

Building Tools for Creative Data Exploration: A Comparative Overview of Data-Driven Design and User-Centered Design

Sara Diamond, Steve Szigeti, and Ana Jofre^(✉)

OCAD University, 100 McCaul St, Toronto, ON M5T 1W1, Canada
{sdiamond, sszigeti, ajofre}@ocadu.ca

Abstract. Visualization scientists seek means to inspire insights from data, which require creative thinking on the part of analysts as well as cognitive reasoning. In information visualization a focus on the user has proven highly effective in the design of usable and engaging interfaces, although it has been argued that such a focus limits innovation in insights about the data and in the creation of metaphors for visualization. If a user-centered design recapitulates existing knowledge, then a design approach which derives exclusively from the data may provide more innovative results. Our approach considers both the designers and the users, whereby our goal is to elicit creativity in both the design of visualization tools and in their application. We compare user-centered design and data-driven design through tool sets that emerged from each of these methods. User-centered design methodologies were used in the creation of a custom interface for editors at a major national newspaper that visualizes measures of each story's popularity. Data-driven design methodologies were used to create a tangible user interface for data visualization. With UCD we built a tool that supported the use of data in editorial decisions and deployed familiar metaphors to encourage significant change in workplace practice. With DDD we unleashed creativity on the part of analysts which resulted in a more innovative approach on the part of designers and a gateway to new user communities. We compare strengths and weaknesses of each methodology through a reflection of our design outcomes.

Keywords: Data visualization · Design · Innovation

1 Introduction

“If we’d ask customers what they wanted, they’d ask for faster horses”, is a popular quote attributed to Henry Ford to illustrate potential problems when involving users in the design process. User centered design often results in successful tool development and refinement. However, such results are incremental changes, not radical innovation [1]. This challenge of innovation in design lies at the center of how we visualize data.

Data visualization requires the adoption of appropriate visual metaphors, where the choice of metaphor is coupled with both qualities of the data and an understanding of potential data use. We use the term metaphor in the context of interface design, where

qualities of a familiar object and its behaviors in the real world are carried into the digital space to facilitate navigation [2], such as the icons for files and folders on a digital desktop. Metaphors in our usage includes familiar ways of representing numerical quantities such as graphs. In a data-driven approach, designers develop metaphors and interactions based on the data and its structure [3], while a user-centered approach first considers user hypotheses and requirements [4]. Innovation and creativity are tightly coupled, and the ability of data visualization tools to provide insight is a potential measure of creativity.

User-centered design (UCD) has been employed in the design of interfaces for decades and is increasingly prevalent within data visualization as it allows for the dynamic management of visual metaphors and underlying data structures within the context or culture of the user [2, 4, 5]. Many UCD approaches seek to understand their users through ethnography studies [6, 7] where users and their practices are studied by designers, and in more mature practice, through participatory engagement, where the intended user plays an active role in the design process [8, 9]. These practices were motivated by the need to capture tacit knowledge about work process within the design of technology through the recognition that systems designed without this knowledge were often flawed and less readily adopted. The best of UCD engages users through charrettes, self-documentation and prototyping.

It is important to recognize two subtle challenges related to a user-centered approach; (i) UCD caters its aesthetics to a particular user, thus reinforcing a normative approach to the data, and (ii) UCD begins with assumptions regarding users' queries of the data. Arguably a data-driven approach, lacking an initial focus on a specific user, may result in tools that are viable for a broader set of potential users and may reveal patterns from the data that do not conform to users' queries or assumptions.

Data-driven design (DDD) begins by considering the data in the absence of knowledge regarding its specific uses in order to foreground relationships within the data that may not have been previously considered. A DDD approach first considers the data and its structure, and then, in subsequent steps, looks at potential applications and user culture [3]. Raw data, in the absence of a known user, suggests the potential of analyses based on visualizations with new aesthetics and interpretations that may not have otherwise been proposed. As Benjamin Fry points out, aesthetics structure experiences in formal perceptual ways and provide interpretive tools with which to construct meanings [10].

While a user-centered approach has resulted in successful visualizations, the results will not necessarily be innovative [1]. Innovative ideas are fragile at their inception, and strong creative visualizations risk rejection [5], or are never invented due to a focus on perceived user needs. An approach induced through data-driven analysis could help guard against one of the dangers of UCD; namely, that questions and metaphors (such as familiar graphs) only recapitulate existing knowledge [5]. This is particularly relevant when the goal of the analytics tools is to prompt insight and creativity rather than to augment rapid decision making.

