Who Needs Stories if You Can Get the Data? ISPs in the Era of Big Number Crunching

Mireille Hildebrandt

Received: 15 March 2011 / Accepted: 8 July 2011 / Published online: 30 July 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract In this article, I will investigate to what extent democracy and the Rule of Law require that ISPs as ‘common carriers’ that provide ‘mere conduit’ pre-empt extensive monitoring of the content they carry. I will trace this duty as a moral duty that is bound up with the framework of constitutional democracy, arguing that such monitoring affords unprecedented data-mining operations that could stifle our account of ourselves as moral agents in the novel infosphere.

Keywords ISPs · Data mining · Privacy · Ontological friction · Ambient Law · The Rule of Law · Deep packet inspection · Data retention · Number crunching

1 Introduction

We can observe the proliferation of what is called ‘big data’ everywhere. From the life sciences to physics and astronomy to criminology, forensics, insurance, the protection of critical infrastructure to marketing and advertising, massive amounts of data are captured, stored and aggregated. This only makes sense (and money, obviously) to the extent that we have the techniques to squeeze meaning out of these numbers. De-contextualized as they are, these data require re-contextualization to make them work for us. And their number implies that this can only be achieved by means of data-mining technologies with enormous

---

M. Hildebrandt
Institute of Computer and Information Sciences (ICIS), Radboud University Nijmegen,
Nijmegen, the Netherlands

M. Hildebrandt
Law Science Technology and Society (LSTS), Vrije Universiteit Brussel, Brussels, Belgium

M. Hildebrandt (✉)
Department of Jurisprudence, Erasmus School of Law, Rotterdam, Netherlands
e-mail: hildebrandt@frg.eur.nl

computing power, way beyond the scope of the human mind. Finding patterns in databases is commonly called Knowledge Discovery in Databases (KDD), alluding to novel ways of abducting knowledge from aggregated machine-readable data. Suggesting that these computational patterns constitute knowledge implies a novel epistemology.

In this contribution, I will trace the ethical and legal consequences of this novel engagement with knowledge production, focusing on the position of Internet Service Providers (ISPs). What interests me here is the question to what extent ISPs should provide built-in protection against extensive monitoring of Internet conduct. In other work, we have coined the articulation of legal protection into the ICT infrastructure as Ambient Law (Hildebrandt and Koops 2010). Instead of inquiring whether ISPs have a moral duty to monitor content in order to fight child pornography, human trafficking, money laundering or other criminal actions, I will raise the opposite issue of whether we have a moral duty to impose a legal duty on ISPs to prevent systematic interception, storage, aggregation and analysis of data that pass through their infrastructure. I will not argue for a moral duty for individual ISPs, based on a utilitarian or deontological framework of moral philosophy, but for a requirement to build in such protection based on the framework of constitutional democracy. The moral duty that informs this requirement hinges on two interrelated moral values, namely privacy and transparency. The first is related to negative and positive freedom in the development of personal identity; the second relates to the transparency of those in authority (and those in power). Both depend on and constitute the framework of constitutional democracy: they are public goods that celebrate a relative and relational individual autonomy as something that is not given, cannot be taken for granted and requires a specific system of checks and balances.

I will start with a reflection on Floridi's notion of 're-ontologizing of the infosphere', to capture the enormity of the transformations that are taking place in our information and communication infrastructures. Next, I will trace the epistemological implications of this re-ontologization, focusing on the advance of big number crunching or knowledge discovery in databases, ending with a discussion of the 'inference problem'. This problem concerns the fact that the emerging infosphere seems capable of anticipating our behaviours, before we become aware of them. This raises a number of ethical issues about the extent to which we are being 'read' since such inferences could dissolve the 'ontological friction' that safeguards our privacy. Finally, I will discuss the role of ISPs as caretakers of the most extensive amount of data that can be mined and used to profile us. So far, most privacy activists have focused on the Big Data stored by commercial giants like Google or governmental agencies that require personal data to e.g. redistribute welfare or claim taxes. However, recently it has become technically feasible for ISPs to inspect the data they 'carry' in a systematic manner, enabling them to become involved in a number of schemes to make business out of data. This article will argue that constitutional democracy requires us to build effective and legitimate fences into the bottom line of the infrastructure that is under the care of the ISPs. This will be the only reliable way to prevent individual citizens from becoming singularly transparent to anybody who can afford to pay for these inferences.

2 Re-ontologizing of the Infosphere: *Umwelt*, *Welt* and What?

To develop an understanding of the role of ISPs in our socio-technical landscape, I think that Floridi's notion of a re-ontologization of the infosphere is on the spot (Floridi 2005). Before developing an ethics of number crunching and deriving a legal obligation for ISPs to prevent unlicensed profiling, we need to frame the transformation of our ICT infrastructure.

Floridi's main point is that the ongoing informatization of human society has not merely augmented or enhanced our agency but has radically changed the ontological structure of the infosphere. This implies that both our selves and our environment are in a process of radical alteration. Though he does not cite Negroponte's shift from atoms to bits (Negroponte 1996), he describes the radical changes in terms of a 'migration of humanity from its *Umwelt* to the infosphere itself' (Floridi 2010: 189). He describes this in function of a thorough digitization of information that amounts to a digitization of information entities that results in a homogenization of the processor and the processed. In turn, this constitutes the informationalization of interactions between different agents and the evolution of new informational agents (both artificial and hybrid).

Though I agree that we are in a state of radical transition of both ourselves and our environment, it is not clear how we should understand the move from an *Umwelt* to 'the infosphere itself'. Opposing the concept of infosphere to that of *Umwelt* suggests a move from the material to the mental or at least to the disembodied. In other work, Turilli and Floridi (2009: 108) have emphasized that information must not be confused with data:

A datum is something that 'makes a difference' and, as such, can be perceived, measured and captured via an interaction. (...) Information is produced through the elaboration of data. Semantic information can be thought of as the result of a set of operations performed by an agent taking raw data as input and producing well-formed, meaningful and truthful data (that is, information) as output (...). Semantic information is not the result of a 'snapshot' or passive observation, but depends on agents' proactive meaningful data elaborations (semanticisation)'.

Information, understood in this way, is a crucial sign of embodied agency. It indicates an entity capable of distinguishing the 'difference that makes a difference' (Bateson 1972). This is inherent in the concept of *Umwelt*, since this necessarily assumes an agent that fits its affordances (Gibson 1986). As such, an *Umwelt* assumes embodied informational entities capable of interacting on the basis of relevant data that are spotted and translated into what is 'actionable' for the agent (Varela et al. 1991). I would think that an *Umwelt* is an infosphere and to the extent that a human observer can speak in a meaningful way about the infosphere, this infosphere is an *Umwelt* for whichever agents are interacting, practicing their autonomy and adapting to relevant changes. Going online, entering cyberspace, playing around in virtual worlds or hanging out in 'the cloud' does not imply disembodiment. It does, of course, imply a translation of one's interactions into machine-readable discrete data, but I would argue that this is a re-embodiment

rather than a disembodiment (Ihde 2002). It is also—to some extent—a homogenization that allows processor and processed to interact seamlessly, but the discipline of human–machine-interfacing confirms that this requires a lot of added energy and complexity to ‘work’. Humans, dogs, plants and machines are embodied differently and this has consequences for the manner in which they constitute and participate in the infospheres (= *Umwelts*) that afford their interaction. I am using the plural here to highlight the fact that there is not one infosphere, despite the fact that the World-Wide Web seems to suggest an overarching global online infosphere. What makes this seemingly unified homogenized infosphere interesting is the ongoing de- and re-contextualization of digital data that it affords (Kallinikos 2006), thus indeed constituting a plurality of permanently reconstituted online *Umwelts*.

