

 Open access • Journal Article • DOI:10.1007/S10676-006-9133-Z

Maintaining the reversibility of foldings: Making the ethics (politics) of information technology visible — [Source link](#)

Lucas D. Introna

Institutions: Lancaster University

Published on: 01 Mar 2007 - Ethics and Information Technology (Kluwer Academic Publishers)

Topics: Ethics of technology, Morality and Information technology

Related papers:

- [Reassembling the Social: An Introduction to Actor-Network-Theory](#)
- [Research Commentary: Desperately Seeking the IT in IT Research--A Call to Theorizing the IT Artifact](#)
- [Morality and Technology](#)
- [Pandora's Hope: Essays on the Reality of Science Studies](#)
- [Human-Machine Reconfigurations: Plans and Situated Actions](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/maintaining-the-reversibility-of-foldings-making-the-ethics-wkthxvte14>

Maintaining the reversibility of foldings: Making the ethics (politics) of information technology visible

Lucas D. Introna

*Department of Organisation, Work and Technology, Lancaster University, Management School,
Lancaster University, Lancaster, LA1 4YX, UK
E-mail: l.introna@lancaster.ac.uk*

Abstract. This paper will address the question of the morality of technology. I believe this is an important question for our contemporary society in which technology, especially information technology, is increasingly becoming the default mode of social ordering. I want to suggest that the conventional manner of conceptualising the morality of technology is inadequate – even dangerous. The conventional view of technology is that technology represents technical means to achieve social ends. Thus, the moral problem of technology, from this perspective, is the way in which the given technical means are applied to particular (good or bad) social ends. In opposition to this I want to suggest that the assumed separation, of this approach, between technical means and social ends are inappropriate. It only serves to hide the most important political and ethical dimensions of technology. I want to suggest that the morality of technology is much more embedded and implicit than such a view would suggest. In order to critique this approach I will draw on phenomenology and the more recent work of Bruno Latour. With these intellectual resources in mind I will propose disclosive ethics as a way to make the morality of technology visible. I will give a brief account of this approach and show how it might guide our understanding of the ethics and politics of technology by considering two examples of contemporary information technology: search engines and plagiarism detection systems.

Key words: disclosive ethics, design, Heidegger, politics, Latour, technology

Technology and morality

Technology has become the default way of ordering and organising our society – maybe it has always been. Nevertheless, its role in organising society has now become so legitimate as to remain almost beyond question and scrutiny. Today, every attempt to ‘modernise’ society is inevitably connected with technology, particularly with information technology. Implicit in this wholehearted embrace of technology as a mode of social ordering is a fundamental belief that technology is something at our disposal, for us to deploy according to our needs. In other words that technology is a neutral means (or medium) that can be shaped and moulded to suite our needs and desires. More specifically that technology is a technical means to achieve certain social ends. Conceived in this way the important ethico-political questions emerge with reference to the social ends that the technology is supposed to serve, or not. From this perspective the ‘medium’ itself, technology as such, remains unproblematised. It is my suggestion

that this view of technology is problematic if we are to understand the ethico-political import of technology. The means/ends distinction – technical means/social ends – that this view assumes is based on a more fundamental assumption. It assumes that there is an ontological separation between the technical world and the social world. In other words that they exist as two *separate* types of realities that ‘impact’ on each other in some way or another when they meet or connect. This assumption may be useful, even pragmatically necessary, for the purposes of designing technology. Nevertheless, when it slips into the background, and we forget it – or worse, forget that we have forgotten it – then it can become a very powerful way of hiding the most important ethico-political aspects of our contemporary, technologically inspired, project of modernisation. It is my claim that this assumption should be made visible again. I will suggest, with Heidegger (1977, p. 4), that the “essence of technology is by no means anything technological”. Unless, we see the intimate co-constitutive relationship between the ‘technical’ and the ‘social’

we will not grasp the full ethico-political import of our often unquestioned embrace of technology. This paper hopes to further contribute to the opening up of the ‘black box’ of the socio-technical relation, and in so doing bring the fullness of the ethico-political nature of technology into view – to disclose what we mostly forget in the flow of getting on with everyday life. I will structure the paper as follows. First, I will discuss the traditional way for articulating the socio-technical relation. Second, I will provide a critique of this position drawing on the work of Heidegger and Latour. I will then proceed to suggest how such a perspective changes the way we frame the ethico-political questions with regard to technology by proposing disclosive ethics as an alternative approach. Forth, I will present two case examples that will aim to make the conceptual discussion more concrete. Finally, I will suggest some ways in which disclosive ethics could become an appropriate way of dealing with the morality of technology in practice.