With both UCD or DDD, data visualization eventually requires the adoption of appropriate visual metaphors, where the choice of metaphor is coupled with both the qualities of the data and an understanding of potential data use. In DDD, designers

involve users at later stages of the design process than in UCD. Namely, in UCD, the user is involved from the beginning, while in DDD, the user becomes involved after a prototype has been produced.

In this work, we compare both approaches to designing visualizations. In part 1, we examine the outcomes of a pilot study in which one team of designers explicitly engages in DDD, while the other team practices UCD to create visualizations. In part 2, we examine two case studies from our research. Our UCD case study is the design of a custom interface for editors at a major national newspaper that visualizes measures of each story's popularity, and our DDD case study is on the design of a tangible user interface for data visualization. The case studies illustrate (1) the creativity of designers working through the two approaches, and (2) how the tools might best support the creativity of the users: data analysts.

While research has considered the shortcomings of UCD as a means to innovate [11], there has been no comparative examination of the different outcomes of UCD and DDD in the domain of data visualization. Our research hopes to contribute an early step in this direction.

2 Part 1: Pilot Study

We begin our work by discussing a pilot study comparing different design approaches while using the same data to reduce a number of potentially confounding variables. The goal was (i) to evaluate the kinds of insights the DDD and UCD approaches provide from the data on the part of analysts, (ii) to compare the metaphors that emerge from these methods, and (iii) to develop a deeper understanding of both methods with a goal of potentially articulating new design methods, including those that support creativity on the part of designers.

In order to compare the outcomes of the UCD and DDD approaches, we assigned a team of graphic designers to each method. Both teams were provided with identical data sets. The UCD team worked closely with NLogic staff in order to understand user needs, whereas the DDD team had no access to existing users of the data.

2.1 Data Sets

Compiled by our industry partner, the data consists of detailed demographic information on radio listeners in Canada, and media consumption habits of the market research participants. Additionally, both teams had access to information regarding participant recruiting and data collection protocols.

2.2 Team Composition

Each of the two teams comprised three graphic designers and each is led by a researcher. The designers varied in their experience with developing metaphors and user interfaces, but all had existing design experience and training. In addition, a third researcher provided guidance to the teams and developed the experimental hypothesis.

2.3 Design Process

Both the UCD and the DDD teams worked through their particular design process in parallel. Figure 1 shows the procedure for both teams. The design process began with data analysis, then a design phase, followed by prototyping and then evaluation. The teams repeated the cycle in order to further refine the designs, aiming towards high fidelity prototypes. In each cycle, the main difference between the two teams occurred during data analysis and evaluation.

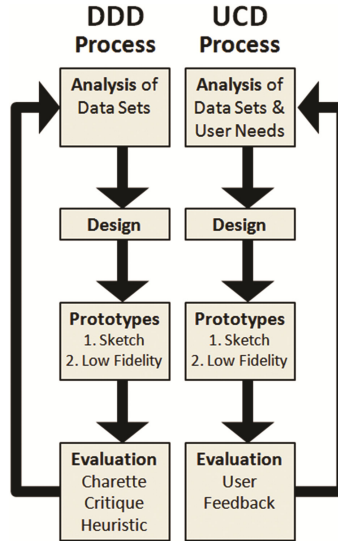


Fig. 1. The UCD and DDD design processes.

In these two phases, the UCD team's focus was on information provided by users, while the DDD team remained isolated from users. During the data analysis phase, the UCD team gathered data on the users, including developing personas, writing and administering questionnaires and conducting task-analysis [5]. The DDD team was provided with the raw data files, and guided on how to identify relationships in the data that would be interesting to visualize by users. Evaluation for the UCD team involves users, while the DDD team relied on the use of charrette's, critiques and heuristic evaluation [2].

Throughout the design process, the UCD and DDD teams did not communicate nor share experiences regarding their work. We believe this allowed us to better understand differences in the two approaches.

2.4 Observations

The themes emerging from the data were divided into three categories: (i) a focus on listener behavior and consumption patterns, (ii) a focus on geographical aspects of the

data, and (iii) a focus on temporal qualities of the data. While these themes are consistent in both UCD and DDD, each team developed distinctly different insights and metaphors.