Instead of moving straight from *Umwelt* to infosphere, we may deepen our understanding by introducing the concept of *Welt*, defined as the semantic universe that builds on spoken and written human language. Ricoeur (1973) has explained how text allows us to distinguish between an ostensive reference (always dependent on the *Umwelt*) and the non-ostensive reference that depends on symbolic language (thus creating a *Welt*). He saliently highlights how precisely the materiality of written text affords the creation of a con-text beyond one’s *Umwelt*, even allowing for a measure of de- and re-contextualisation that is only possible when a text is read in a radically different context (in a different time and/or place). Ong (1982); Eisenstein (2005); Goody and Watt (1963); Ihde (1990) and many others have traced the *material* affordances of the script and the printing press as preconditional for the kind of world we now inhabit, thus calling attention to the fact that moving from an actual situation to its context to more abstract thought depends on the embodiment of human speech by means of stone engravings, clay inscriptions, paper writings and finally the movable type printing press. Higher levels of abstraction thus depend on the advent and proliferation of text. Taking this into account we can understand the re-ontologization of the infosphere as a move from the *Welts* of modern times to an infosphere that allows for a higher level of abstraction. I would not qualify this as ‘the infosphere itself’ since this could refer to the homogenization of an infosphere without interfaces, beyond the need to translate, enabling seamless interaction. Such homogenization would take us beyond any ontological friction, seemingly ruling out any obstruction that stands between us and ever more salient and more relevant information. Though the description of the re-ontologized infosphere could invite such an idealistic but unseemly interpretation, this cannot be Floridi’s own. Below I will discuss his concept of privacy as an ‘ontological friction’ that safeguards a person’s autonomy in terms of informational self-constitution, as well as several warnings against the threat of a unified language or indiscriminate transparency. The question remains how the migration from *Umwelt* to *Welt* continues in the age of ubiquitous computing, taking note of the fact that the emergence of a *Welt* has not replaced our situatedness in an *Umwelt* but rather extended and multiplied our situatedness.

In the next section, I will discuss the epistemological implications of big number crunching or knowledge discovery in databases as the most salient transformation of the socio-technical landscape.

3 Epistemological Implications of Knowledge Discovery in Databases

3.1 Mediated Perception and Narrative Cognition

Floridi (2010: 255) suggests that ‘we are, to the best of our knowledge, the only semantically structuring structures in the infosphere’.

As such, we give and make sense of what we experience by objectifying it. The distance from the world that makes a rich cognitive life possible is also the price imposed by the reification of the world.

This dovetails with the position of the human observer—already mentioned above—as preconditional for the attribution of meaning, which is made possible by the introduction of human language and its externalization and objectification in the script. It confirms Ricoeur’s description of the emergence of a *Welt* in the face of the distantiation (cf. Geisler 1985) that is inherent in human language and its sedimentation in the script. It clarifies how the concept of an infosphere and notions such as moral agency and moral patients depend on a third person perspective that allows us to talk about things, others and our selves, allowing us to not only reach high levels of abstraction *but to also tell stories about blame and responsibility*. Blaming another or calling another to account (Butler 2005) implies a distantiation, a view from elsewhere, an eye for what could have been different. It disrupts the immediacy of a conscious awareness that is not also conscious of its own consciousness (Plessner 1975; Cheung 2006), setting ‘us’ apart from a number of other moral agents—thus actually creating the possibility of describing other information entities as moral patients or moral agents. The distantiation promotes a perception mediated by language and a narrative cognition.

Ricoeur (1992) has linked this distantiation to our narrative self-identity. Our need to tell stories about the events in our lives that make a difference is typical for our way of knowing both the world and our selves. There is a subtle tension between ambiguity and certainty in a story; it challenges us to imagine other courses of action but it also returns us to the legitimate expectations that constrain our actions. This tension seems to be constitutive for our personal identity, challenging us to reinvent both our selves and the world while still sustaining continuity with former ‘selves’ and the webs of meaning that create common ground with our fellows. One might be tempted to speak of a balance instead of a tension but this assumes too much; the struggle between disruptive reinvention and submissive continuity requires a persistent effort, and its outcome cannot be taken for granted.

The question I want to raise is how data-mining operations on a grand scale—i.e. super crunching (Ayres 2007)—will impact our mediated perception and narrative cognition. What does it mean that ‘knowledge’ is mined from massive amounts of data? If the distantiation inherent in human language affords the contestation of knowledge claims, how will knowledge that is achieved by means of data science (KDD, machine learning, neural nets or multi-agent systems) afford its own scrutiny? Is a similar critical distantiation inherent in the statistical calculations made possible by a computing power that is not available to the human mind?

3.2 Number Crunching: KDD

KDD is a process of inferring or abducting knowledge from aggregated data. It is also called pattern recognition, because it detects patterns between data in a database. It depends on five subsequent steps that are reiterative, thus creating a permanent feed-back and allowing a persistent fine-tuning of the patterns found. The first step concerns translating the flux of real-time events into discrete machine-readable data. There is an element of objectification and reification here, highlighting our inability to make sense of the world in flux. As Floridi (2010: 255) mentions when discussing our need for semantic structuring:

We freeze changes into state- or phase-transitions and modular events and transform patterns and structures into objects and properties, finally privileging a naive ontology of sufficiently permanent things and qualities.

However, this step is not made by us. It is implemented by computing systems, translating events into bits that are without meaning until they are translated again into stuff we can 'read'. The second step consists of the aggregation of the data, making it possible to 'mine' them in search of relevant patterns. This is done during the third step, which is data analysis or data mining (Fayyad et al. 1996), using algorithms, heuristics, neural nets or other computational techniques to find clusters, association rules or other correlations between the data. When speaking of KDD, we must take into account that the amount of data that can be 'mined' this way is incredibly high: Ayres (2007: 11) speaks of datasets 'measured not in mega- or gigabytes but in tera- and even petabytes (1,000 terabytes)'. To detect patterns in such datasets is not possible for the naked human eye. Especially when using unsupervised learning techniques or bottom-up algorithms the patterns found are not so much the confirmation of a correlation that was first hypothesized. These patterns are novel hypothesis, 'discovered' by the computational techniques used to find structures in the dataset. This is why the subtitle of Ayres book on *Super Crunching* is: *why thinking-by-numbers is the new way to be smart*. Anderson (2008) actually speaks of 'The End of Theory', announcing that this type of knowledge construction will make 'the Scientific Method obsolete'. Though such exuberant enthusiasm can easily be criticized for being naive and even dangerous, I think we cannot deny that data mining provides a new type of knowledge, requiring a novel epistemology.

