Crim Law and Philos (2008) 2:163–180
DOI 10.1007/s11572-007-9042-1

ORIGINAL PAPER

# Ambient Intelligence, Criminal Liability and Democracy

**Mireille Hildebrandt**

Published online: 27 October 2007

**Abstract** In this contribution we will explore some of the implications of the vision of Ambient Intelligence (AmI) for law and legal philosophy. AmI creates an environment that monitors and anticipates human behaviour with the aim of customised adaptation of the environment to a person's inferred preferences. Such an environment depends on distributed human and non-human intelligence that raises a host of unsettling questions around causality, subjectivity, agency and (criminal) liability. After discussing the vision of AmI we will present relevant research in the field of philosophy of technology, inspired by the post-phenomenological position taken by Don Ihde and the constructivist realism of Bruno Latour. We will posit the need to conceptualise technological normativity in comparison with legal normativity, claiming that this is necessary to develop democratic accountability for the implications of emerging technologies like AmI. Lastly we will investigate to what extent technological devices and infrastructures can and should be used to achieve compliance with the criminal law, and we will discuss some of the implications of non-human distributed intelligence for criminal liability.

**Keywords** Criminal liability · Agency · Distributed intelligence · Ambient intelligence · Democracy · Rule of law · Causality

## Introduction

In modern legal systems criminal liability is based on the concept of agency, which is connected with such notions as individual freedom, intentional action and individual accountability. The same concept of agency is crucial for theories of democracy, since the achievement of a shared policy depends on representation, deliberation and participation by individual human agents. Thus, criminal liability in a constitutional democracy can be

M. Hildebrandt (✉)
Erasmus University Rotterdam, Rotterdam, The Netherlands
e-mail: hildebrandt@frg.eur.nl

M. Hildebrandt
Faculteit Rechten en Criminologie, Vrije Universiteit Brussel, Pleinlaan 2, Brussel 1050, Belgium

said to depend in two ways on a concept of human agency: first, because decisions about the scope of the criminal law demand democratic consent that presumes human agency and second, because criminal liability invokes the censure of punishment and such censure does not make sense if the addressee of the violated legal norm is not an agent in the sense referred to above.

In this article I will explore some of the issues arising from the emergence of distributive (human and non-human) intelligence. How should we understand criminal liability if behaviour is initiated or controlled by intelligent machines, and how does such machine-induced violation of the criminal law reflect on our understanding of democracy? How should we, for instance, conceptualise the link between human freedom and causality—crucial in the case of criminal liability for harm caused—in a situation that does not permit a strict separation of causal determination and free will?

To prepare for this exploration I will first discuss the vision of Ambient Intelligence and the technology of autonomic computing that is a precondition of its realisation (Section "Ambient Intelligence: When Things Come Alive"). After that I will briefly move into philosophy of technology to elucidate the relationship between humans and the technological infrastructures they create and depend on. In doing this I aim to destabilize the Cartesian divide between causal determinism and mental voluntarism, paradoxically seeking a more fuzzy and more precise perspective on human and non-human action. The point of this exercise will be to argue for a conceptualisation of technological normativity, as being pertinent for democratic accountability in the information age (Section "Human and Non-human Actants"). In the last section I will assess the implications of this analysis for criminal liability. First, I will discuss to what extent new technologies can and should be used to achieve compliance with the criminal law. Second, I will briefly assess some of the implications of non-human distributed intelligence for criminal liability (Section "Technological Normativity and Criminal Liability").

## Ambient Intelligence: When Things Come Alive

### Interconnectivity of Things

The European Commission and commercial business enterprises are investing substantial research funding into what has been called Ambient Intelligence (AmI).[1] AmI denotes an 'intelligent' environment that 'knows' your preferences and adapts itself to cater to your inferred wishes. The 'intelligence' of such an environment depends on ubiquitous computing, which is computing that is invisibly embedded in the environment and no longer requires deliberate human input. Instead of intentional programming of the environment by its users, the environment anticipates preferences and adjusts to them even before the user becomes aware of them.[2] *All things* in the environment are tagged with wireless devices that contain information that can be read by reading devices, which are connected to online databases.[3] Together with sensors and cameras the movements, temperature, sound and other relevant data of things and persons are continuously collected, stored and processed in order to detect certain patterns that allow this technological infrastructure to anticipate

[1] ISTAG (2001), Aarts and Marzano (2003).

[2] Lessig (1999), p. 154.

[3] See the Communication of the European Commission to the European Parliament (2007) and the draft version of the European Policy Outlook RFID (2007).

how one prefers to live. This anticipation is based on correlating seemingly trivial data; for instance, a correlation may be found between the length of pauses in one's key-stroke behaviour and the act of making coffee. The correlation may be provide a surprising accuracy in predicting at precisely which point you will get up and move into the kitchen (or go down to the coffee-bar). Such pattern recognition depends on the real time accumulation of data in online databases, allowing what is called 'knowledge discovery in data bases' (KDD),[4] thus gaining an unprecedented access to the habits and life-styles of individual persons based on their previous behaviour patterns. Interestingly machines may thus come to 'know' things about a person that she was not aware of herself.

The claim that the environment has a certain type of intelligence is based on the fact that it is capable of learning and adapting.[5] The 'it' of the previous sentence refers to a network of interconnected things that are communicating data, allowing real time awareness of states, movements and interactions and real time anticipation of habits that are translated into preferences. The awareness, the anticipation and the adaptation of the environment reside in the interconnections of the network: the intelligence does not arise in one particular thing but emerges in the permanent interplay of things turned into machines. Such an environment has been called *The Internet of Things*, because it seems to turn the offline world online.[6]

## Autonomic Computing

A crucial technology for AmI is what has been called 'autonomic' computing, because of the way in which the collection and aggregation of data, the search for significant and relevant patterns and the subsequent adaptation of the environment all take place without human intervention. The networked environment depends on programs that allow the machines that perform all these actions to make inferences and decisions without disturbing the human users of this environment with queries about their needs. The term 'autonomic' suggests that the external environment is regulated by AmI in the same way that our internal environment is regulated by the autonomic nervous system.[7] The autonomic nervous system is continuously monitoring all our vital bodily functions, in order to continuously adapt the temperature, heart rate, hormonal levels etc. It is crucial for our survival and we cannot begin to imagine its complexities. Computer scientists suggest that AmI will function in a similar way: monitoring and adapting the environment in an unobtrusive way, as it were 'under the skin', without bringing the decisions that need to be taken to our conscious attention.