2.5 UCD Results

As expected, by working closely with users of existing NLogic information visualization tools, the early designs of the UCD team represented incremental improvements to the existing tools which were intended to correct shortcomings with existing data visualization metaphors and interfaces. Even when encouraged to try more unorthodox designs, the sketches remained tightly coupled with the results of user task analysis.

For example, Fig. 2 shows the use of a tree metaphor which is not only familiar in the field of information visualization, but the suggested interaction between user and prototype was identical to actions described in the needs assessment. This prototype changed one means of analyzing the data for another by altering the aesthetics of the interface and making improvements in usability, but failed to evoke new perspectives on the data. This is an important point, since the incorporation of a different metaphor ideally reveals something in the data that the previous metaphor did not.

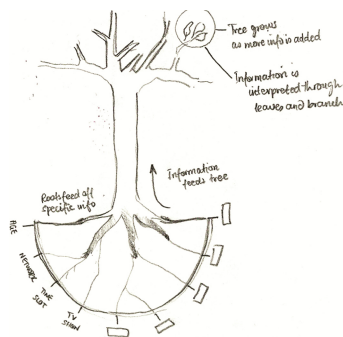


Fig. 2. UCD sketch: analogy of tree, where dynamic queries are performed through the roots, which in turn make the tree grow based on what it's been fed with.

2.6 DDD Results

In contrast to the UCD team prototypes, the prototypes that emerged from the DDD team made use of more granular data and considered the nature of how the data itself was collected. The designs introduced metaphors intended to present data about radio stations' listener demographics to a general audience in a novel way (for example via physical radio sets blaring in the hallways of the NLogic offices, as showed in Fig. 3). Some of the prototypes propose visualizations that allow technical staff at NLogic to monitor panelists' activity and identify potential anomalies and outliers in the data.

2.7 Existing Design Bias

During the first phase of sketching most of the designers came up with “typical” UI visualizations, including traditional widgets such as buttons and checkboxes to navigate the data. We believe this speaks to the existing preconceptions and training of many designers. Designers are trained to consider and imagine the user and hence bring assumptions about users to their work even in the absence of a clear definition of user needs. This bias can limit initial attempts at innovation and experimental design in both a UCD and DDD approach. Creative methods from the design world, such as probes and brainstorming techniques, have been created to subvert existing bias and focus on the user’s immediate needs and expectations rather than the potential of a design [12].

2.8 Conducting Research with an Industry Partner

Conducting research in the context of an industrial partner presents several challenges. Sedlmair et al. [13] discuss issues that also arose in our work such as access to data, selecting an evaluation context, finding domain expert participants, confidentiality of information and engaging with complex work processes. Additional issues pertain to conducting concurrent design approaches involving the same lead researchers and representatives of the industrial partner who had contact with both design teams. While working on the same data visualization project, our comparative study imposed two separate approaches, which can cause potential confusion from our partners’ point of view when exchanging information with one or the other team. It is essential to ensure that the representative of the industrial partner both have the time to devote to concurrent projects, and perfectly understand the design protocols of each.

2.9 Keep the Design Teams Separated, yet Synchronized

To guarantee a valid design protocol within the two design teams, they must not share information between each other. This can become an issue, when sharing the same work space, and organizational measures should be put into place so as to minimize the odds of the designers mistakenly exchanging information across teams. At the same time, we argue that it is preferable that the teams remain more or less synchronized in the overall design process and also for final evaluation with users.

2.10 Filtering Out Information from Users

While working with a large, complex database comprising many data tables and relations, it was necessary to maintain contact with our industrial partners, asking for clarifications and additional material (for example, the data collection protocol employed). In their answers to our questions, our partners sometimes provide user-contextual information that, if communicated to the DDD team, potentially introduces additional bias.

The design problems the two teams chose to address differed. On the one hand, the problems expressed by current users of the visualization tools employed to analyze the data shaped the design work of the UCD team, whereas the DDD team found problems

to address via the data tables themselves. The raw data, in the absence of a known user, suggest visualizations, new aesthetics, uses and interpretations that may not have otherwise been proposed. We consider this a measure of creativity. While we did not expect to find examples of radical innovation (and we have not), we have found differences in outcomes from the UCD and DDD teams that are worth exploring further.