The fourth step is that of interpretation: what do these correlations mean. Do they imply causality and if so, in which direction? Might they be spurious, depending on other factors not yet mined or interwoven in more complex ways with factors not taken into consideration? In the case of behavioural data, one can ask whether the patterns indicate unconscious motivations, particular emotions, deliberate choices, socio-economic determinations, geographic or demographic influences, every time raising questions about our assumptions regarding cognition and perception. For instance, affective computing claims to 'map' our emotional states in relation to a number of factors, recasting our sense of being in control of our selves (Picard 1997). Brain scans are correlated with thoughts and emotions, according to some researchers potentially revealing our innermost life (Kamitani and Tong 2005). What interests me here is that the findings of data-mining operations require interpretation since just like in the case of text *these transcriptions do not speak for themselves*. We

must learn to ‘read’ them, and before all, we must learn to accept that they cannot be without bias, thus requiring a new critical stance capable of providing alternative interpretation of the same patterns as well as alternative computing techniques to mine the same dataset (Sculley and Pasanek 2008). What should worry us here is that further automation, as envisioned in e.g. autonomic computing, will remove this step. The narrative of autonomic computing is that the complexity of interacting computer networks will be such that no amount of human programmers can find all the bugs, repair system break-downs or sit down to interpret the results of real-time data-mining operations. This implies that the interpretation will be performed by the machines.

The fifth step concerns the application of the knowledge that was mined. Interestingly, this application can be used to check whether the ‘knowledge’ fits the new data; it allows modulating or even rejecting the inferred correlations. Within the context of science (brain imaging, human genetics) this will allow for more accurate models of explanation, within the context of security and commerce it should allow for more accurate measurements of the effectiveness of specific safety regulations or advertising strategies. In fact, the fifth step is entirely geared to fine-tune the consequences of targeted interventions. If certain Web-surf behaviour can be correlated with specific buying behaviour, this will allow businesses to personalize their service. Note that, behavioural advertising is probably the forerunner of more widespread targeted servicing that could eventually lead to smart environments like Ambient Intelligence and the Internet of Things, which are supposed to proactively cater to its users’ inferred preferences (Van den Berg 2010).

3.3 The Inference Problem

To the extent that smart infrastructures base their engagement with users on inferred preferences, we encounter a novel problem, which Dwyer (2009) has saliently coined as ‘the inference problem of pervasive computing’. It relates to the fact that in the narrative of proactive computing, the networked environment is always one step ahead of us because it is continuously ‘reading’ us. In my opinion, this is the most far-reaching implication (affordance) of the re-ontologization of the infosphere. Whereas, we have been used to *things* that do *not* anticipate our behaviours and to *people* whose anticipations we *can* guess to some extent, we are not at all accustomed to networked things that anticipate us without us having a clue as to how we are being anticipated. This relates to the problem of ‘double contingency’ put forward by Parsons’ and Luhmann: I can only act if I can anticipate how you will interpret my action, and the same goes for you (Vanderstraeten 2007). This either leads to the deadlock of a vicious circle, or creates the ambiguity and the openness of human society. One way to solve the problem is to take ‘the intentional stance’ towards others, assuming that they will act on reasons that we can follow thus allowing us to anticipate them and vice versa (Dennett 2009). The problem is that we may find it difficult to attribute reasons and meaning to the actions of the smart environment, either because we follow Searle’s Chinese Room argument (Searle 1980) and believe that machines cannot (yet) produce meaning or because we have no access to the sense that machines make of our behaviours. As to the first, this touches upon the epistemological implications of the novel knowledge production:

how can we perceive and cognize how we are being read? What would it mean if we take an intentional stance towards the smart environment even if we agree with Searle that the Smart Room does not ‘understand’ us the way we expect another person to understand us? As to the second, this connects with the fact that much of the knowledge we are talking about is proprietary or secret for security reasons; either hidden as a trade secret or as an intellectual property, or confidential because it is used to predict criminal behaviour (Harcourt 2007).

The most important point to be made here is that the re-ontologization that comes from the capture, storage, aggregation and analysis of ‘big data’ is not so much about our personal data proliferating ‘everware’ (Greenfield 2006). It is about the fact that some of our trivial data can be matched against knowledge that we are not aware of. The concept of personal data is not very relevant here: first, because any seemingly trivial data may become personal data once it is correlated with some profile at some point in the future (Hildebrandt 2008b), and second, because it is not our personal data that will make the difference but the knowledge mined from huge datasets. These datasets are composed of personal and other data from a massive amount of others, and these data may even be anonymized before mining them. The consequences of the data deluge do not relate to personal data but to inferred knowledge claims. Rouvroy (2011) and Rouvroy and Berns (2010) have described what Floridi may term the re-ontologization of the infosphere as a ‘statistical governance of the real’. The landscape that will surround us in the case of ubiquitous proactive computing will consist of networked artefacts that change their behaviour on the basis of a hidden complexity of continuous monitoring and real-time computations, meaning that we are being ‘read’ by way of statistics. At the same time, this landscape will be adapting itself based on the same real-time statistical calculus. This implies an unprecedented subliminal reconfiguration; instead of regulating our lives on the basis of written (legal) rules that we can see and contest, this infrastructure would regulate our life world implicitly. As it were from under the skin, below the threshold of consciousness (Hildebrandt 2011).

4 Some Ethical Implications of the Novel Epistemology

4.1 Re-ontologizing Privacy: Are We Being Read?

I am not sure whether it makes sense to speak of re-ontologizing privacy, because it might confuse those who think of an ontology as something that concerns what is given as a substance. Building on Floridi’s own discussion of our need for ‘a naive ontology of sufficiently permanent things and qualities’ (Floridi 2010, see above) I would, however, suggest that re-ontologizing of the infosphere demands a re-ontologization of privacy. Compared with e.g. *re-thinking* privacy speaking of a re-ontologization acknowledges that privacy requires integration into the informational infrastructure of our *Umwelt*, *Welt* and infosphere. In the end, we want the substance of privacy, not merely the idea. This dovetails with the observation that privacy is an affordance of the ICT infrastructure of the printing press, as Stalder and others (Hildebrandt and Koops 2010; Stalder 2002; Hildebrandt 2008a) have suggested, whereas it is unclear whether privacy will also be an affordance of the emerging

smart infrastructure that builds on real-time pervasive monitoring and proactive subliminal interventions.

Nevertheless, in order to re-ontologize privacy as an affordance of the novel infosphere we need to re-think what privacy means in the era of Ambient Intelligence and other types of smart environments. The most salient novelty of such environments is their capacity to infer our future behaviours and their capacity to act on these inferences without disturbing us with boring requests for confirmation or consent. The point is not that the relevant software ‘has our data’, but that the networked infrastructure will ‘read’ us. Whether they read us right or wrong is not even the point, because—paraphrasing Merton—‘if machines defines a situation as real, it is real in its consequences’. In the previous infosphere, we had two defences against being read by others: first, we could anticipate the way we would be read and adapt our behaviour; second, we could seek ways to shield ourselves from being read.

