merging with VR/Augmentation can mediate group experiences. The idea of greater immersion can further extend our knowledge of sensory input.

The chapter also looked at how AI can become our detractors but also, our guides through the information overload that results from being overly connected. The human identity and what it means to be human could be opened further to change with improvements of sensory connections. The mass collection of individuals' data and metric could lead to AI clones that duplicate a person in a cyber-world – whether or not such a person exists in reality.

The gap between human activity and machine interaction is decreasing as technologies advance. The immersion leads to transparent interactions and enhancement of sensory input. The direction this takes is guided by the users' needs and imagination which which are further ignited as possibilities are realised. Likewise, the fact that HCI appears to present itself in waves matches these relationships between users, technology, and information streams. What is made is guided not so much by contextual menus and the perception of particular interfaces, but the relationships within an immersive environment which engages increasingly more senses/feeling than previously explored.

## References

- Adler E, Hoon M, Mueller K et al (2000) A novel family of mammalian taste receptors. *Cell* 100(6):693–702
- Aghajan H, Lopez-Cozar Delgado R, Augusto J (eds) (2009) *Human-centric interfaces for ambient intelligence*. Amsterdam, Academic Press/Elsevier
- Augusto J, Callaghan V, Kameas A et al (2013) Intelligent environments: a manifesto. *Hum Cent Comput Inf Sci* 3:12. Springer. [www.hcis-journal.com/content/3/1/12](http://www.hcis-journal.com/content/3/1/12). doi: <https://doi.org/10.1186/2192-1962-3-12>
- Balkwill L, Thompson W (1999) A cross-cultural investigation of the perception of emotion in music: psychophysical and cultural cues. *Music Percept Interdiscip J* 17:43–64. <https://doi.org/10.2307/40285811>
- Braun M, Cheok A (2014) Towards an olfactory computer-dream interface. *ACM Press*, pp 1–3. <https://doi.org/10.1145/2663806.2663874>
- Bregman A (1990) *Auditory scene analysis: the perceptual organization of sound*. MIT Press, Cambridge, MA/London
- Brixen E (2007) Techniques for the authentication of digital audio recordings. In: *Audio engineering society convention 122*. Audio Engineering Society

- Chen LY, Tee BC, Chortos AL, Schwartz G et al (2014) Continuous wireless pressure monitoring and mapping with ultra-small passive sensors for health monitoring and critical care. *Nat Commun* 5:5028. <https://doi.org/10.1038/ncomms6028>.
- Chinellato E, del Pobil AP (2009) The neuroscience of vision-based grasping: a functional review for computational modeling and bio-inspired robotics. *J Integr Neurosci* 8(2):223–254
- Cremer L (1948) *Die wissenschaftlichen Grundlagen der Raumakustik*, Bd. 1. Hirzel-Verlag, Stuttgart
- Doty R, Shaman P, Applebaum S, Giberson R et al (1984) Smell identification ability: changes with age. *Science* 226(4681):1441–1443
- Downey L (2015) Well-being technologies: meditation using virtual worlds. PhD thesis, College of Engineering and Computing Nova Southeastern University
- Fletcher H, Munson WA (1933) Loudness, its definition, measurement and calculation. *Bell Labs Tech J* 12:377–430
- Fourier J (1822) *Théorie analytique de la chaleur*. Firmin Didot Père et Fils, Paris
- Groen F, Pavlin G, Winterboer A, Evers V (2017) A hybrid approach to decision making and information fusion: combining humans and artificial agents. *Robot Auton Syst* 90:71–85. <https://doi.org/10.1016/j.robot.2016.08.009>
- Harrison S, Tatar D, Sengers P (2007) The three paradigms of HCI. In: *Alt. Chi. Session at the SIGCHI Conference on human factors in computing systems* San Jose, California, USA, pp 1–18
- He W, Boesveldt S, de Graaf C, de Wijk R (2014) Dynamics of autonomic nervous system responses and facial expressions to odors. *Front Psychol* 5. <https://doi.org/10.3389/fpsyg.2014.00110>
- Herz R, McCall C, Cahill L (1999) Hemispheric lateralization in the processing of odor pleasantness versus odor names. *Chem Senses* 24:691
- Herz R, Eliassen J, Beland S, Souza T (2004) Neuroimaging evidence for the emotional potency of odor-evoked memory. *Neuropsychologia* 42:371–378. <https://doi.org/10.1016/j.neuropsychologia.2003.08.009>
- Hou C, Wang H, Zhang Q, Li Y, Zhu M (2014 Aug 6) Highly conductive, flexible, and compressible all-graphene passive electronic skin for sensing human touch. *Adv Mater* 26(29):5018–5024. <https://doi.org/10.1002/adma.201401367>
- ISO/IEC JTC1/SC29/WG11 (1999) *Information Technology- Coding of audiovisual objects—ISO/IEC.D 4496 (Part 3, Audio)*
- ISO226:2003E, *Acoustics – normal equal-loudness-level contours*
- Jones S (2016) Doing the right thing: computer ethics pedagogy revisited. *J Inf Commun Ethics Soc* 14(1):33–48. <https://doi.org/10.1108/JICES-07-2014-0033>
- Jones S, Hara S, Augusto J (2015) eFRIEND: an ethical framework for intelligent environments development. *Ethics Inf Technol* 17(1):11–25. <https://doi.org/10.1007/s10676-014-9358-1>
- Juhani P (2014) *The eyes of the skin: architecture and the senses*. Wiley, Chichester
- Kajstura M, Trawinska A, Hebenstreit J (2005) Application of the electrical network frequency (ENF) criterion. *Forensic Sci Int* 155(2):165–171
- Kim DH, Lu N, Ma R, Kim YS, Kim RH et al (2011) Epidermal electronics. *Science* 333:838–843
- Kobayashi Y, Habara M, Ikezaki H, Chen R et al (2010) Advanced taste sensors based on artificial lipids with global selectivity to basic taste qualities and high correlation to sensory scores. *Sensors (Basel)* 10(4):3411–3443. <https://doi.org/10.3390/s100403411>
- Kumar R, Bhondekar AP, Kaur R, Vig S, Sharma A, Kapur P (2012) A simple electronic tongue. *Sens Actuators B* 171–172:1046–1053
- Lantham R (2017) *The Coming Technological Singularity: How to survive in a post-human era*. In: *Science fiction criticism: an anthology of essential writings*. Bloomsbury Academic, London. ISBN: 9781474248624
- Li X, Staszewski L, Xu H, Durick K, Zoller M, Adler E (2002) Human receptors for sweet and umami taste. *Proc Natl Acad Sci USA* 99(7):4692–4696. <https://doi.org/10.1073/pnas.072090199>
- Linden DJ (2015) *Touch: the science of hand, heart and mind*. Viking Press, Toronto