Importantly, the DDD team tended towards designs that communicated the statistical limitations of the NLogic dataset, which the designs of the UCD did not. This was a finding of particular interest to our industrial partner, who expressed struggles with reminding their clients of the limits of the data itself. Gathered from a large population sample, the increasing subdivision of the data results in sample sizes that are statistically insignificant (for example, from Canadian women ($n = 1200$) to women in Toronto ($n = 200$) to women between the ages of 18 and 24 ($n = 42$) to women who listened to over 3 h of radio per day ($n = 8$)). Clients had been considering the data at such a granular level, that the perceived correlations between data points was no longer valid.

In subsequent projects, described below, we continue to consider the differences between a UCD and a DDD approach to addressing design challenges. Work with The Globe and Mail (the Sophi Heads-Up Display) represents a UCD approach, whereas the Tangible User Interface research project emerged directly from our work with NLogic, continuing a DDD approach.

3 Part 2: Case Studies

3.1 User-Centred Design Case Study: The Sophi Heads-up Display

The Sophi Heads-up Display (HUD) is an analytics tool developed with a user-centered design approach for editors at Canada's The Globe and Mail newspaper. It overlays relevant data about articles' performances onto The Globe and Mail website, allowing news editors to easily make data-driven decisions throughout their workday. This technology represents a significant change within the editorial practices of newspapers. The Sophi HUD is a sign of the significant disintermediation occurring within media industries where data analysis increasingly supports traditional editorial decision-making. Such systems are intended to facilitate editorial decisions which are no longer purely intuitive but engaged with data regarding readership. The Sophi HUD shows article performance over time in relation to reader engagement with content and related advertisements. It introduced alerts on rapidly trending data, alerts editors to articles that may be underperforming in relation to editorial expectations, provided access to more granular data regarding an article (such as the number of readers and length of time of engagement over time) in order to help identify underlying reasons for the article's performance, and introduces a chart the editor can use to view and compare the popularity of all articles appearing on The Globe and Mail website that were published in the previous 24 h.

Design of the Sophi HUD was informed by intensive collaboration with The Globe and Mail employees. Data on editorial practice and requirements emerged from 20 semi-structured interviews with a broad representation of editorial team members. This data was supplemented by participant observation through job shadowing, cataloguing the

data analysis tools that they were currently using, and undertaking design work in situ with users. These practices recognize that work environments influence cognitive processes and must be considered when designing a software system [14]. Our close interaction with editors and other stakeholders (including the data analytics team at the Globe and Mail) provided us with the context from which the requirements emerged, allowing us to modify our initial designs with a better understanding of the challenges faced by users of our proposed system.

The Sophi HUD has undergone a number of iterations, each informed by user feedback generated from interviews, training sessions, participatory design charrettes, and informal requests. The Globe and Mail has integrated the Sophi HUD into a number of editors' daily routines, and it continues to evolve as users learn from the data presented to them. The editorial team adopted the new tool at a time of significant transition in the use of data in making editorial decisions. The goal of this technology is to facilitate new kinds of insights, using data as well as "gut instincts". The Sophi HUD supports the creative problem solving requirements that are fundamental to editorial tasks, but does not provide a new way of thinking or metaphor regarding through the nature of how a digital newspaper is experienced by its readers.

The UCD approach was appropriate in this project because the users had very clear requirements and needs based on the data the organization collected. Innovation was not the goal; instead, we focused on developing improvements over tools that the editorial team had utilized, but which had not adequately addressed editorial needs. Both Chartbeat [15] and Parse.ly [16] are web analytics software products that are familiar to the editorial team users. Composed of a dashboard and a heads-up display system, these two products provided the users with examples of data visualization and analysis tools that could positively support editorial responsibilities. Knowledge of these software products, in particular their strengths and weaknesses, informed the design of the Sophi HUD. Further research would warrant a DDD approach to see if different kinds of realizations emerge from the data, once the editors become more familiar with the Sophi HUD.