*Morality is a human problem
not a technological one*

The most common view of technology is that it is an artefact, tool, or system that is designed to be available for humans to achieve their objectives and outcomes – to e-mail, to write, to store, to manipulate, to interact, and so forth. This view is rooted in our everyday intuitions about the world in which our tools are seen as something distinctly separate from us. In this view tools are seen as objective and neutral ‘technical things’ (separate from us) that we can draw upon, or not, to achieve our particular ends. This relationship between us and our tools is often expressed as a *means/ends relationship* where technology is designed – based on some technical rationality – as means (or tool) to achieve particular ends. Some of these tools might be useful and others not. However, when users take up tools or artefacts (e-mail, word processor, mobile phone, etc.) these tools will tend to have an *impact* on the way they do things. For example if I send you an e-mail I write and communicate differently to sending you a paper-based letter. Thus, according to this view, we need to understand the impact that ICT has on social practices as these tools are taken up and used in these everyday practices and situations. For example, how will communication with mobile phones change or impact on, our social interaction, and possibly our social relationships?

In posing these questions on the impact of this or that technology this view does not primarily concern itself with the prior process of how this or that particular technology was developed – why and how did it

come about in the first instance, and why not something else? These questions are seen as ‘technical’ questions that are the domain of designers and engineers (working within the assumption of technical rationality). It is taken for granted that the particular technology is the outcome of a more or less technical rational attempt by designers and engineers to solve concrete practical problems by the ‘technical’ means available to them – obviously within certain economic constraints. From this perspective the relevant question is how society uses these technical means – for what purposes – and how this usage impacts on, and changes, social practices. In thinking through these questions – of the social or ethical significance of technology – it is normally assumed that the particular technology – mobile phones in this case – operates in a more or less uniform manner in different social settings. In other words, it assumes that a particular technology has certain *determinate* effects on, or in, the context of its use. This way of conceptualizing ICT leads to questions such as: what is the impact of the Internet on education? Or what is the impact of closed circuit television (CCTV) on privacy?

The fundamental (or ontological) assumption that the ‘technical’ and the ‘social’ are ontologically different types of reality allows the tool view to retain a distinction between the ‘technical’ means and the ‘social’ ends – or between facts and values as critiqued by Latour (2002). For example, for adherents of the tool view, technology is neutral and it is humans that bring (good or bad) intentions or values to technology when they use it. This view is often expressed in the debate on guns in the following manner: ‘it is people that are bad not guns’. Similarly, for them the Internet is a neutral (technical) medium that can be used for good (education, community, commerce, etc.) or for bad (pornography, gambling, terrorism, etc.).

Much of the ethical and policy debate about information technology has been informed by this, what I refer to as ‘tool view’ of technology. In these debates the concern is to regulate or justify the conduct or practices that the technology now makes possible (or prevents). These policies are seen, and presented as, ways to regulate or balance competing rights or competing values as these are made possible (or not) through the new technology. For example, what sort of policies do we need to protect our children when they go on the Internet, as the Internet now allows them to go ‘anywhere’? How would these policies affect the right to free speech? Or, what sort of policies do we need to secure the rights of producers of digital products? How would these policies affect the right of society to a reasonable access to these products? Furthermore, these debates are most often directed at an institutional level of discourse –

that is, with the intention to justify (or not) the policies or conduct for governments, organizations, and individuals. In these debates, on the impact of ICT and new media, ethics and ethicists are primarily conceived of as presenting arguments for justifying a particular policy (or set of policies) that balance certain values or rights, over and against other possible policies. In presenting these arguments ethicists normally apply ethical theories (such as consequentialism, utilitarianism, deontological ethics, to name three) to new cases or problems presented by the use, or perceived impact, of the particular technology.