The first defence is perhaps hardly a matter of privacy. It engages our positive freedom to act, which depends on the double contingency mentioned above. Only if we have a clue as to how our actions will be interpreted, can we reach out and perform an action. Guessing how others will read us is mainly an intuitive, implicit process that ‘runs’ in the background, allowing us a smooth interaction. Making these intuitions explicit, however, will allow us to contest and resist them, choosing a course of action that goes against the grain of what is expected. The introduction of the script, which forces us to pay conscious attention to what we ‘read’, has at some point triggered the introduction of written law, *which thus in the end ‘affords’ its own contestation*. This is called the paradox of the *Rechtsstaat* or the advance of the Rule of Law as compared with the Rule by Law; it is a historical artefact made possible by the ICT infrastructure of the printing press (Hildebrandt 2008a). The problem of smart ubiquitous pervasive proactive environments is that we do not—yet—know how to anticipate their anticipation. This will make it hard to contest their inferences, except in an entirely intuitive manner. We lose the ability to consciously reflect on how we are being profiled and this will rule out contestation based on deliberation.

The second defence concerns our negative freedom, our ability to ‘unplug’ from social intercourse. This seems directly related to the idea of informational privacy. As Floridi (2005) has noted the right to informational privacy is often justified as a proto-ethical value that enables the realization of other ethical values (reducing privacy to a condition for other values, such as freedom, autonomy or fairness) or as a relation of ownership (whereby a person owns her personal data). Tavani (2008) has responded with the observation that most privacy theories actually fall within the category of either *control* theories or *restricted access* theories, arguing that a combination of both aspects would provide a more comprehensive theory of privacy. I refer to Solove (2002); Cohen (2000); Nissenbaum (2010) for more fuzzy (and therefore more precise?) conceptions. However, they all worry about control over or access to personal data, and more precisely they are all concerned with personal identifiability. The challenge we face is how to ‘unplug’ from being ‘read’, which is a much more extensive and a very different attack on our privacy. The point, moreover, is not merely that we must re-think privacy as the right not to be read all the time by an infrastructure capable of disseminating these ‘readings’ way beyond our circle of trust. The point is that we need an infrastructure that affords us a

measure of control over who can read us, without putting the burden of protection entirely on the shoulders of individual ‘users’.

To summarize, we need to design the upcoming proactive infosphere in a way that allows us to anticipate how we are being categorized, while at the same time also giving us space to be ‘unread’. This would constitute a re-ontologization of privacy, rather than a mere re-thinking of privacy.

4.2 Ontological Friction in the Infosphere: Who Is Transparent for Whom?

Floridi (2005: 186–87) has developed the notion of privacy as a function of ‘the ontological friction in the infosphere’:

The ontological features of the infosphere determine a specific degree of ‘ontological friction’ regulating the information flow within the system. ‘Ontological friction’ refers here to the forces that oppose the information flow within (a region of) the infosphere, and hence (as a coefficient) to the amount of work required for a certain kind of agent to obtain information (also, but not only) about other agents in a given environment, e.g. by establishing and maintaining channels of communication and by overcoming obstacles in the flow of information such as distance, noise, lack of resources (especially time and memory), amount and complexity of the data to be processed etc. Of course, the informational affordances and constraints provided by an environment are such only in relation to agents with specific informational capacities.

This notion is pivotal for an adequate understanding of privacy, because it does not start from individual users that control ‘their’ information, but from an infosphere that has as an *affordance* a measure of opacity of individual citizens. Floridi (2005: 190) is quite optimistic about how the novel infosphere affects this affordance:

It [re-ontologized infosphere, mh] has led not only to a huge expansion in the flow of personal information being recorded, processed and exploited, but also to a large increase in the types and levels of control that agents can exercise on their personal data.(...): At their point of generation, digital ICTs can foster the protection of personal data, e.g. by means of encryption, anonymization, password-encoding, firewalling, specifically devised protocols or services, and, in the case of externally captured data, warning systems. At their point of storage, legislation, such as the Data Protection Directive passed by the EU in 1995, guarantees that no ontological friction, already removed by digital ICTs, is surreptitiously reintroduced to prevent agents from coming to know about the existence of personal data records, and from accessing them, checking their accuracy, correcting or upgrading them or demanding their erasure. And at their point of exploitation—especially through data-mining, -sharing, -matching and -merging—digital ICTs could help agents to control and regulate the usage of their data by facilitating the identification and regulation of the relevant users involved.

In everyday practice, however, the Directive seems impotent to deliver the substance of the access and control rights it has attributed in the form of written law.

Our personal data proliferate and we have absolutely no clue as to which cloud is holding them or who is inspecting them for reasons of either profit or security (Koops and Leenes 2005). This might not be much of a problem to the extent that their proliferation easily turns these data into noise, re-establishing the ontological friction that constitutes privacy. However, the KDD technologies described above have the power to turn data into the currency they now present, creating an entirely new type of transparency and introducing an entirely different ontological friction. The transparency regards an invisible visibility. Whereas we cannot see how we are being profiled due to the proprietary regime on data mining and databases (creating a novel ontological friction), the data controllers who initiated these techniques have the means to make us visible as correlated humans (creating novel information entities that represent us as e.g. potential customers, terrorists, credit risks, healthcare consumers). The point is that in a constitutional democracy citizens should be shielded with a default of opacity, whereas those in authority should operate from a default of transparency (Gutwirth and De Hert 2008). The present re-ontologization of the infosphere has the opposite effect. So, while Floridi's concept of ontological friction in the infosphere is a rich analytic tool to observe changes in the affordances of the infosphere, its application should take into account that privacy is no longer about personal data but about statistically inferred pre-emptive personal profiles.

Floridi (2005: 194) provides a second analytical tool to discuss privacy in the re-ontologized infosphere. He advocates a re-interpretation of privacy:

Such re-interpretation is achieved by considering each person as constituted by his or her information, and hence by understanding a breach of one's informational privacy as a form of aggression towards one's personal identity.

Though he then links the notion of privacy with the notion of personal identity in a way that is problematic for many reasons (arguing in favour of biometric identification, at 197–8, but see e.g. Rouvroy (2007)), the link between informational privacy and personal identity is important (Hildebrandt 2006). Agre and Rotenberg (2001: 7) have defined the right to privacy as:

The freedom from unreasonable constraints on the construction of one's identity.

This is interesting for six reasons. First, this definition acknowledges that identity is not given, but the ephemeral result of a dynamic process. Second, it confirms negative freedom (*freedom from*) as the core of privacy, while—third—not every constraint is seen as a violation but only unreasonable constraints. Fourth, the constraints concern positive freedom (*freedom to*), thus providing a broader perspective to the concept of freedom instead of reducing it to the freedom to act arbitrarily. Fifth, the definition links privacy with the development of one's identity, while—sixth—understanding identity as inherently relational. This highlights the importance of context: who we are is co-determined by the context we engage with. If the context takes away the possibility to foresee how we are being anticipated, the freedom to develop our identity is at stake.