Living in a environment that is aware of me to an extent previously reserved for animistic worldviews in ethnographic narratives may be experienced as if things have become alive. To provide the reader with a taste of the vision of AmI I will quote Kevin Kelly, co-founder of *Wired Magazine*, who seems to believe that we are on the verge of a 'fantastic' future:

[4] Custers (2004), Hildebrandt (2006b).

[5] Ambient Intelligence cannot be equated with Artificial Intelligence (AI) as criticised by, e.g., Dreyfus in his groundbreaking (1972). The reason is 2-fold: first, AI research has incorporated some of the criticisms and downsized some of its claims regarding machine intelligence, second, in this case AI techniques are used to create AmI, which is more contextual, user-centric and practice oriented than 'old-style' AI.

[6] ITU (2005).

[7] Kephart and Chess (2003).

> My daily travels should cross an ecology of machines, all alive in their interactions with me and each other. (…) Imagine if our toaster remembered who we were, and had a message from our refrigerator about the availability of bread, and was updated with the most recent smarts via the Webmaster, and was repaired by itself, and kept track of my calories, and was in communication with all the other toasters in the world to better itself—why, the world of machines would be liberated at last.[8]

On the basis of their interconnections things begin to communicate with us in a direct way, not via the keyboard of a computer, not via deliberate input. These things know what I will probably want, not what I think or say that I want, and it may be that this probability is more accurate than my own conscious reflection about my own habits, life-style, objectives or intentions. My own actions will be tuned *to and by* this environment in a persistent process of mutual adaptation. Some have predicted that this customisation of the environment will normalise us into certain types of people without ever having asked for our advice, locking us into a comfortable cage that is based on our previous behaviour patterns.[9] Others have argued that the availability of data from past interactions will undermine our right to oblivion, providing an ever increasing resource of information about our inclinations that can be held against us at any point in time.[10] Still others have warned that the filtering of undesired states of affairs, news, people and situations will threaten the basic tenets of a viable democracy because we are no longer confronted with dissent, with the unexpected, with everything that causes growth and calls for compromise.[11] In this contribution we will focus on another issue that may arise from the emergence of autonomic computing, as a precondition for AmI. What happens to criminal liability if it is unclear who has been violating a legal norm: a thing, a network of interconnected machines, their programmers, the service providers that initiated the network to gain a profit or the user who acted upon decisions taken without her conscious awareness?

## Human and Non-human Actants

### What Things Do[12]

From the perspective of Enlightenment philosophy, criminal liability presumes an individual human subject capable of intentional action. Whilst civil liability has worked with strict liability for a long time, the default position for criminal liability—at least in theory—still seems to be that both wrongfulness and culpability are necessary. To suggest that things could act wrongfully and be culpable for doing so seems open to ridicule, reminding one of medieval trials in which dogs were punished for criminal conduct. Even if one acknowledges that in some modern legal systems organisations with legal personality may

[8] Quoted in: Aarts and Marzano (2003), p. 259. Wired Magazine (online version: http://www.wired.com/wired/) is a famous magazine which reports on how technology affects culture, politics and economy. It tends to creative techno-optimism—as may be guessed from this quotation from its founding father.

[9] Lessig (1999), p. 154. Cp. Brey (2006).

[10] Solove (2004).

[11] Sunstein (2001).

[12] This is the title of a book on non-human action that combines and critically assesses the work of continental and Anglo-American philosophers of technology: in particular Don Ihde and Bruno Latour. Verbeek (2005). Cf. Ihde (1990).

be held accountable by means of criminal liability, at first glance it still makes no sense to punish a machine or a software program. Used to think in terms of retribution and/or utility, one could conclude that retribution does not work because so far machines have no intentions, while the arguments based on utility will not provide reasons for punishment because so far machines do not learn as a consequence of punishment, nor will other machines be deterred by the punishment. In fact we cannot even begin to think what could count as punishing a machine (hitting or detaining it doesn't seem to make sense). The censure of punishment appeals to a reasonable human person, not to a mindless machine.

In a world partly ruled by autonomic computing, the theoretical framework underpinning individual criminal liability crumbles. The separation between subjects and objects is distorted, and the idea that our actions are informed by conscious reflection loses its hold.[13] There are three possible reactions to this situation. First, one can simply reject a world ruled in part by autonomic computing because it would imply the end of human civilisation. Second, one can simply deny that subjects and objects can no longer be separated, and invent complex reasonings to detect a human subject to whom one can attribute criminal liability somewhere in the chains of interconnected things and humans. Third, one can make a first attempt to rethink democracy and the rule of law in a way that allows for the fact that things have become 'actants' in ways that should be taken seriously.[14] In this article I will follow the third approach, without entirely discrediting the other two. If we imagine one were to live in a world entirely ruled by autonomic computing, I would agree that human society as we know it today would be eradicated and replaced by post-human intelligence. The scenarios for such a world have already been written, and one of the reasons to endorse the third approach is precisely to acknowledge the existence of such doom scenarios and to prevent their realisation.[15] On the other hand, if AmI turns out to be a bridge too far, a vision that is not realised in any substantial way, the second approach may suffice for the occasional strictly regulated occurrence of autonomic computing. In the present situation we cannot be sure what will happen; the only serious option is to speculate about potential implications and to invent ways of dealing with them. Adopting the third approach means that we want to retain the achievements of constitutional democracy, acknowledging the fact that it is a historical artefact that cannot be taken for granted. This is the default position from which we must undertake to rethink criminal liability, acknowledging that the transition from a written tradition to a tradition infused by advanced information and communication technologies (ICTs) calls for a radical rethinking of both democracy and the rule of law.[16] To this end we will leave for the moment the domain of legal philosophy, to assess some of the findings within the domain of philosophy of technology, which has developed interesting insights, well beyond scientistic techno-optimism and existentialist or phenomenological pessimism.

[13] This—in itself—is not a new situation. In the philosophy of mind swords have been crossed on this issue. Brain scientists have detected many 'autonomic' processes supposedly 'causing' our actions before conscious reflection comes in. On the complex relationship between conscious reflection and intentional action in the case of human subjects see Bayne (2006). Bayne critically discusses the conclusions drawn from the experiments by Libet, who is often claimed to have proven that brain states precede the conscious decision to act, supposedly meaning that they in fact 'cause' our actions. The fact that we experience our actions as the result of our own free will is then judged to be an illusion, see Haggard and Libet (2001). A more interesting position is developed by brain scientist and philosopher Varela in Varela et al. (1991).

[14] Cp. Bourcier (2001).

[15] Garreau (2005).

[16] About this transition Lévy (1990). About the possible consequences for the articulation of law, cp. Hildebrandt (2007, to be published).