To summarise: the tool view of technology tends to allocate all agency to human beings. Its understanding of the human/technology relationship is anthropocentric in as much as it is the human that does things with or by means of the neutral technical capability – tools at the disposal of human ends.

Questioning the means/ends distinction

The assumed separation between us and our tools has been a focus of critique for constructivists for some time now (Bijker et al. 1987; Callon 1987; Law 1991; Latour 1991; Pfaffenberger 1992; Berg and Lie 1995; Bijker 1995; Bowker and Star 1999). They argue that the taken for granted assumption (or ontology) of the subject/object dualism (also expressed as the technology/society dualism) leads us to inappropriate conclusions about the ethical import of technology – for them means and ends cannot be separated as suggested in the tool view. They argue that this view does not take into account that the technology does not simply appear but is the outcome of a complex and socially situated development and design process. By socially situated I mean subject to the aspirations, interests, power, values, assumptions, beliefs, etc., of a diverse set of potential stakeholders – such as financiers, technologists, users, markets, to name a few. In this development and design process *many alternative options become excluded* or closed off in favour of the technology that is now available – obviously with important implications. In other words there are many cultural, political and economic forces that shape the particular options suggested as well as the way the selected options become designed and implemented (Bijker et al. 1987). Thus, it is not only technology that ‘impacts’ on society; technology itself is already the outcome of complex, subtle, and situated social processes. Moreover, they argue that when we look at the actual uses of particular technologies we discover that users interpret and use them in many diverse and often unexpected ways – leading to many and diverse unintended consequences. Indeed, the degree to which the tech-

nology/society distinction is useful at all (ontologically or analytically), as well as ‘where’ the ethics and politics of technology is located, varies between the different constructivist authors. I will not take up this debate here.¹ Rather I want to focus on the work of Bruno Latour as my position is most closely aligned with his views, especially his more recent work.

For Latour (1999, 2003) the question of ‘where’ the politics (and ethics) is located – i.e. in the human or in the tools – as such is not a relevant question because, for him, the social and the technical are a unity *from the start* – they have never been otherwise. For him, to account for humans and non-humans in ways that would suggest that they are separately already what they are – as ‘social’ and ‘technical’ – and then we ‘add’ them together to ‘make’ a socio-technical world would simply be wrong. Latour (2003) suggests that both humans and non-humans share a common and ongoing co-constitutive history: “Humans and non-humans are engaged in a history that should render their separation impossible.” (p. 39). More than that, they do not merely share a common history; they are each other’s common history: “A body corporate is what we and our artefacts have become. We *are* an object institution” (Latour 1999, p.192, emphasis added). In this ‘object institution’ – that has never been otherwise – it may not be possible to simply allocate intentionality and properties this way or that way: “Purposeful action and intentionality may not be properties of objects, but they are also not properties of humans either. They are properties of institutions [collectives of humans and non-humans], apparatuses, or what Foucault called *dispositifs*” (Latour 1999, p. 192). It is clear from these comments that Latour is talking about the human/non-human relationship as a fundamental co-constitutive unity in ways very similar to Heidegger (1962, 1977). One of the important consequences of Latour’s and Heidegger’s position is that our encounter with technology – as a possibility to do this or that – is always conditioned from the very beginning, and in some fundamental way, by the horizon of intelligibility within we already find ourselves – it is already in some strict sense *pre-given*, and as such pre-understood as an inherent part of our way

¹ Refer to Brey, Philosophy of Technology meets Social Constructivism, *Techné. Journal of the Society for Philosophy and Technology*, 2(3/4): 56–79, 1997 or Sismondo, Some Social Constructions. *Social Studies of Science*, 23(3): 515–553, 1993, also see the debate between Kling, and Woolgar and Grint in Volumes 16 and 17 of *Science, Technology & Human Values* in 1991 and 1992. See also Winner, Upon Opening the Black Box and Finding it Empty: Social Constructivism and the Philosophy of Technology. *Science, Technology, and Human Values*, 18(3): 362–378, 1993.

of being (Introna and Ilharco 2006). They – the tools – constitute us as much as we constitute them; differently stated they fold into us as much as we fold into them. Let us consider an example.