Returning to Floridi's proposition that a person is constituted by her information, I would fine-tune (or overturn?) this conception by arguing that a person is constituted

by her anticipation of what information people have on her and how this will impact their ‘reading’ of her. Though I am definitely constituted by my genes and by the information processes they coordinate at the level of the cell, I am not necessarily constituted by information about my genes. Note that without this information I am still constituted by my genes. I may, however, be constituted by information about my genes to the extent that this influences my sense of self, or my expectation of how others define me (e.g. my partner, or my insurance company). This explains why flows of information, rather than information by itself, co-constitute my identity. This also explains why the state of being ‘unplugged’ from social interaction (stopping, obfuscating, or filtering the flow of information) is necessary to reset one’s boundaries (Altman 1975), to recoup one’s internal coherence in the face of conflicting ‘readings’ by different persons or organizations, and also to reconstitute the self beyond the triggers of the environment. This refers to the way Floridi defines the autonomy of an agent: being ‘able to change state without direct response to interaction’ (Floridi and Sanders 2004: 9). Though his concept of agency refers to agency at the highest level of abstraction, meaning that it is not specific for human autonomy, we can confirm that at the level of human agency the abilities to interact and adapt must be coupled with the ability to reset oneself. The positionality of double contingency, typical for agents interacting and adapting beyond the *Umwelt* in a *Welt*, requires a measure of ‘unreadability’ to achieve autonomy. As in the case of all informational entities, both the construction and the maintenance of complexity require a sustained effort. Autonomy cannot be taken for granted.

5 A Legal Obligation for ISPs to Sustain Ontological Friction

5.1 Defining ISPs

Having traced some of the crucial epistemological and ethical implications of big number crunching, we must now investigate the role of ISPs as providers of Big Data and, thus, as enablers of extensive and detailed profiling. In this section, I will clarify what is meant with ISPs; in the next, I will investigate to what extent the present legal framework enables number crunching; and finally, I will discuss how the duty to pre-empt profiling is derived. This will allow me to conclude about the necessity to impose a legal obligation on ISPs that reinstates the ontological friction that is constitutive of constitutional democracy.

ISPs have been defined as (Perset 2009: 11):

Internet access providers, which provide subscribers with a data connection allowing access to the Internet through physical transport infrastructure.

They form a part of the larger group of ‘Internet intermediaries’ that function ‘(1) to provide infrastructure; (2) to collect, organise and evaluate dispersed information; (3) to facilitate social communication and information exchange; (4) to aggregate supply and demand; (5) to facilitate market processes; (6) to provide trust; and (7) to take into account the needs of both buyers/users and sellers/advertisers’ (Perset 2009: 6). Apart from internet access providers, examples of such intermediaries are

data processing and web hosting providers, search engines, portals, internet advertising networks, certification authorities, cloud computing platforms, internet payment networks and participative networked platforms such as Wikipedia, virtual worlds, social networking sites and more (Perset 2009: 7).

Though the difference between content and access providers may not always be obvious, I will focus on access providers that provide what has been called ‘mere conduit’. Practically speaking, this means broadband service providers (Yemini 2008). It implies that I will focus on the so-called ‘common carrier’ function of ISPs that highlights their role as preconditional for the functioning of a host of other applications. Providers of virtual worlds, wikis or social networking sites provide a structured and often proprietary framework for users to publish and co-create content, which depends on internet access. Their role extends far beyond the physical infrastructure that allows users to connect and communicate. For this reason, the consequences of their behaviour are different, entailing a different type of responsibility. Cloud computer platforms, however, may come close to ISPs in providing physical infrastructure for online activities. To the extent that software applications, audio and video will move from personal computers to the cloud (software as service and web streaming) we may need to involve them in our account of ISPs and the ethics of number crunching. As some authors (e.g. Carr 2008) have noted, this move will radicalize computing, data mining and the control of the web. It will influence, reinvent or even terminate net neutrality, requiring serious analysis in line with that of the responsibility of ISPs.

Within the context of this article, I understand ISPs as information entities, as agents and as moral agents, building on Floridi’s framework of Information Ethics (Floridi and Sanders 2004). Since I advocate not merely an ethical but also a *legal* obligation to maintain technical obstructions against systematic monitoring of content, it is important to note that ISPs are legal subjects in most—if not all—jurisdictions. They can in fact be held accountable in a court of law. Indeed, law preceded the discussion of agency for nonhuman entities, notably for corporations, funds and even ships, by taking a pragmatic view of legal personhood for non-natural persons (Wells 2001; French 1979). To the extent that these entities should be able to act in law (e.g. to contract) or should be held accountable for harm caused (e.g. in the case of tort or crime) positive law has attributed legal personhood to a variety of non-natural persons, thereby allowing second and third parties as well as victims to sue such an entity in a court of law. As Dewey (1926) has noted the attribution of legal personhood is artificial, but not imaginary, just like an artificial lake is not an imaginary lake. Once again paraphrasing Merton (1948), we can state: ‘if the law defines a situation as real, it is real in its (legal) consequences’.

5.2 ISPs Monitoring of Traffic Data and Content

I will now discuss existing and emerging incentives for ISPs to monitor the data packets they transport. This is important because once they get involved in systematic monitoring they can provide interested parties with Big Data to an extent that might even dwarf Google’s databases. At the moment of writing this article, one of the major Dutch ISPs, KPN, has announced that it is already using

Deep Packing Inspection to determine the Internet data traffic of its customers.¹ It claims to use this to check whether customers are using instant messaging or applications, in order to charge them for using applications like, for instance, Skype. This could jeopardize net neutrality, but in the end it will also enable extensive monitoring of the content of internet traffic.

Let me begin with a citation from a striking argument for net neutrality, based on the individual right to freedom of expression that has been put forward by Yemini (2008: 15):

In order to 'reach' the logical and content layers, one has to 'pass through' the physical layer; whoever controls the physical layer, unless restricted by law, becomes a gatekeeper for all other layers; and scarcity of physical layers means more control, and ability to realize that control, for fewer gatekeepers.

It should be obvious that the capacity of ISPs to track and trace our online behaviour is unprecedented. At this moment, a major amount of tracking is already taking place, first and foremost in order to provide good service to clients and to protect them from security risks. This involves monitoring to detect malware, to prevent spam and to manage bandwidth. Second, the Data Retention Directive 2006/24/EC requires monitoring and logging of traffic data as a means to improve law enforcement. Third, some ISPs may inspect data packets in order to detect IP infringements or illegal content,² though this is controversial and will easily violate privacy as well as data protection legislation if not warranted by law.