In an attempt to side-step the Cartesian divide between active intentional subjects and passive material objects, the anthropologist of science and technology Bruno Latour speaks of human and non-human 'actants'.[17] This is not to claim that objects have intentions or that humans are determined by non-humans, but to provide a vocabulary that allows one to think in terms of the actions of both humans and non-humans. We can benefit from such a terminology in so far as it provides us with a generic concept of action that describes what humans and things actually do, not why they do it. To prevent confusion with the traditional philosophical concept of agent, he calls humans and non-humans actants instead of actors.[18] The question of what counts as an action is answered pragmatically rather than scientistically or moralistically since it concerns the impact of a certain behaviour rather than its cause or reason (anticipation of effects).

Following the cyberspace philosopher Pierre Lévy, we can define actants as 'everything that makes a difference within a network', while every actant 'defines itself by the difference it produces'.[19] One of the key messages of Lévy's work is an emphasis on the collective nature of the intelligence that emerges within the Internet: the intelligence does not reside in one personal computer, but depends on the interconnections. The same can be said for the Internet of Things as it turns the offline world online: 'smart things' depend on the network in which they participate to be able to create 'knowledge' and adapt. This raises the question of who is acting: one particular human or non-human, or the network as a whole. One could of course claim that this is not very new, because in cases of alternative or double causality similar questions have been raised. A person who was exposed to asbestos while also having smoked all his life may develop long cancer, and even though we can calculate statistically the chances of incurring lung cancer either by exposure to asbestos or by smoking, we cannot establish what actually caused this particular case of long cancer. This problem of liability—framed in terms of causality—is complex because we cannot close the gap between knowledge at the level of epidemiology and knowledge at the level of a particular instance. Lawyers have found interesting ways of attributing both causality and liability in such cases, usually located in a shift in the burden of proof, especially in cases of tort. One could also argue that a solution in all such cases is simple: as in the case of defective products or faulty services, one blames either the user (victim) or the person or corporation that designed it (product liability). But this response is inadequate, for two reasons. First, many of the problems that have been posed by less advanced technological artefacts and infrastructures, have been dealt with by means of a (re)distribution of risk. Such distribution confirms the fact that causality is a matter of *attribution* rather than the discovery of something 'out there' waiting to be discovered. Smart things seem to extrapolate the causal uncertainty to a point that requires a new awareness of the often rhetorical use of concepts like causality. Rather than merely confirming that smart technologies do not pose new questions, I would claim that the distributed intelligence required in AmI highlights the urgency of these questions. Second, however, the multi agent systems (MASs) we are discussing here may in fact raise new

[17] Latour (1999).

[18] In philosophy an agent could be defined as a subject capable of action. When acting the agent becomes an actor. The notion of action is traditionally connected with intention. This is, however, not necessary. One could discriminate between action as a generic notion and intentional action as a particular type of action. Within the field of computer science software programs or nodes in a network are often called agents. See footnote 19 below.

[19] Lévy (1990) (translation mh), p. 157.

questions.[20] Systems that learn in unpredictable ways and generate solutions that even computer engineers cannot explain seem to have acquired a new type of agency, not fully determined by the intentions of the designers. The effects they produce cannot be reduced to being unintended side-effects, caused by unexpected use or other factors. The emergent properties of MASs are generated by the technologies, indeed their unintended consequences were intended as such.[21] Even if tort liability could be extended to create a kind of strict liability for those that design, sell or use such technologies, criminal law would betray the presumption of innocence in attributing guilt without establishing who is to blame. The other option, leaving harm effected by distributed intelligence outside the scope of criminal law, would not solve but displace the problem: it could create a market for such technologies outside the censure of criminal law, leaving the regulation of such technologies to whoever can afford the risk of tort liability.

Returning to the question of who is acting, one particular human or machine and/or the network as a whole, we will thus follow Lévy by defining actants as those humans or non-humans ‘who make a difference’. This will allow us to regard both the network as a whole and its individual constituents as actants, depending on the level at which effects have been generated. It may be tempting to equate this criterion of agency with the usual attribution of causality (effective cause), but at the level of a multi-agent system (MAS) effects emerge in a way that cannot easily be defined in terms of causality. In the fields of brain science, artificial intelligence and philosophy of mind, the fact that networks consisting of individual nodes generate effects that depend on the connectivity between the nodes (emergence) is discussed as a paradigm for understanding intelligence (connectionism) and it seems that attributing causality to one of the nodes in such cases basically makes no sense.[22] In discussing the collective intelligence that emerges from advanced ICT applications, Lévy speaks of a cognitive ecology in which there are no causes and mechanical effects, but occasions and actors:

> Technological innovations make possible or condition the appearance of this or that cultural form (no modern science without printing, no personal computer without a microprocessor), but they do not necessarily determine them. It is a bit like in the domain of biology: a species does not derive from an environment. Obviously there would not have been any fish without water, but the sea was not obliged to create vertebrates, it could also have contained only algae and molluscs.[23]

Lévy’s position clearly avoids the mechanical causality that we tend to attribute to material things in sharp contrast with the mental freedom that we associate with individual human subjectivity. Instead of turning to causes and reasons, he directs our attention to what matters most—the difference a thing or non-thing makes within its environment; he argues that the link between the environment and an organism is neither determined nor

[20] Within computer science, a definition of a MAS could be: ‘An agent can be a physical or virtual entity that can act, perceive its environment (in a partial way) and communicate with others, is autonomous and has skills to achieve its goals and tendencies. It is in a multi-agent system (MAS) that contains an environment, objects and agents (the agents being the only ones to act), relations between all the entities, a set of operations that can be performed by the entities and the changes of the universe in time and due to these actions’, see Ferber (1999), as discussed in Rouchier (2001).

[21] Cf. the integration of computer science and sociological perspectives in: Meister et al. (2007).

[22] See again Varela et al. (1991) for a further elaboration of this position and of what they call ‘enaction’, being their paradigm for understanding the embodied mind.

[23] Ibid (translation mh), pp. 169–170. Cp. Friedrich Nietzsche in *Die Fröhliche Wissenschaft*, 1882, nr. 217: ‘Before the effect one believes in other causes than after the effect (translation mh)’.

*un*determined, but fundamentally *under*determined.[24] It may be the case that in rethinking criminal liability we need standards such as these to detect points of reference for the eventual attribution of criminal responsibility. Shifting from causes and reasons to 'making a difference' can highlight the anticipation of the consequences of one's actions as the most adequate criterion for criminal liability. An exploration of whether and how such attribution of criminal liability could be conceptualised, especially in the case of non-human intelligence, will be launched in Section "Attribution of Criminal Liability in the case of Networked Human-Non-human Actions". Before moving into the debate on what non-human action could mean for democracy and criminal liability, in Section "Democratic Criminal Liability and the Right to Violate the Criminal Law", I will explore the way technologies regulate our lives in the sense that they induce or enforce and inhibit or rule out certain choices of actions.