When a consultant takes up a mobile phone the consultant acquires a certain capability (to contact and be contacted) but that is not all that happens. We need to look at this seemingly obvious statement of ‘consultant,’ ‘mobile phone’ and ‘taking up’ more closely. The mobile phone only becomes ‘a mobile phone’ when taken up by the consultant. When it lies on the table it is a potential to be ‘a mobile phone’, but it only becomes an actual ‘possibility for contacting’ when it is picked up and one dials the number, and, of course, there is sufficient credit on the account, and it is possible to get a signal, and so forth. In taking up the mobile phone both the mobile phone and the consultant are transformed – i.e. reconstituted. The mobile phone is no longer ‘merely’ an object and the consultant becomes a human that embodies the possibility to contact and be contacted at a distance. With the mobile phone in hand, the world, now and in the future, becomes revealed in new ways (for example, a person ‘far off’ is suddenly ‘near’); previous possibilities not visible or evident at all suddenly emerge as possible possibilities – some intended, some unintended. What this example shows is that the mobile phone and the consultant are each other’s constitutive condition for being what they are as ‘mobile phones’ and as ‘consultants.’ Obviously the mobile phone is just one of a multiplicity of constitutive relations or conditions that constitutes the horizon of possibilities for a person to ‘be a consultant’. However, in my example thus far I have limited our discussion to the relationship between the person and mobile phone. As such I have not yet revealed the full constitutive horizon at work in the example. Recognizing the possibilities or affordances² of the mobile phone draws on a prior familiarity with a world where things like mobile phones and the act of phoning by using a device already makes sense – we can imagine many worlds where it would not make sense. If we were to locate the mobile phone in a culture where such practices do not exist at all the mobile phone will not even show up as ‘a mobile phone’. It might just show up as an odd and strange object lying there – in the way we sometimes encounter archaeological objects from alien cultures. Thus, for the example to make sense at all – also for us as readers – it draws on an already present familiarity of a world in which such things and practices already makes sense (Heidegger, 1962,

² Affordances are the perceived properties of an artefact that suggests how it should or could be used (Norman 1988).

pp. 97–98) – a way of being that is already pre-given. Latour tries to capture this idea of a human world *already implied* in our encounter with the tool with the notion of a ‘fold’.³

I would like to define the regime proper to technology by the notion of fold... What is folded in technical action? Time, space and the type of actants. ...[one might say a world]. If, for pedagogical reasons, we would reverse the movement of the film of which this hammer is but the end product, we would deploy an increasing assemblage of ancient times and dispersed spaces: the intensity, the dimension, the surprise of the connections, invisible today, which would thus have become visible, and, by contrast, would give us an exact measure of what this hammer accomplishes today. (2002, pp. 248–249)

In this co-constitutive relationship of simultaneous giving (of possibilities) and taking (of opportunities) – as Latour reminded us – the “intentionality [at work in this relationship] may not be properties of objects, but they are also not properties of [the] humans either” (Latour 1991, p. 192). This view of technology takes agency to be *simultaneously human and material from the start*. It is not anthropocentric. It treats technology as material culture that is not neutral but the very condition of our way of being, what we are as contemporary human beings – it constitutes us as much as we constitute it. In this co-constitutive relationship means and ends exist simultaneously and constitute each other’s possibility to be. In other words our tools, like us, and simultaneously with us, are important political sites.⁴ It is these political sites – a co-constitutive nexus of tool and human relationships – that we ought to take very seriously if we really want to take up the problem of the ethics of technology. How will we do it?