- McGann JP (2017) Poor human olfaction is a 19th-century myth. *Science* 356(7263). doi:<https://doi.org/10.1126/science.aam7263>
- Moseley R (2016) Deep immersion with Kasina: an exploration of meditation and concentration within virtual reality environments. In: *Proceedings: the computing conference (2017)*
- Moskvitch K (2017) The Internet of things 2.0: when things start to listen. *Eng Technol* 12(1):63–65
- Nyquist H (1928) Certain topics in telegraph transmission theory. *Trans Am Inst Electr Eng* 47:617–644
- O’Connell M (2017) *To be a machine*. Granta Books, p 153. ISBN: 9781783781966.
- Prescott TJ, Diamond ME, Wing AM (2011 Nov 12) Active touch sensing. *Philos Trans R Soc Lond Ser B Biol Sci* 366(1581):2989–2995. <https://doi.org/10.1098/rstb.2011.0167>
- Ryan EB, Cacioppo T (1981) Dimensions of sociolinguistics In: Scherer KR, Giles H (eds) *Social markers in speech*. Cambridge University Press, 1979, p 395. *Lang Soc* 10:443. doi:<https://doi.org/10.1017/S0047404500008873>
- Scherer K (2003) Vocal communication of emotion: a review of research paradigms. *Speech Commun* 40(1–2):227–256. ISSN: 0167-6393. [https://doi.org/10.1016/S0167-6393\(02\)00084-5](https://doi.org/10.1016/S0167-6393(02)00084-5)
- Sheutz M, Arnold T (2016) Are we ready for sex robots. In: *Proceedings HRI ‘16 The eleventh ACM/IEEE International conference on human robot interaction, Christchurch, New Zealand* IEEE Press Piscataway, NJ, USA, pp 351–358. ISBN: 9781467383707
- Suzuki R, Homma S, Matsuura E, Okada K (2014) System for presenting and creating smell effects to video. *ACM Press*, pp 208–215. doi:<https://doi.org/10.1145/2663204.2663269>
- Toko K (2000) Taste sensor. *Sensors Actuators B* 64(1–3):205–215
- Tootell R, Tsao D, Vanduffel W (2003) Neuroimaging weighs in: humans meet macaques in “primate” visual cortex. *J Neurosci* 10:3981–3989
- Turkle S (2011) *Alone together*. Basic Books. ISBN: 9780465031467
- Villarreal B, Gordillo JL (2016) Bioinspired smell sensor: nostril model and design. *IEEE/ASME Trans Mechatron* 21:912–921. <https://doi.org/10.1109/TMECH.2015.2487260>
- Wang X, Dong L, Zhang H, Yu R et al (2015) Recent progress in electronic skin. *Adv Sci* 2(10):1500169. <https://doi.org/10.1002/advs.201500169>
- Weiser M (1999) The computer for the 21st century. In: *ACM SIGMOBILE Mobile computing and communications review – Special issue dedicated to Mark Weiser, vol 3, Issue 3*, ACM New York, NY, USA, pp 3–11. doi: <https://doi.org/10.1145/329124.329126>