## Technological normativity

Legal normativity is often conceived of as a set of prescriptions and prohibitions (primary rules), combined with rules of competence (secondary rules about who can recognize, change or adjudicate the content of such primary rules).[25] In a constitutional democracy legal normativity is embedded in democratic procedures via the legislator, and constitutes a set of checks and balances that complements democratic enactment with case-to-case interpretation in court, which allows law to be attuned to actual legitimate expectations within civil society. Inspired by Searle's speech act theory, legal theory distinguishes constitutive from regulative rules; regulative rules concern a given behaviour that exists independently from the rule, constitutive rules concern a state of affairs that is called into existence if the rule is followed.[26] If I do not fasten my seat-belt I may violate a legal norm, and criminal liability may be attributed for this, but I can still drive my car. If, however, I do not register my 'marriage' with the civil registry, this does not result in criminal liability, but I will simply not be married. If one violates a regulative rule, this does not make it impossible to behave in the way that is prohibited; if one violates a constitutive rule the result one aimed for cannot be achieved. The distinction is relevant because criminal liability is at stake mainly in the case of violating regulative rules.[27]

Technological normativity does not consist in a set of prescriptions and prohibitions, combined with rules of competence, issued by a democratic legislator or established in a court of law. Technological normativity is not issued, enacted or established by legal authorities that are constrained by democratic accountability. This in itself is enough to raise one's eyebrows at the thought that there can be such a thing as technological normativity. What could normativity be, other than the deliberate regulation of human

[24] It may be interesting to discriminate—with Lévy—between types of action that merely consist of the realisation of a predefined possibility (implying mechanical application) and those that involve the actualisation of an underdetemined possibility (implying a measure of creativity and unpredictability). See Lévy (1998).

[25] Hart (1994/1961). Rules of competence are ultimately a matter of practice, an example of legal authorities declaring themselves legal authorities (Münchhausen is mythical but very real). Cp. Hart's ultimate rule of recognition.

[26] Searle (1969), Mittag (2006).

[27] One could think of a situation in which the violation of a constitutive rule results in damage because the objective of the rule is not achieved. If 'causing' such damage is criminalised the violation of this constitutive rule does lead to criminal liability.

interaction? If the reader can exercise some charity in following the argument, I will explain the advantages of a generic concept of normativity, covering both legal and technological normativity.[28] Instead of reserving the concept of normativity for more or less explicit rules that guide our interactions, I suggest a generic concept of normativity defined as 'the way humans or non-humans constrain human and non-human interaction'. This will allow us be explicit about the fact that and the way in which technological devices and infrastructures constrain our actions by inducing or enforcing and inhibiting or ruling out certain types of behaviour. Without a generic conception of normativity the similarity between technological and legal normativity may be invisible and we may fail to acknowledge the implications of the introduction of new technologies for our freedom to act in one way or another. By developing such a generic notion we may also become more precise in discriminating technological from legal normativity. For this reason I argue that anybody or anything that constrains our actions has a normative impact on our behaviour, because certain choices of actions are restricted or created. In fact, like legal normativity, technological normativity can be regulative or constitutive; it can regulate existing behaviour or it can be a precondition of certain behaviour. For instance, my AmI car may check my alcohol intake when I get into the driver's seat. Whenever I have overstepped the maximum allowed for driving, it may flash an irritating red light on the dash board as long as I drive. This technology is regulative of driving under the influence of alcohol, because it inhibits driving under the influence of alcohol. If the smart car is designed in a way that rules out driving whenever I have had too much alcohol, the technology has become constitutive of driving the car: I cannot drive it unless I abstain from a certain intake of alcohol. It is constitutive because driving is now dependent on me following a certain rule of behaviour, just like the constitutive rules for marriage.[29] Interestingly, a legal norm that prohibits driving after a specified alcohol intake will always be regulative. This is because actually driving a car cannot depend on compliance with a legal rule. This means that technological normativity can achieve a measure of compliance with certain rules of behaviour that is not within the reach of legal normativity.[30] It also means that technological normativity could be used to rule out certain types of criminal liability, because the actions they concern are no longer possible: the inability to drive under the influence of alcohol is just one example of such removal of the possibility of criminal conduct.

## Technological Normativity and Criminal Liability

Two types of implications derive from the analysis of technological normativity. First, the use of technological normativity may allow us to deprive people of the possibility of certain criminal conduct. Second, the emergence of an intelligent environment will cause severe problems for the attribution of criminal liability to a person, since the criminal action may have surfaced within an imbroglio of humans and non-humans.

[28] I am appealing to Davidson's principle of charity (or rational accommodation), claiming that the argument is coherent and refers to a reality that may overtake us sooner rather than later.

[29] Which could be entered into in earlier times without complying with such rules.

[30] Actually, it would be more precise to note that contemporary legal normativity, articulated in the technology of the script, cannot enforce compliance in the way that some technologies could. For an analysis of the transition from oral law to written law cp. Hildebrandt (2007).

## Democratic Criminal Liability and the Right to Violate the Criminal Law

### *Code as Law*

In 1999 Lawrence Lessig published *Code and Other Laws of Cyberspace*. In this book he defends the proposition that computer code constrains the interactions on the Internet in a constitutive way: creating a new space for interactions the regulation of which depends on the architecture of the space. He convincingly demonstrates that the present architecture of the Internet *makes a difference*: it could have been designed in other ways which would have created another space with other possibilities for regulating the behaviour of its users.[31] His conclusion is that computer code is a kind of law and should be used in a more deliberate way to create the kind of online environment that we should want. His *Code as Law* movement has gained momentum as many lawyers and social engineers recognise the normative impact of ICT, arguing that in a democracy we should use the most effective instruments to reach the goals we agree on. If the market is a more effective instrument we should prefer it to legal regulation, and in so far as computer code is even more effective this should do the job.