³ The notion of the fold is not new. Indeed Latour acknowledges his debt to Deleuze. For a more detailed discussion of Deleuze’s notion of the fold refer to Tom Conley, *Folds and Folding*. In C.J. Stivale, editors, *Gilles Deleuze: Key Concepts*. pp. 170–181. McGill-Queen’s University Press, Montreal, 2005.

⁴ I use the term ‘site’ here in the sense that Schatzki, *The Site of the Social. A Philosophical Account of the Constitution of Social Life and Change*. University Park Penn, Pennsylvania University Press, (2002; Schatzki, *The Sites of Organizations. Organization Studies*, 26(3): 465–484, 2005 uses it. For him a site “are arenas or broader sets of phenomena as part of which something – a building, an institution, an event – exists or occurs.” Or “A site is a type of context [which] can loosely understood as an arena or set of phenomena that surrounds or immerses something *and enjoys powers of determination with respect to it*” (p. 468, emphasis added).

Disclosive ethics: to maintain the reversibility of foldings

To maintain the reversibility of foldings: that is the current form that moral concern takes in its encounter with technology. (Latour 2002, p. 258)

Folded into – or enclosed in – the ongoing co-constitutive horizon or nexus of human and technology relationships are (un)intentions, (im)possibilities, (dis)functions, affordances/prohibitions that renders possible some ways of being and not others, that serves the (il)legitimate interests of some and not others. In the case studies below I will show how this nexus constitutes some websites/pages as desirable to search engines and not others; and some students as plagiarists and not others. I want to show how this political site is constituted through algorithms and practices in a manner that makes it impossible to simply trace and tie this politics down to this or that particular intention or agency, human or tool. I want to show that it is in fact the outcome of the ongoing and *simultaneous* play of human and material agency. These examples will show that ethics is not something that comes afterwards – when we already have ‘humans’ and ‘technologies’ – rather it is always and already present in a fundamental way. Every conceiving, shaping, materialisation and taking up of this or that artefact or technology is already rendered possible and meaningful (or not) by an already taken for granted way of being. As Latour (2002) rightly suggests:

“Morality is no more human than technology, in the sense that it would originate from an already constituted human who would be master of itself as well as of the universe...Morality and technology are ontological categories ...and the human comes out of these modes, it is not at their origin.” (p. 254)

The *telos* of this ethico-political site – the nexus of ongoing human and technology relations—is closure. For technology to function as ‘technology’ – we may also say for politics to function ‘as politics’ – it seeks closure – or ‘enrolment’ in the actor network theory language. This process of closure or enrolment can best be understood as a process of hegemonisation. Critchley (2004) suggests that:

Politics is the realm of the decision, of action in the social world, of what Laclau, following Gramsci, calls hegemonisation. Hegemonisation is understood as actions that *attempt to fix the meaning of social [or socio-technical] relations*. ... If we can see politics with the category of hegemony, and in my

view it is best conceived with that category, then politics is an act of power, force and will that is contingent through and through. Hegemony reveals politics to be the realm of contingent decisions by virtue of which subjects (understood here as persons, parties or social movements) *attempt to articulate and propagate meanings of the social [and the socio-technical]*. At its deepest level, the category of hegemony discloses the political logic of the social. (p. 115, emphasis added)

This politics, or hegemonisation, is necessary, it cannot be avoided. Decisions (and technologies) need to be *made* and programmes (and technologies) need to be *implemented*. Without closure technology (politics) cannot be effective as a programme of action or ordering. Obviously, if the interests of the many are included – in the *enclosure* as it were – then we might say that it is a ‘good’ politics (such as democracy). If the interests of only a few are included we might say it is a ‘bad’ politics (such as totalitarianism). Nevertheless, all technological/political events of enclosing are always in a sense violent as they always include (certain intentions, possibilities or affordances) and exclude others as the very condition of the operation of hegemonisation. It is this ongoing, and often implicit, operation of hegemonisation – of inclusion and exclusion – inherent in all political sites, which is the concern of disclosive ethics. When making this claim it is clear that for me ethics (with a small ‘e’) is not first and foremost ethical theory or moral reasoning about how we ought live (Caputo 1993). It is rather the ongoing questioning of the actual operation of hegemonic closure in which the interests of some become in/excluded as an implicit part of the material operation of relations of power – in codes, plans, programmes, technologies and the like. More particularly, I am concerned with the way in which the interests of some become excluded through the operation of closure, or hegemonisation, as an implicit and essential part of the *actual configuration and actual operation* of the ethico-political site of technology.