As Ohm (2009) explains, it is only recently the case that ISPs have the technical means to actually inspect the greater part of the content that passes on their network. In an even more vehement argument than that of Yemini, he proposes an alignment of privacy and net-neutrality arguments. Ohm (2009: 1420) declares somewhat dramatically:

Internet Service Providers (ISPs) have the power to obliterate privacy online. Everything we say, hear, read, or do on the Internet first passes through ISP computers. If ISPs wanted, they could store it all, compiling a perfect transcript of our online lives. In fact, nothing in society poses as grave a threat to privacy as the ISP, not even Google, a company whose privacy practices have received an inordinate amount of criticism and commentary. Although Google collects a vast amount of personal information about its users, an ISP can always access even more because it owns and operates a privileged net-work bottleneck, the only point on the network that sits between a user and the rest of the Internet. Because of this fact about network design, a user cannot say anything to Google without saying it first to his ISP, and an ISP can also hear everything a user says to any other websites like Facebook or eBay, things said that are unobtainable to Google. The potential threat to privacy from unchecked ISP surveillance surpasses every other threat online.

¹ See <http://english.capital.gr/News.asp?id=1194719> (last visited 14th May 2011).

² See for instance, DtecNet (at <http://dtecnet.com/>), which sells software solutions to ISPs that provide 'unique filtering solutions specifically targeting copyrighted content. With real-time traffic monitoring, the system allows ISPs to drastically reduce network congestion caused by distribution of illegal content on P2P protocols'.

There are various reasons why ISPs will be induced or even enforced to monitor the data packets they have under their care to a much further extent than at the present moment. The main incentives emerge in the context of business models and forensic investigations. In the case of business models we may expect that profits can be made by having clients pay for greater speed or special access to particular sites in case of congestion; this could easily violate net-neutrality (equal access to the internet for all). Also, ISPs could get involved with targeted advertising based on Web-surf behaviour, which could further erode the privacy of their clients. In many ways, ISPs would have access to far more data than even Google does, and we should have no doubt about the extent to which Google could 'read us' (Conti 2009). In the case of forensic investigation, the transnational struggle against child pornography, human trafficking, terrorism and anti-money laundering will continue to feed the agenda for more access to more behavioural data. As mentioned above, the Data Retention Directive 2006/24/EC requires Member States to impose a duty on providers of electronic communications to retain traffic data for not less than 6 months and for not more than 2 years. Though it does forbid Member States to base any obligation to retain the content of a communication on the Directive, traffic data easily contain many indications of content (e.g. in the subject header of an email) posing a much bigger threat for privacy than for instance the traffic data of telephone calls.

Another incentive that may warrant more intrusive monitoring derives from the ease with which Internet users can post all kinds of incorrect, sloppy, slanderous or otherwise irresponsible information. One could say that, worse than mere noise, such 'information' is not even meant to be truthful, coherent or respectful. From the perspective of Information Ethics, such 'information' could actually increase entropy and disrupt carefully constructed domains of knowledge, by disseminating incorrect rumours about whatever anybody likes to engage with. Hate speech or quasi-concerned comments that evoke hatred against pedophiles, Muslims or Christians as a group, or against specific politicians, journalists or other individuals can have far-reaching consequences for a person's reputation, capability to find work, insure herself or engage in public debate. Referring to Sunstein (2009), Fish (2011) notes that:

Rather than producing truth, the free and open marketplace of the Internet 'will lead many people to accept damaging and destructive falsehoods,' and unless there is 'some kind of chilling effect on false statements,' the 'proper functioning of democracy itself' may be endangered.

This has led to a call to hold ISPs accountable for the harm caused by the content they disseminate (Levmore and Nussbaum 2011; Fish 2011). This is an intelligible reflex, but the ensuing obligation for ISPs to monitor the data they 'carry' is very problematic. The European legislator seems to have been aware of this. Under the heading of 'No general obligation to monitor' art. 15 of the eCommerce Directive 2000/31 EC stipulated that

1. Member States shall not impose a general obligation on providers, when providing the services covered by Articles 12 [mere conduit, mh], 13 [caching, mh] and 14 [hosting, mh], to monitor the information which they transmit or store, nor a general obligation actively to seek facts or circumstances indicating illegal activity.

2. Member States may establish obligations for information society service providers promptly to inform the competent public authorities of alleged illegal activities undertaken or information provided by recipients of their service or obligations to communicate to the competent authorities, at their request, information enabling the identification of recipients of their service with whom they have storage agreements.

Though the need to fight hate speech, dangerous or incorrect information, child pornography and international terrorism may benefit from a duty to remove content that causes moral harm (conform also recital 48 of the same Directive that requires Member States to apply duties of care to detect and prevent certain types of illegal activities), good arguments can be given to couple this duty with a complementary obligation that requires ISPs to prevent *systematic* deep packet inspection. Removing content should then always be based on *incidental* deep packet inspection, conditional on specific charges of illegal content. As we have seen the Data Retention Directive that was enacted 6 years after the eCommerce Directive already overrules Art. 15 by imposing a general obligation to not only monitor but to actually retain traffic data. Even if ISPs do not—yet—use or sell these data for data mining, they are now forced to build an infrastructure that makes them available at the request of the competent national authority (which is not necessarily a judge). And though the Data Retention Directive stipulates that the retained data may only be provided in specific cases and in accordance with the necessity and proportionality requirements of e.g. the European Convention of Human Rights, the mere fact that such data are somehow kept in store for further inspection should worry us. Especially considering the kind of knowledge that can be inferred from this data, it becomes ever more important to stave off further obligations to retain data. In point of fact, we should reconsider to what extent this obligation must be mitigated or withdrawn. The judgements of the Constitutional Courts of various Member States provide ample indications of the highly problematic nature of these measures from the perspective of a substantive conception of the Rule of Law (De Vries et al. 2011 forthcoming). Considering the extent to which citizens could be ‘readable’ as a consequence of inferences drawn from the Big Data held by ISPs we need to consider imposing a legal duty to pre-empt systematic monitoring.

5.3 ISPs in a Constitutional Democracy

Having substantiated the preconditional role played by ISPs as enablers of a new type of transparency that will reduce the ontological friction that provides us with privacy, I will now clarify how the obligation to pre-empt this is derived.

In moral theory the repertoire for ethical duties is often reduced to two types of justification: deontological theories based on Kantian notions of human autonomy (which is assumed as given for a transcendental subject) (Rawls 2005) and utilitarian theories based on calculations of cost and benefit (differentiating between e.g. rule and act utilitarianism (Edmundson 2004), or plain and average utilitarianism (Rawls 2005)). A third way to understand ethical action has been located in virtue ethics, which focuses on the development of character and personality as an anchor for

practical wisdom. I would like to opt-out of this restricted framework and opt-into an ethics inspired on the work of pragmatists such as e.g. Peirce, Dewey, Holmes, G.H. Mead, Toulmin and Haack. Though both utilitarianism and pragmatism are consequentialist theories, they differ greatly in terms of their level of analysis. The pragmatist maxim looks at the meaning of words in terms of the consequences of their usage, coming close to Wittgenstein's understanding of language (Taylor 1995). Deontological, utilitarian as well as virtue ethics all build on some form of methodological individualism.