### *Democracy and the Right to Violate the Criminal Law*

Meanwhile the first criticisms have been published.[32] The main critique consists of two different arguments: first it is seen as a problem that technological normativity—as distinct from legal normativity—lacks democratic legitimacy; second, it is claimed that technological normativity is inherently deterministic, thus curtailing the freedom provided by legal regulation. The two arguments can stand on their own feet but their interrelation forms a third argument against the use of *code as law*: if technology can indeed enforce compliance with its norms and is developed and introduced without public consultation of those who will suffer or enjoy its consequences, then the use of technology contradicts the basic tenets of constitutional democracy.[33]

The first argument seems to take for granted that technological normativity does not fall within the scope of democratic regulation. The argument can be read in three ways: first, it can indicate that at this moment the development and introduction of new technologies take place in the sphere of private enterprise and independent scientific research, neither of which endorse democratic procedures to decide whether, which and how technologies should be developed; second, it can imply that the development and introduction of new technologies cannot and/or should not be brought under the regime of democratic decision making processes, but must be left to scientists and commercial enterprise; third, the argument can be turned round and used to plead for integration of technological innovation into the workings of constitutional democracy, to ensure that those who will suffer or enjoy the consequences at least have a say in the matter. The first way of reading the argument is rather matter of fact and does not—in itself—argue

[31] Lessig (1999), chapters 3 and 4. For instance, the anonymity and subsequent lack of accountability of behaviour on the Internet is a consequence of its design.

[32] Tien (2004), Brownsword (2005).

[33] The idea that democracy is at stake wherever people suffer or enjoy the indirect consequences of an action can be found in Dewey (1927).

against using *code as law*, though it may call for acceptance on the condition that it has survived democratic contestation. On the second way of reading the argument, it claims either that scientific research and commercial enterprise *should* not be regulated and that emerging technologies do not belong in the domain of public decision-making, or that they *cannot* be regulated because they escape human determination. I would agree that scientific research and commercial enterprise escape human determination in a strong sense, but would add that both scientific research and commercial enterprise depend on the legally constituted political framework that provides for their relative autonomy. A balanced relative autonomy does not preclude democratic participation in the decision-making processes regarding emerging technologies. On the third way of reading the argument, it calls for action, requiring the reinvention of democracy with regard to technological innovation. This seems the best way to read the argument, taking into account the state of the art and the need for relative autonomy in scientific research and commercial enterprise.[34]

The second argument seems to take for granted that technological normativity determines our behaviour instead of allowing us to resist its normative impact. This is not a fact: no such general statement about the nature of Technological normativity can be made, since its nature and implications depend on the particular technology in its particular environment.[35] As indicated above, a technology can be regulative of existing behaviour, or constitutive of a specific behaviour. In many cases this depends on the design of the technology (as in the case of the smart car that monitors your alcohol intake), but it may also depend on the way users have integrated the technology into their lives (since the mobile phone became the widespread artefact it is today, it has become constitutive of specific types of communication not even conceived of by its makers).[36] This argument can be read as pointing to an undesirable infringement of human freedom, to the extent that a particular technology does indeed enforce or preclude certain behaviour. However, if such behaviour is considered criminal, one could plead that it creates freedom for the potential victim as long as it only prohibits what we should not want to do anyway.[37] This raises the interesting question whether we have a (moral) right to violate the criminal law. Should I be granted the choice to drive a car while under the influence of alcohol, thus increasing the risk of killing others on the way? Should I be granted the choice to drive in the wrong direction, thus increasing the risk of collision? Is the function of the criminal law both to prevent crime and to keep open the possibility of committing crime, for instance because this provides humans with the opportunity to practice their moral judgement? I do not think the criminal law is intended to supply openings to practice one's moral judgement, because we cannot imagine explaining this to the victim. The criminal law censures wrongful actions that cannot be otherwise prevented, and one of the reasons why such actions cannot always be prevented is that it is not always clear in advance when something will count as a wrongful action. Criminal liability is established by a court of law, because we cannot presume a person to be guilty until we have assessed the evidence and listened to the story of the defendant who may plead that what he did should not count as the particular criminal

[34] For an out-of-the-box way of rethinking democracy cp. Latour (2004).

[35] Verbeek (2005), pp. 6–9.

[36] 'Mobile phone operators report their biggest profits from the runaway success of short text messages', Damina Mycroft, 'Intrinsic and Extrinsic Intelligence', Aarts and Marzano (2003), p. 256.

[37] At least, this is what Montesquieu suggested in his *The Spirits of the Laws*, XI, 3, Paris 1748: 'liberty can consist only in the power of doing what we ought to will, and in not being constrained to do what we ought not to will'. Translation available at: http://www.agh-attorneys.com/4_charles_montesquieu_SOTL.htm

act with which he was charged.[38] Constitutive technological normativity thus seems to contradict the presumption of innocence: the machine or the software program just rules out certain behaviour, without providing an opportunity to contest the behaviour's wrongfulness. The machine is not interested in right or wrong, it just precludes certain actions, because someone programmed it that way, perhaps because he was legally obligated to do so.

This brings in the third argument: if the introduction of a specific technology has not been open to democratic contestation, especially by those who will suffer the consequences, then the use of technological normativity circumvents the checks and balances that are characteristic of legal normativity in a constitutional democracy. This implies a democratic deficit to add to the deficit detected above, concerning the absence of contestation in a court of law. At this point we need to discriminate between two arguments. First, one may reject the use of constitutive technological normativity to enforce compliance with a legal norm that has been established in a democratic way, because it is not possible to contest that one's behaviour would have counted as criminal. This is an argument related to the rule of law, which depends on the judgement of an independent court. Second, one may reject the use of constitutive technological normativity per se, if the choice of this technological enforcement mechanism has not been made in a democratic way. This is an argument related to democracy.

As a conclusion we would claim that while computer code generates a kind of normativity similar to law, it lacks—precisely because it is NOT law—both democratic legitimacy and the possibility of contesting its application in a court of law. This is a major deficit in the relationship between law, technology and democracy.[39]

## *A Relational Theory of Law and a Pluralist Conception of Technology*

The idea that technology is never neutral does not imply that it is either good or bad; it merely suggests that it will always have a normative impact on our lives because it will constrain our actions in one way or the other. The point is that the moral evaluation of a particular normative impact cannot be made in advance or out of context: it will depend on the script that is inscribed into the design of the technology and on the way it has been integrated into the lifeworld of humans and non-humans.[40]

The idea that 'technology is neither good nor bad, but never neutral'[41] is part of a pluralist conception of technology.[42] Instead of endorsing a substantive conception (which takes Technology to be either the best thing that ever happened to us, or the end of human

[38] An example of such a plea is the—unsuccessful—plea of a husband who claimed that his raping his wife *did not count as* rape under the common law. See (1995) 21 EHRR 363, wherein the Court ruled that in this case the husband having sex with his wife in fact *counted as* rape.

[39] To remedy this deficit we may have to learn to articulate some of the protective aspects of law in the technologies we aim to regulate, cf. Hildebrandt (2008, to be published).

[40] About the notion of a script, see Latour (1993). About the fact that the actual normative impact is constituted by the way humans actually attach themselves to a technology, see Verbeek (2005), p. 217 on the multistability of technological mediation and the interpretive flexibility this provides. The fact that a moral evaluation cannot be made in advance does not imply that we can sit back until the consequences are in force. It rather calls for speculative exploration of such consequences and should involve some kind of democratic participation of those who may suffer or enjoy them.