Before we proceed we must pause to consider an important and interesting question. If the design, implementation and use of technology are in fact political programmes why do we accept its hegemonic consequences so readily? It is not possible to address this question in full here. Nevertheless, a few comments can be made. First, I would suggest that we *delegate* decisions and actions to technology because it may be more convenient, necessary, or more morally desirable. For example we allow traffic lights to make decisions on who may go and who needs to stop because it is more convenient and it is necessary

(where will we find all the people to do it?) and it is morally more desirable (it treats every driver equally in some sense). As Latour (1992, p. 234) suggested: “we have been able to delegate to non-humans not only force but also values, duties and ethics”. Second, we are often unaware of what we have delegated, where, and for whom, as the delegation process is mostly an implicit outcome of ordinary pragmatic ‘technical’ decisions made as an implicit part of the design, implementation and use context – thus, hegemonisation is an implicit part of the ongoing co-constitution of technology in the world. In addition to this we must acknowledge that we always delegate more than we think since the multistable nature of artefacts means that they may become used in ways never anticipated by the designers or originators. Third, and most importantly, I think we do not realise that in delegating ‘values, duties and ethics’ to technology we are simultaneously changing the conditions that constituted us as the sort of beings that we are – in our contemporaneous way of being. In other words we are not aware of the co-constitutive relation that we have with technology. As such we can very easily see, understand and comprehend what we have gained (usefulness, efficiency, convenience, etc) but we find it difficult to comprehend the more or less subtle changes in the co-constitutive horizon of our ongoing way of being. This awareness often only emerges over longer historical epochs, as suggested by Heidegger (1977). The practical and the co-constitutive play themselves out in different horizons of intelligibility. The practical is in the urgency of getting things done (decisions made, technologies implemented), the co-constitutive emerges in the critical reflective historical interpretation of our ongoing way of becoming. Indeed bringing these two horizons together – without falling into essentialism or into a form of descriptive particularism—is the important challenge for us to make sense of the ethico-political site of technology. This is what disclosive ethics aims to contribute to.

As those concerned with disclosive ethics, we can see the operation of this ‘closure’ or ‘enclosure’ in many related ways. We can see it operating as already ‘closed’ from the start – where the voices (or interests) of some are shut out from the design process and use context from the start. We can also see it as an ongoing operation of ‘closing’ – where the possibilities for suggesting or requesting alternatives are progressively excluded. We can also see it as an ongoing operation of ‘enclosing’ – where the design decisions become progressively ‘black-boxed’ so as to be inaccessible for further scrutiny. And finally, we can see it as ‘enclosed’ in as much as the artefacts become subsumed into larger socio-technical net-

works from which it becomes difficult to ‘unentangle’ or scrutinise. Fundamental to all these senses of closure is “the event of closure [as] a delimitation which shows the double appartenance of an inside and an outside...” (Critchley 1999, p. 63).

Obviously at a certain level the development and design of technology is rather a pragmatic question. However, it is my contention that many seemingly pragmatic or technical decisions may have very important and profound consequences for those excluded – as I will show below. This is the important task of disclosive ethics. Not merely to look at this or that artefact but to trace all the moral implications (of closure) from what seems to be simple pragmatic or technical decisions – at the level of code, algorithms, and the like – through to social practices, and ultimately, to the production of particular social orders, rather than others. For disclosive ethics it is the way in which these seemingly pragmatic attempts at closing and enclosing connect together to deliver particular social orders that (ex)cludes some and not others – irrespective of whether this was intended by the designers, or not. Indeed it is my argument that in the ongoing design of complex socio-technical sites these political decisions – of (ex)inclusion or (en)disclosing of possibilities or intentions – often do not surface for consideration as such. They often emerge as technical or economic choices that ought to be subjected to technical or economic rationality. Furthermore, these political possibilities often emerge as a systemic effect or outcome where it is difficult to trace and locate a particular ‘author’ or designer who intended it as such. Thus, the task of disclosive ethics is, as Latour (2002) proposes, the “*reopening of the tombs in which automatisms have been heaped*” and to work for the maintenance of “*the reversibility of folding.*” (p. 258). The notion of reversibility should not be taken literally. It is not to ‘go back’ to a more original position. It is rather an ideal of putting into practice the conditions that will facilitate openness rather than closure. I will return to how one might do this in the last section. Disclosure always has the aim to make the ethics/politics of technology more or less explicit so that it might come up for scrutiny. I would suggest that this task is especially difficult with information technology.