3.2 Data-Driven Design Case Study: A Tangible User Interface for Radio Listenership Data

Emerging directly from our initial study with NLogic, the DDD group developed a series of metaphors making use of physical interfaces with the data. Sketches were developed which included proposals for a wall of transistor radios (where the position of the antenna and the volume represented different data points, see Fig. 4), a national map with 3D data representations (where the volume and height of the objects represented data points connected to specific geographic locations, see Fig. 5) and a table with mechanical functions that would raise or lower various columns, each representing an aspect of listenership data, inspired by the inForm interface from MIT [17]. While none of these sketches led to prototypes, when shared with potential users (our industrial partners) the notion of improving collaboration emerged as a design challenge. While the industrial



Fig. 6. The Sophi Heads-Up display.

partners explained that many collaborations take place remotely by means of telecommunications, they identified a need for a tool encouraging people to not only work on the same problem, but to work on it together at the same time in the same room (Fig. 6).

To address this design challenge, we created a prototype for a Tangible User Interface (TUI) for interactive data visualization. In this instance DDD processes focused on giving the designers the largest creative range possible, outside of immediate user needs. Research indicates that graspable tangible interfaces measurably increase collaborative behavior when compared to screen-based interfaces [18–22]. Our hybrid system combines a tabletop graspable user interface with a two-dimensional screen display; the users interrogate the data by manipulating tokens on the tabletop and the screen displays the results of the user’s query. The tokens are tracked by means of a camera placed discretely beneath the transparent tabletop. The bottom of each object is marked with a fiducial marker, and the camera placed below the table captures the image of the fiducial markers in real time. The fiducial markers are read using open-source reactIVision software [23]. The reactIVision software outputs the position of the markers, if they are in the field of view of the camera, and this information is input into our software, which constructs the visualizations from a (user-provided) database (Fig. 7).

While the first prototype used NLOGIC’s radio station listener demographic data to create a market research package for advertisers targeting radio air time [24, 25], this system can be (and has been) generalized to query various data sets. The software is designed to allow users to provide their own formatted data, and there are no specialized hardware requirements. The system consists of our software, a webcam, and fiducial markers with which to tag the tokens. The fiducial markers can be printed and attached to any object the user wishes to use to explore their data set.

We believe that the TUI combines embodied cognitive processes, with multiple perspectives on the data from collaborating users, thus opening up the possibilities of new discoveries about the data. As an outcome of a DDD approach, TUI represents an innovation that did not emerge from any existing analytics tools available to the user. Only by removing an understanding of the requirements and needs of users,




Fig. 7. Schematic of interactions with the tangible interface

was the design team able to propose conceptual tools that open a window on different interactions between - and with - data.

4 Discussion and Conclusion

Different design challenges require different approaches. In order to marshal the creativity of designers and the creativity of users, it is important to understand the broader requirements of a project: is innovation a necessary component or is there instead a need for incremental improvement over existing processes?

When the design challenge is focused on the improvement of an existing set of tools, then a UCD approach is most effective. Incremental improvement requires an acknowledgement that existing insights emerging from data sets should not be replaced, but rather that they need to be observed more efficiently or in a more nuanced manner. In other words, existing approaches provide cognitive aids that can be improved through the design process.

Where new or different insights are required, a UCD approach may prove less successful because of the fundamental focus on perceived user needs, and the constraints then returned to the analysts in the tool provided. Articulating user needs risks the elimination of innovative designs, particularly if there are existing data visualization tools used to analyze the data set.

However, while research has considered the shortcomings of UCD as a means to innovate [26], we have found it to be an effective means for creating tools that can support changes within workplace practices and allow prospective users to feel ownership and engagement with the technology. The Sophi HUD was found to be immediately useful for the editorial team at the Globe and Mail. Various users at the newspaper have adopted this visualization platform, and we continue to work closely with them to develop improvements and to meet their emerging needs. Sophi HUD emerged from close

consultation with users in order to address the shortcomings of existing tools. As a result, workplace practices have changed dramatically as editors increasingly understand the relationship of readership data to their particular practice. The Sophi HUD has successfully supported the cognitive requirements of its users, proven by its adoption and continued use.

While the Sophi HUD was successfully designed to optimise the needs of a specific set of users, a key outcome of the DDD approach appears to be versatility of use. The TUI system fits a variety of users and situations. Since the initial design did not respond to a set of requirements and needs, there are fewer limits on how the tool may be applied. For example, the TUI system has been adopted by school teachers for classroom use. By not focusing a narrow set of user needs, we were able to create a software system that allows users to explore different data sets in different environments and situations. Whereas the UCD approach resulted in a tool with limited use cases, the DDD approach provides cognitive support for a wide variety of users and data sets. The application of the TUI is bound only by the creative capacities of potential users, which are broad.