[41] Kranzberg (1986).

[42] Verbeek (2005), p. 11.

civilisation), or an instrumentalist conception (which takes technologies to be neutral tools to implement goals), we will argue for a pluralist conception that takes technological instrumentality seriously without taking its moral status for granted. A pluralist conception of technology avoids thinking in terms of Technology as some continental philosophers tended to do;[43] it rather invites empirical speculation about the normative implications of specific technologies.[44] At the same time it avoids seeing technologies as nothing but tools to reach certain goals, in which case the only criterion for evaluation would be the effectiveness and efficiency of a specific tool to achieve a particular objective. AmI enabling technologies may be used to enforce compliance with legal norms by rendering certain criminal actions impossible, but to evaluate whether this is a good thing we must look into the question of whether we lose some of the fundamental principles embedded in the criminal law on the way.

This brings us to the question of what should inform our choice between legal and technological instruments. Should we take a substantive view of legal normativity and proclaim that legal instruments are the only way to achieve democratically established policy goals, rejecting the use of technological devices or infrastructures, because legal instruments are always better, having an intrinsic value not inherent in technological instruments? Or, should we take an instrumentalist view of legal normativity and admit that sometimes technological tools are more effective and efficient means to the objectives established by our democratic legislature? As in philosophy of technology, legal philosophy has different conceptions of law and legal normativity. These conceptions are not equivalent with those of technological normativity, but they reveal interesting similarities. To demonstrate the relevance of the discussion I will focus on the conception of criminal law. First we have a critical conception of criminal law, which regards law as an instrument to limit the power of government, not as an instrument to achieve policy objectives. The idea is that governments will punish anyway, they do not need the law for that. It is only under the rule of law that punishment is brought under the regime of a criminal law that stipulates a number of restrictions on criminal liability and its establishment, notably requiring the verdict of an independent and impartial judge. This conception of the criminal law compares well with the substantivist understanding of technology, which regards technology as an autonomous force that is considered either beneficial or detrimental to human society. In both cases the value of Law or Technology is taken for granted, since it does not depend on the particular law or legal framework, technological device or technological infrastructure that is to be evaluated. Second, we have an instrumentalist conception of criminal law, which regards law as an instrument to achieve the policy objectives of the legislator. The idea is that law can be an effective and efficient means to certain states of affairs, especially since the criminal sanction can induce compliant behaviour. In this case legal instruments are taken to be neutral with regard to the purpose they serve. This conception seems complementary to the instrumentalist conception of technology, implying that legal and technological tools can be freely exchanged, depending on the efficacy and efficiency of the one or the other.

[43] Ibid. Verbeek provides a critical reconstruction of the work of Jaspers, Marcuse, and Heidegger before moving into Ihde, Latour and Winner.

[44] Empirical speculation involves actual knowledge of a specific emerging technology and of the environment into which it may be introduced, it involves educated guessing and trained intuition instead of both general deductive reasonings and reductive quantitative social research. Cp. the way Isabelle Stengers uses the term speculation in: Stengers (2000).

Third, we will now introduce a relational conception of law and legal normativity that compares well to a pluralist conception of technology, though—other than this—it is implied in the legal framework that constitutes constitutional democracy, thus implying a conception that is both normative and descriptive.[45] In a relational conception of law the critical function of legal normativity is integrated with its instrumental functions: the criminal law constitutes the competence to establish criminal liability, thus being instrumental to preventing criminal conduct and censuring its occurrence;[46] its critical function resides in the way it constitutes this competence, for instance in the fact that criminal liability needs to be established in a court of law to allow contestation of the validity of the allegedly violated legal norm and/or the violation itself. The instrumental and critical functions of the criminal law can thus be distinguished but should not be separated, as this would allow trade offs that risk the effectiveness of either aspect of the law. This relational conception of law does not operate in a politically neutral manner: it does not see law as an instrument to be used at will by any type of government. In this sense it is not merely a descriptive conception, as a pluralist conception of technology seems to be, even if it is based on the fact that technology is never neutral. The non-neutrality of the latter concerns its normative impact, the non-neutrality of the former concerns its positive entanglement with a specific type of society (constitutional democracy). However, the question whether a particular criminal law does integrate its instrumental and protective aspects cannot be answered in the abstract; it demands investigation. This also implies that the question of whether an integration of legal and technological normativity, for instance by the use of a specific technological infrastructure to support the workings of the criminal law, cannot be answered in advance. It will need careful study of the normative implications of the technology, paying attention to the integration of instrumental and protective aspects. If a technological infrastructure rules out behaviour we consider criminal this may imply that one can no longer contest the relevant legal norm in court. In the case of the prevention of traffic accidents by means of 'smart' cars I think that we should be willing to live with the fact of 'missing' the chance to drive well above the speed limit, or the chance to cross a road without checking for oncoming traffic. However, autonomic computing could also be used to preclude certain actions of specific individuals who are calculated to be more prone to negligent or wrongful behaviours. In that case I have serious doubts as to whether such autonomically implemented actuarial justice can still count as criminal law in a constitutional democracy.

## Attribution of Criminal Liability in the Case of Networked Human-Non-human Actions

Imagine that my fridge tells my toaster that it has run out of bread, after which the toaster tells the fridge that I will need a piece of toasted bread at 7 am tomorrow morning, leading to an order being placed with a supermarket that will deliver my preferred type of bread in time for the best price. This does not necessarily constitute an action. It does with regard to the legal consequence produced by the contract that is implied in the transaction, but it

[45] This conception of law has been developed by Foqué and 't Hart (1990). See also Hildebrandt (2006a).

[46] This instrumentality differs from instrumentalism the way the pragmatist consequentialism of Peirce and Dewey differs from a utilitarian pragmatism in which means and ends are separated. In that sense prevention regards the legitimate anticipation of the consequences of one's actions rather than an amoral calculation of costs and benefits. This means that we agree with Duff's rejection of utilitarian consequentialism (Duff 2001, pp. 3–19), without reverting to a Kantian position in which only intentions count.

does not with regard to the fact that my neighbour has been found dead the following morning. However, it may still turn into an action in that sense if there is a relationship between the delivery of the bread and the death of my neighbour. If, for unknown reasons, the bread was delivered to my neighbour's house around the time of his death, the chain of interactions may become relevant and may qualify as an action. The fact that the adaptations of my smart house are not even *inter*active but mainly *pro*active—in order not to disturb me with the trivialities of the household—could then raise cumbersome questions as to 'who did it'.[47]

The bread that I ordered may have contained an ingredient to which my neighbour was fatally allergic. As his own smart home was in the habit of ordering his own bread supply (which evidently did not contain such ingredients) around the same time as the delivery of my loaf, my neighbour trusted this to be his and—having worked until 3 am on a paper on moral philosophy—sat down to have some. One can object that this case does not pose any new questions; the same could have happened if I checked my fridge, ordered bread by phone and went off to sleep. However, in that case the 'fault' could more easily be traced: either I gave a wrong house number or the supermarket made a mistake. If I knew about my neighbour's predicament, I might be charged with culpable homicide (or if there is motive, with murder). What if I ordered via the Internet and the web server had a bug, causing a mix up? What if my fridge communicates with other fridges and actually knows that my neighbour is allergic to my bread, and—when registering the delivery of the bread with the neighbour—should have warned someone or something about the imminent danger?