In the case of ISPs' monitoring capacities, such a restricted ethics will not do. The impact of the re-ontologization of the infosphere on societal checks and balances is not merely a matter of enlarging or reducing the scope of human autonomy. Its impact concerns asymmetries in access to and understanding of data, information, knowledge and meaning. This requires an ethics that takes into account the consequences for the polity, for civil society, for distributive justice, reciprocal power relationships and e.g. the state's capacity to legislate and provide the legal certainty that allows individual citizens to act as relatively autonomous agents. As argued in the preceding section, ISPs constitute the enabling information and communication infrastructure of our societies. They form the bottom line of a host of other applications that are turning into the critical infrastructure of everyday life. To some extent they build on, partner with, replace and transform the information and communication infrastructure of the printing press. Having different affordances they will induce different behaviour patterns, and according to some authors even different perception and different cognition (Carr 2010; Tapscott 2009). This will affect the constitution of authors and actors as well as the polity they constitute. From the perspective of democratic theory and the rule of law, this implies that the duties resting on ISPs cannot be based on their individual ethical virtue, categorical duties or calculations of cost and benefits. Rather, these duties derive from the framework of constitutional democracy, i.e. the historical artefact that sustains the tension between the institution of self-rule and the constraints imposed on democratic majorities by the protection of human rights and constitutional arrangements that produce the autonomy they aim to protect. The argument is based on a substantive conception of the Rule of Law (or the 'Rechtsstaat'), which defines the Rule of Law in terms of an effective commitment to human rights and the separation of powers. This substantive conception provides legal protection not afforded by the formal conception, which merely checks whether the government has adhered to certain procedural standards when governing, legislating and adjudicating (Grote 1999; Koetter 2010). A substantive conception of the Rule of Law can, thus, prevent democracy from turning into the tyranny of majority rule, as it will safeguard the liberties of individual citizens and their rights to dissent and to develop new majorities.

Precisely because ISPs provide the novel ICT infrastructure of current democracies we cannot be indifferent to the way this infrastructure is designed. Their role as scaffolding of a viable constitutional democracy actually demands that they do not merely refrain from monitoring their data packets, but actually build in the legal protection against such monitoring at the level of the technical infrastructure.

6 Conclusions: Who Needs Stories if You Can Get the Data?

In this contribution, I have focused on how ISPs may violate the identity of individual human agents by obstructing their ability to anticipate how they are being read and their ability to shield themselves from being ‘read’ by their smart environment. My argument builds on the assumption that in a constitutional democracy, privacy requires *a default of transparency of data miners* and *a default of opacity of individual persons*.

In his provocative article in *Wired Magazine*, Anderson (2008) claimed that data mining will ultimately refute theory. We could paraphrase him by saying: who needs theory if you can get the data? Similarly, we may assume that statistical inferences could one day refute the stories a person tells to give an account of herself. At some point the police, the judge, the insurance agent, your webshop, Google Analytics, your doctor and your tax collector will point out that you match an aggregated profile that probably (sic!) provides a better prediction of your future behaviour than whatever you tell them. In this contribution I have argued that this is not science fiction nor unavoidable. Information and communication infrastructures do not entirely determine our interaction, though they often determine the bandwidth. Though we invent them, they in turn invent us. Insofar as we begin to notice how different ICT designs have different implications, constitutional democracy requires to design them in a way that protects individual persons against statistical determinations they cannot foresee. This requires us to build in the necessary ontological frictions to guarantee a measure of opacity for citizens and to remove the ontological frictions that prevent us from anticipating how we are anticipated. This duty cannot be put on the shoulders of individual ISPs that must survive in an incentive structure that forces their hand to monitor our web behaviours. Such a duty must be decided upon within the framework of democratic participation and constitutional constraints, safeguarding a level playing field tuned to the novelties of a computational infosphere.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Agre, P. E., & Rotenberg, M. (Eds.). (2001). *Technology and privacy: the new landscape*. Cambridge: MIT Press.
- Altman, I. (1975). *The environment and social behavior. Privacy personal space territory crowding*. Monterey: Brooks/Cole.
- Anderson, C. (2008). The end of theory: the data deluge makes the scientific method obsolete. *Wired Magazine*, 16(7).
- Ayres, I. (2007). *Super crunchers: why thinking-by-numbers is the new way to be smart*. New York: Bantam Books.
- Bateson, G. (1972). *Steps to an ecology of mind*. New York: Ballantine.
- Butler, J. (2005). *Giving an account of oneself* (1st ed.). New York: Fordham University Press.
- Carr, N. G. (2008). *The big switch: rewiring the world, from Edison to Google*. New York: W.W. Norton.
- Carr, N. (2010). *The shallows: what the internet is doing to our brains*. New York: W.W. Norton.
- Cheung, T. (2006). The language monopoly: plessner on apes, humans and expressions. *Language & Communication*, 26, 316–330.

- Cohen, J. (2000). Examined lives: informational privacy and the subject as object. *Stanford Law Review*, 52(1373–1437).
- Conti, G. (2009). *Googling security: how much does Google know about you?* Reading: Addison-Wesley.
- De Vries, K., Bellanova, R., De Hert, P., & Gutwirth, S. (2011). The German Constitutional Court Judgement on data retention: proportionality overrides unlimited surveillance (doesn't it?). In S. Gutwirth, Y. Poullet, P. De Hert, & R. Leenes (Eds.), *Privacy and data protection: an element of choice*. Springer: Dordrecht. forthcoming.
- Dennett, D. (2009). Intentional systems theory. In *Oxford handbook of the philosophy of mind* (pp. 339–350). Oxford: Oxford University Press.
- Dewey, J. (1926). The historic background of corporate legal personality. *The Yale Law Journal*, 35(6), 655–673.
- Dwyer, C. (2009). The inference problem and pervasive computing. In *Proceedings of Internet Research 10.0.*, Milwaukee, WI.
- Edmundson, W. A. (2004). *An introduction to rights (Cambridge introductions to philosophy and law)*. Cambridge: Cambridge University Press.
- Eisenstein, E. (2005). *The printing revolution in Early Modern Europe* (second ed.). Cambridge: Cambridge University Press.
- Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P., & Uthurusamy, R. (Eds.). (1996). *Advances in knowledge discovery and data mining*. London: AAAI Press.
- Fish, S. (2011). Anonymity and the dark side of the internet. *The New York Times* (3 January).
- Floridi, L. (2005). The ontological interpretation of informational privacy. *Ethics and Information Technology*, 7(4), 185–200.
- Floridi, L. (2010). The philosophy of information as a conceptual framework. *Know Techn Pol*, 23, 253–281. doi:10.1007/s12130-010-9112-x.
- Floridi, L., & Sanders, J. W. (2004). On the morality of artificial agents. *Minds and Machines*, 14(3), 349–379.
- French, P. A. (1979). The corporation as a moral person. *American Philosophical Quarterly*, 16(3), 207–215.
- Geisler, D. M. (1985). Modern interpretation theory and competitive forensics: Understanding hermeneutic text. *The National Forensic Journal*, III (Spring), pp. 71–79.
- Gibson, J. (1986). *The ecological approach to visual perception*. New Jersey: Lawrence Erlbaum Associates.
- Goody, J., & Watt, I. (1963). The consequences of literacy. *Comparative Studies in Society and History*, 5 (3), 304–345.
- Greenfield, A. (2006). *Everyware. The dawning age of ubiquitous computing*. Berkeley: New Riders.
- Grote, R. (1999). Rule of Law, Rechtsstaat and 'Etat de droit'. In C. Starck (Ed.), *Constitutionalism, universalism and democracy—a comparative analysis* (pp. 269–306). Baden-Baden: Nomos.
- Gutwirth, S., & De Hert, P. (2008). Regulating Profiling in a Democratic Constitutional State. In M. Hildebrandt & S. Gutwirth (Eds.), *Profiling the European Citizen. Cross-disciplinary perspectives* (pp. 271–302). Dordrecht: Springer.
- Harcourt, B. E. (2007). *Against prediction: profiling, policing, and punishing in an actuarial age*. Chicago: University of Chicago Press.
- Hildebrandt, M. (2006). Privacy and Identity. In E. Claes, A. Duff, & S. Gutwirth (Eds.), *Privacy and the Criminal Law* (pp. 43–58). Antwerpen-Oxford: Intersentia.
- Hildebrandt, M. (2008a). A Vision of Ambient Law. In R. Brownsword & K. Yeung (Eds.), *Regulating Technologies*. Oxford: Hart.
- Hildebrandt, M. (2008b). Profiles and Correlatable humans. In N. Stehr & B. Weiler (Eds.), *Who owns knowledge? Knowledge and the Law* (vol. 265–285). New Brunswick: Transactions Books.
- Hildebrandt, M. (2011). Autonomic and autonomous 'thinking': preconditions for criminal accountability. In M. Hildebrandt & A. Rouvroy (Eds.), *Law, human agency and autonomic computing. The Philosophy of Law meets the Philosophy of Technology*. Abingdon: Routledge.
- Hildebrandt, M., & Koops, B. J. (2010). The challenges of Ambient Law and legal protection in the profiling era. *Modern Law Review*, 73(3), 428–460.
- Ihde, D. (1990). *Technology and the lifeworld: from garden to earth (The Indiana series in the philosophy of technology)*. Bloomington: Indiana University Press.
- Ihde, D. (2002). *Bodies in technology (Electronic mediations v. 5)*. Minneapolis: University of Minnesota Press.
- Kallinikos, J. (2006). *The consequences of information. Institutional implications of technological change*. Cheltenham: Edward Elgar.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the brain. *Nature Neuroscience*, 8(5), 679–685. doi:10.1038/nm1444.