The strength and weakness of each system (Sophi HUD versus TUI) is not necessarily found in the ultimate usefulness of the system – but rather in the appropriateness of the design processes to specific design challenges. Sophi HUD addresses a clear set of requirements. Editors have previously worked with similar software and were able to articulate the pros and cons of Chartbeat and Parse.ly. Thus, the design of the Sophi HUD incorporates the strengths of existing software, removes unnecessary functionality, and produces a bespoke tool that is tightly coupled with user needs and served effectively in a significant work force transition. The various methodologies incorporated in the UCD approach provided the designers with objectives and a clear direction in terms of constrained but effective creativity.

The DDD approach resulted in systems that users found to be a significant innovation over existing tools. It is doubtful that the TUI system would have emerged from a UCD approach, which insists on a tight connection between designer and user and a clarity of user requirements. The TUI represents innovation and a leap forward in terms of tools that the industry partner (NLogic) would incorporate into their work practice. However, the shift to using a tool that encourages curiosity and collaboration requires a fundamental shift in terms of work practice. The Sophi HUD was deemed trustworthy because it incorporated existing user relationships with a data visualization tool; the TUI asks the user to abandon their current work practice in exchange for both data insights that may be radically different from existing software applications, and for facilitated collaboration.

DDD is not an approach that eliminates the need for involving a user in the design process. We undertook user testing and feedback into various iterations of the TUI. In this process the design team identified a series of new uses for the TUI well beyond its original context. DDD re-orders the steps of designing for a user group – instead of gathering requirements and needs and then looking for means to address those needs, the designer first considers the data and the potential relationship of data points to propose tools of analysis which are then refined based on requirements and needs. The user should never be removed from the design process, but where they are

positioned in the initial design stage may be the difference between iterative and innovative design outcomes.

The TUI has proven valuable to a broad range users and analysts, specifically in the field of education where the goal is to encourage free-ranging exploration of a wide range of data sets in a collaborative environment. On the other hand, the Sophi-HUD rapidly became an indispensable tool in the editorial room at The Globe and Mail. We will continue to explore and compare both the UCD and DDD design methods to better understand the context and applications in which each of these approaches best supports creativity and insight.

Acknowledgement. The authors wish to thank Bianca Channer, Lan-xi Dong, Alex Gebhardt, Stephen Keller, Lauren Livingston, Anton Mwewa, Anne Stevens, Borzu Talaie and Phuong Ha Vu. We thank our industrial partners Jonathan Lee and David Phillips from NLogic and Gordon Edall from the Globe and Mail for their time and valuable input. We also note with thanks the important contributions of Fanny Chevalier (to the pilot study) and Jessica Peter (to the Sophi HUD case study). This research is supported by the BRAIN Alliance via the Ontario Research Fund - Research Excellence (ORF-RE) program. Additional funding provided by the NSERC Strategic Partnership program, MITACS, and the OCE-VIP program.

References

1. Norman, D.A., Verganti, R.: Incremental and radical innovation: design research vs. technology and meaning change. *Des. Issues* **30**(1), 78–96 (2013)
2. Rogers, Y., Sharp, H., Preece, J.: *Interaction Design: Beyond Human - Computer Interaction*, 3rd edn. Wiley, Chichester (2011)
3. Tufte, E.R.: *The Visual Display of Quantitative Information*, 2nd edn. Graphics Press, Cheshire (2001)
4. Espinosa, O.J., Hendrickson, C., Garrett, J.H.: Domain analysis: a technique to design a user-centered visualization framework. In: *1999 Proceedings of IEEE Symposium on Information Visualization, (Info Vis 1999)*, vol. 144, pp. 44–52 (1999)
5. Gould, J.D.: Human-computer interaction. In: Baecker, R.M., Grudin, J., Buxton, W.A.S., Greenberg, S. (eds.) *The Art of Human Computer Interface Design*, pp. 93–121. Morgan Kaufmann Publishers Inc., San Francisco (1995)
6. Button, G.: The Ethnographic Tradition and Design. *Des. Stud.* **21**(4), 319–332 (2000)
7. Salvador, T., Mateas, M.: Introduction to design ethnography. In: *CHI 1997 Extended Abstracts on Human Factors in Computing Systems, CHI EA 1997*, pp. 166–167. ACM, New York (1997)
8. Bjögvinsson, E., Ehn, P., Hillgren, P.-A.: Design things and design thinking: contemporary participatory design challenges. *Des. Issues* **28**(3), 101–116 (2012)
9. Schuler, D., Namioka, A.: *Participatory Design: Principles and Practices*. CRC Press, Boca Raton (1993)
10. Reas, C., Fry, B.: *Processing: A Programming Handbook for Visual Designers and Artists*. MIT Press, Cambridge (2007)
11. Greenberg, S., Buxton, B.: Usability evaluation considered harmful (some of the time). In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2008*, pp. 111–120. ACM, New York (2008)