What if my fridge should have warned someone or something? The process of machine to machine communication (M2M) could have resulted in the ringing of an alarm, or in automatic closure of the fridge in which the bread was put. How can we detect what or who made a difference here, if the intelligence of the environment in preventing the death of my neighbour arises from the interactions between machines, not even from one particular machine? How to attribute criminal liability in the case of distributed intelligence?

Referring to Latour, Felix Stalder writes: 'If action, distributed along a chain of humans and non-humans, is not qualified more richly than as effecting something somewhere, then intentionality, accountability, and responsibility for this action are equally distributed along the same chain. (…) Conceptualizing agency as a distributed effect is a very powerful analytical strategy but politically difficult because of the immanent danger of equalizing humans and machines to the point where responsibility and accountability for action vanishes'.[48] The interesting point made by Lévy above was that an action is not defined by its causes or reasons, but by the difference it makes. In reply to Stalder this means that not just anything that is effected anywhere should always counts as an action. Whether something counts as an action depends on whether it makes a difference that is relevant to the issue at stake, in this case the death of my neighbour. This may be a positive action, leading to the bread being delivered to the wrong address, or a negative action, consisting of the wrongful failure to prevent the neighbour from eating the fatal bread. Evidently more than one action can constitute criminal liability for the death, meaning that more than one 'node' in the network of distributed intelligence can be held liable for the

[47] Interactive computing depends on deliberate human-computer interaction, proactive computing does not. It builds on the real time monitoring and machine anticipation of inferred human preferences. Cp. Tennenhouse (2000).

[48] Stalder (2000–2003).

same event.[49] The most pertinent issues derive from the fact that we may have to denote the network itself as the author of the incriminated action, because the nodes had no control over the chain of events that brought about the relevant effect. This raises two questions: (1) does proactive computing excuse human nodes from a failure to prevent violation of the criminal law, because they are not aware of the effects the network may produce; and (2) in which case could a network constitute the kind of subject that can be called to account in a court of law and censured if found guilty? It seems to me that these questions are interrelated: if we cannot attribute criminal liability for wrongful actions because the responsibility is diffused beyond measure, we should think twice before introducing the technological infrastructure that enables such unaccountable consequences.

One solution is the introduction of strict liability, but while this may be an adequate solution for tort or breach of contract, it seems inapt for criminal liability. If we cannot establish criminal guilt—consisting of wrongful and culpable behaviour—for any specific human or non-human component of a network, it does not make sense to censure a component for what has been brought about by the joint action of the network. Theories on collective moral responsibility that build on individual moral responsibility as the 'natural' type of responsibility seem to miss the point here,[50] as distributed intelligence precisely consists of an intelligence that emerges within the network without being traceable to one or more of its constituent parts. A smart fridge is not smart in itself; it becomes smart because of the interconnections within the network it co-constitutes. In fact, seeking to 'blame' one of the nodes would be like seeking the particular part of the human brain that 'caused' an incriminated action, instead of recognising the emergence of individual agency in a succession of brain states that somehow constitute a specific action.[51] This calls for reflection on the possibility of censuring the subject that emerges from distributed non-human intelligence, boiling down to the question of whether punishment would then make sense. If we take punishment in a constitutional democracy to have a communicative purpose, rather than a strict deontological objective, a utilitarian goal or an expressive function, two conditions must be met in order to subject an individual to punishment. First, the subject must be sensitive to censure and second, the subject must be capable of standing trial. In both cases it must be presumed that the subject of punishment can be involved in what Antony Duff calls '*reciprocal* and *rational* engagement'.[52] If a subject is not sensitive to censure it cannot learn from being punished; if a subject cannot take the stand in a court of law it cannot contest the incrimination, which would turn the punishment into discipline. Within the scope of this article I cannot elaborate on the point at which a networked environment would acquire the kind of agency presumed by the criminal law. It seems to me that artificial intelligence in itself does not qualify as such agency, even if some kind of consciousness

[49] A common definition of a node can be found at http://en.wikipedia.org/wiki/Node_(IT): A node is a device that is connected as part of a computer network. (…) Nodes can be computers, personal digital assistants (PDAs), cell phones, or various other network appliances, (…). We are using the term in a more generic way, integrating social network theory with network theory; a node in our case is the nexus of different lines of communication in a networked environment, meaning that nodes can also be human individuals.

[50] Miller (2006).

[51] This brings us well into the nexus of the philosophy of mind and the neurosciences, cp. Varela et al. (1991).

[52] Duff (2001), p. 79.

would emerge.[53] Animals have consciousness but we do not consider them fit to be subjected to legal punishment, because we have no indication that they can reflect on their actions as their own actions. Their consciousness is an awareness of the environment, without the concomitant awareness of this awareness which is typical of the human sense of self. Helmuth Plessner actually took this to be the crucial difference between humans and non-human life forms: the self-consciousness of the human person creates a distance between the self, the world and the self itself, condemning humans to what he called indirect directness, natural artificiality and a utopian position.[54] To be sensitive to censure, rather than mere discipline, a subject needs to be conscious of its self, allowing the kind of reflection that can lead to contestation or repentance in the case of a criminal charge. It would be interesting to investigate to what extent non-human intelligent networks may develop such self-awareness, and we should not be surprised when even more complex forms of self-awareness develop, for instance a distributed awareness or a sense of being distributed in different locations.

## Conclusive Remarks

In this contribution I have explored some of the implications of non-human distributed intelligence for democracy and criminal liability. To challenge conventional notions of human agency I have presented the vision of AmI and the implied processes of autonomic computing, claiming the need for alternative conceptions of the relationship between law and technology if we want to maintain a viable constitutional democracy. After this I made a modest attempt to explore the implications of non-human distributed intelligence for the attribution of criminal liability.