- Koetter, M. (2010). Rechtsstaat und Rechtsstaatlichkeit in Germany. In M. Koetter, & G. F. Schuppert (Eds.), *Understandings of the Rule of Law in various legal orders of the World*. Available at <http://wikis.fu-berlin.de/download/attachments/24511234/Koetter+Germany.pdf>.
- Koops, B.-J., & Leenes, R. (2005). 'Code' and the Slow Erosion of Privacy. *Michigan Telecommunications and Technology Law Review*, 12(1), 115–189.
- Levmore, S. X., & Nussbaum, M. C. (2011). *The offensive Internet: speech, privacy, and reputation*. Cambridge: Harvard University Press.
- Merton, R. K. (1948). The self-fulfilling prophecy. *The Antioch Review*, 8(2), 193–210.
- Negroponte, N. (1996). *Being digital* (1st Vintage (Bookst) ed.). New York: Vintage Books.
- Nissenbaum, H. F. (2010). *Privacy in context: technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.
- Ohm, P. (2009). The rise and fall of invasive ISP surveillance. *University of Illinois Law Review* (5), 1417–1496.
- Ong, W. (1982). *Orality and literacy: the technologizing of the word*. London: Methuen.
- Perset, K. (2009). The economic and social role of Internet intermediaries. OECD Directorate for Science, Technology and Industry.
- Picard, R. (1997). *Affective computing*. Cambridge: MIT Press.
- Plessner, H. (1975). *Die Stufen des Organischen unter der Mensch. Einleitung in die philosophische Anthropologie*. Berlin: Walter de Gruyter.
- Rawls, J. (2005). *A theory of justice* (original edn.). Cambridge: Belknap.
- Ricoeur, P. (1973). The Model of the text: meaningful action considered as a text. *New Literary History*, 5(1), 91–117.
- Ricoeur, P. (1992). *Oneself as Another* (K. Blamey, Trans.). Chicago: The University of Chicago Press.
- Rouvroy, A. (2007). *Human Genes and Neoliberal Governance. A Foucauldian Critique*. Abingdon: Routledge-Cavendish.
- Rouvroy, A. (2011). Technology, Virtuality and Utopia: Governmentality in an Age of Autonomic Computing. In M. Hildebrandt & A. Rouvroy (Eds.), *The Philosophy of Law Meets the Philosophy of Technology. Autonomic Computing and Transformations of Human Agency*. Abingdon: Routledge. forthcoming.
- Rouvroy, A., & Berns, T. (2010). Le nouveau pouvoir statistique. *Multitudes* (40), 88–103.
- Sculley, D., & Pasanek, B. M. (2008). Meaning and mining: the impact of implicit assumptions in data mining for the humanities. *Literary and Linguistic Computing*, 23(4), 409–424.
- Searle, J. (1980). Minds, brains, and programs. *The Behavioral and Brain Sciences*, 3(3), 517–557.
- Solove, D. J. (2002). Conceptualizing privacy. *California Law Review*, 90, 1087–1156.
- Stalder, F. (2002). The Failure of Privacy Enhancing Technologies (PETs) and the Voiding of Privacy. *Sociological Research Online*, 7(2). Available at <<http://www.socresonline.org.uk/7/2/stalder.html>>.
- Sunstein, C. R. (2009). *On rumors: how falsehoods spread, why we believe them, what can be done* (1st ed.). New York: Farrar, Straus and Giroux.
- Tapscott, D. (2009). *Grown up digital: how the net generation is changing your world*. New York: McGraw-Hill.
- Tavani, H. T. (2008). Floridi's ontological theory of informational privacy: some implications and challenges. *Ethics and Information Technology*, 10(2–3), 155–166.
- Taylor, C. (1995). To Follow a Rule. In C. Taylor (Ed.), *Philosophical arguments* (pp. 165–181). Cambridge: Harvard University Press.
- Turilli, M., & Floridi, L. (2009). The ethics of information transparency. *Ethics and Information Technology*, 11(2), 105–112.
- Van den Berg, B. (2010). *The situated self: identity in a world of ambient intelligence*. Nijmegen: Wolf Legal Publishers.
- Vanderstraeten, R. (2007). Parsons, Luhmann and the Theorem of Double Contingency. *Journal of Classical Sociology*, 2(1), 77–92.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind. Cognitive science and human experience*. Cambridge: MIT Press.
- Wells, C. (2001). *Corporations and criminal responsibility* (2nd edn., [Oxford monographs on criminal law and justice]). Oxford: Oxford University Press.
- Yemini, M. (2008). Mandated network neutrality and the first amendment: lessons from Turner and a new approach. *Virginia Journal of Law & Technology*, 13(1), 1–38.