12. Gaver, W.W., Boucher, A., Pennington, S., Walker, B.: Cultural probes and the value of uncertainty. *Interactions* **11**(5), 53–56 (2004)
13. Sedlmair, M., Isenbergh, P., Baur, D., Butz, A.: Information visualization evaluation in large companies: challenges, experiences and recommendations. *Inf. Visual.* **10**(3), 248–266 (2011)
14. Suchman, L.: *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge University Press, Cambridge (2007)
15. Chartbeat | Building the Attention Economy (2017). <https://chartbeat.com/>. Accessed 7 February
16. Parse.ly | Audience Data & Content Analytics for Digital Media (2017). <https://www.parse.ly.com/>. Accessed 7 February
17. Follmer, S., Leithinger, D., Olwal, A., Hogge, A., Ishii, H.: inFORM: dynamic physical affordances and constraints through shape and object actuation. In: *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, UIST 2013*, pp. 417–426. ACM, New York (2013)
18. Antle, A.N., Wise, A.F.: Getting down to details: using theories of cognition and learning to inform tangible user interface design. *Interact. Comput.* **25**(1), 1–20 (2013)
19. Kim, M.J., Maher, M.L.: The impact of tangible user interfaces on spatial cognition during collaborative design. *Des. Stud.* **29**(3), 222–253 (2008)
20. Schneider, B., Jermann, P., Zufferey, G., Dillenbourg, P.: Benefits of a tangible interface for collaborative learning and interaction. *IEEE Trans. Learn. Technol.* **4**(3), 222–232 (2011)
21. Schneider, B., Sharma, K., Cuendet, S., Zufferey, G., Dillenbourg, P., Pea, A.D.: 3D tangibles facilitate joint visual attention in dyads. In: *International Conference on Computer Supported Collaborative Learning (CSCL)*, pp. 158–165 (2015)
22. Ma, J., Sindorf, L., Liao, I., Frazier, J.: Using a tangible versus a multi-touch graphical user interface to support data exploration at a museum exhibit. In: *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction, TEI 2015*, pp. 33–40. ACM, New York (2015)
23. Kaltenbrunner, M., Bencina, R.: . reactIVision: a computer-vision framework for table-based tangible interaction. In: *Proceedings of the 1st International Conference on Tangible and Embedded Interaction, TEI 2007*, pp. 69–74. ACM, New York (2007). doi:[10.1145/1226969.1226983](https://doi.org/10.1145/1226969.1226983)
24. Jofre, A., Szigeti, S., Keller, S.T., Dong, L.-X., Czarnowski, D., Tomé, F., Diamond, S.: A tangible user interface for interactive data visualization. In: *Proceedings of the 25th Annual International Conference on Computer Science and Software Engineering, CASCON 2015*, pp. 244–247. IBM Corp., Riverton (2015)
25. Jofre, A., Szigeti, S., Tiefenbach-Keller, S., Dong, L.-X., Diamond, S.: Manipulating tabletop objects to interactively query a database. In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA 2016*, pp. 3695–3698. ACM, New York (2016). doi:[10.1145/2851581.2890260](https://doi.org/10.1145/2851581.2890260)
26. Crabtree, A., Rodden, T., Tolmie, P., Button, G.: Ethnography considered harmful. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2009*, pp. 879–888. ACM, New York (2009). doi:[10.1145/1518701.1518835](https://doi.org/10.1145/1518701.1518835)