## References

Aarts, E., & Marzano, S. (Eds.) (2003). *The new everyday. Views on ambient intelligence*. Rotterdam: 010 Publishers.

Bayne, T. (2006). Phenomenology and the feeling of doing: Wegner on the conscious will. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior? An investigation of the nature of volition*. Cambridge, MA: MIT Press.

Bourcier, D. (2001). De l'intelligence artificielle à la *personne virtuelle*: émergence d'une entité juridique? *Droit et Socie´te´, 49*, 847–871.

Brey, P. (2006). Freedom and privacy in ambient intelligence. *Ethics and Information Technology, 7*, 157–166.

Brownsword, R. (2005). Code, control, and choice: Why East is East and West is West. *Legal Studies, 25*(1), 1–22.

Communication of the European Commission to the European Parliament (2007). *RFID in Europe: Steps towards a policy framework*, COM2007(96) final of March 2007.

Custers, B. (2004). *The power of knowledge. Ethical, legal, and technological aspects of data mining and group profiling in epidemiology*. Nijmegen: Wolf Legal Publishers.

Dewey, J. (1927). *The public & its problems*. Chicago: The Swallow Press.

Dreyfus, H. (1972). *What computers can't do*. New York: Harper and Row.

[53] In the community of researchers on Artificial Intelligence (AI) and Computational Intelligence (CI) claims concerning the emergence of artificial consciousness are paramount. CI is defined as involving 'iterative development or learning (…). Learning is based on empirical data and is associated with non-symbolic AI, scruffy AI and soft computing. Methods mainly include: neural networks, fuzzy systems and evolutionary computations', see http://en.wikipedia.org/wiki/Artificial_intelligence.

[54] Plessner (1975).

Duff, R. A. (2001). *Punishment, communication, and community*. Oxford: Oxford University Press.

European Policy Outlook RFID (2007). Draft version, the working document for the expert conference on 'RFID: Towards the Internet of Things' in June 2007 in Berlin.

Ferber, J. (1999). Multi-agent system: An introduction to distributed artificial intelligence. Harlow: Addison Wesley Longman.

Foqué, R., & 't Hart, A. C. (1990). *Instrumentaliteit en rechtsbescherming*. Arnhem: Gouda Quint.

Garreau, J. (2005). *Radical evolution. The promise and peril of enhancing our minds, our bodies—and what it means to be human*. New York: Doubleday.

Haggard, P., & Libet, B. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies, 8*(11), 47–63.

Hart, H. L. A. (1994/1961). *The concept of law*. Oxford: Clarendon Press.

Hildebrandt, M. (2006a). Trial and 'fair trial': From peer to subject to citizen. In A. Duff, L. Farmer, S. Marshall, & V. Tadros (Eds.), *The trial on trial. Judgment and calling to account* (Vol. 2, pp. 15–37). Oxford: Hart.

Hildebrandt, M. (2006b). From data to knowledge: The challenges of a crucial technology. In *DuD—Datenschutz und Datensicherheit* (Vol. 30, No. 9, pp. 548–552).

Hildebrandt, M. (2007). Technology and the end of law. In E. Claes & B. Keirsbilck (Eds.), *The limits of (the rule of) law*. Oxford: Hart (to be published).

Hildebrandt, M. (2008). A vision of ambient law. In R. Brownsword & K. Yeung (Eds.), *Regulating technologies*. Oxford: Hart (to be published).

Ihde, D. (1990). *Technology and the lifeworld. From garden to earth*. Bloomingron: Indiana University Press.

ISTAG (2001). *Scenarios for ambient intelligence in 2010*. Information Society Technology Advisory Group, available at: http://www.cordis.lu/ist/istag-reports.htm.

ITU (2005). *The internet of things*. Geneva: International Telecommunications Union (ITU).

Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer, 36*, 41–50.

Kranzberg, M. (1986). Technology and history: 'Kranzberg's laws'. *Technology and Culture, 27*, 544–560.

Latour, B. (1993). *La Clef de Berlin et autres leçons d'un amateur de sciences*. Paris: La Découverte.

Latour, B. (1999). *Pandora's hope. Essays on the reality of science studies*. Cambridge: Harvard University.

Latour, B. (2004). *Politics of nature. How to bring the sciences into democracy* (trans. by Catherine Porter). Harvard University Press.

Lessig, L. (1999). *Code and other laws of cyberspace*. New York: Basic Books.

Lévy, P. (1990). *Les technologies de l'intelligence. L'avenir de la pensée à l'ère informatique*. Paris: La Découverte.

Lévy, P. (1998). *Becoming virtual. Reality in the digital age*. New York: Plenum Trade.

Meister, M., Schröter, K., et al. (2007). Construction and evaluation of social agents in hybrid settings: Approach and experimental results of the INKA project. *Journal of Artificial Societies and Social Simulation, 10*(1), http://www.jasss.soc.surrey.ac.uk/10/1/4.html.

Miller, S. (2006). Collective moral responsibility: An individualist approach. *Midwest Studies in Philosophy, 30*(1), 176–194.

Mittag, M. (2006). A legal theoretical approach to criminal procedure law: The structure of rules in the German code of criminal procedure. *German Law Journal, 7*(8), 637–646.

Plessner, H. (1975). *Die Stufen des Organischen under der Mensch. Einleitung in die philosophische Anthropologie*. Frankfurt: Suhrkamp.

Rouchier, J. (2001). Review of Jacques Ferber's multi-agent system: An introduction to distributed artificial intelligence. *Journal of Artificial Societies and Social Simulation, 4*(2), http://www.jasss.soc.surrey.ac.uk/4/2/reviews/rouchier.html.

Searle, J. R. (1969). *Speech acts, an essay in the philosophy of language*. Cambridge: Cambridge University Press.

Solove, D. J. (2004). *The digital person. Technology and privacy in the information age*. New York: New York University Press.

Stalder, F. (2000–2003). Beyond constructivism: Towards a realistic realism. A review of Bruno Latour's Pandora's Hope. *The Information Society, 16*, 245–247.

Stengers, I. (2000). Another look: Relearning to laugh. *Hypatia, 15*(4), 41–54.

Sunstein, C. (2001). *Republic.com*. Princeton: Princeton University Press.

Tennenhouse, D. (2000). Proactive computing. *Communications of the ACM, 43*(5), 43–50.

Tien, L. (2004). Architectural regulation and the evolution of social norms. *International Journal of Communications Law & Policy* (9), Special Issue on Cybercrime.

Varela, F. J., Thompson, E., et al. (1991). *The embodied mind. Cognitive science and human experience*. Cambridge: MIT.

Verbeek, P.-P. (2005). *What things do. Philosophical reflections on technology, agency and design*. Pennsylvania State University Press.