Information technology and closure

Having argued that the co-constitutive nexus of human and technology relationships is a political site, I now want to claim that the political site of information technology (in the form of software algorithms) is, in a sense, of a different order

(Graham and Wood 2003). I want to contend that scrutinising information technology is particularly problematic since information technology, in particular algorithms, is what I would term an *opaque* technology as opposed to a *transparent* technology (Introna 1998). Obviously I do not see this distinction as a dichotomy but rather as a continuum. As an attempt to draw this distinction some aspects are highlighted in Table 1 below.

Facial recognition algorithms are a particularly good example of a opaque technology (Introna and Wood 2004). The facial recognition capability can be imbedded into existing CCTV networks, making its operation impossible to detect. Furthermore, it is passive in its operation. It requires no participation or consent from its targets – it is “non-intrusive, contact-free process” (Woodward et al. 2003, p. 7). Its application is flexible. It can as easily be used by a supermarket to monitor potential shoplifters (as was proposed and later abandoned, by the Borders bookstore), by casinos to track potential fraudsters, or used for identifying ‘terrorists’ at airports (as is currently in operation at various US airports). However, most important of all is the obscurity of its operation.

Most of the software algorithms at the heart of information technology such as facial recognition systems are propriety software objects. Thus, it is very difficult to get access to them for inspection and scrutiny (Introna and Wood 2004). More specifically, however, even if you can go through the code line by line, it is almost impossible to inspect that code *in operation*, as it becomes implemented through multiple layers of translation for its execution. At the most basic level we have electric currents flowing through silicon chips, at the highest level we have programme instructions, yet it is almost impossible to trace the connection between these as it is being executed. Thus, it is virtually impossible to know if the code you inspected is the code being executed,

when executed. In short, *software algorithms are operationally obscure*.

It is my argument that the opaque and ‘silent’ nature of digital technology makes it particularly difficult for society to scrutinise it. Furthermore, this inability to scrutinise creates unprecedented opportunities for this silent and ‘invisible’ micro-politics to become pervasive (Graham and Wood 2003). *Thus, a profound sort of micro-politics can emerge as these opaque (closed) algorithms become enclosed in the social-technical infrastructure of everyday life*. Paradoxically, we tend to have extensive community consultation and impact studies when we build a new motorway. However, we tend not to do this when we install CCTV in public places or when we install facial recognition systems in public spaces such as airports, shopping malls, etc. Most informed people tend to understand the cost (economic, personal, social, environmental) of more transparent technologies such as a motorway, or a motorcar, or maybe even cloning. However, I would argue that they do not seem to understand the ‘cost’ of the more opaque information technologies that increasingly pervade our everyday life – maybe it is because these technologies have been ‘closed’ from the start. In the next section I will discuss and disclose two such technologies: search engines and plagiarism detection systems.

Disclosing the politics of information technology: two case examples

Search engines

A detailed discussion of the political site of search engines is beyond the scope of this paper. I will just draw a few rough outlines to indicate some of the political/ethical strategies that emerge in the search engine site (refer to Introna and Nissenbaum (2000) for a more detailed discussion).

No search engine can cover the entire internet, not even the entire publicly accessible web.⁵ In 1999 Lawrence and Giles (1999) estimated that the web consists of 800 million unique pages and that the best search engines indexed no more than 16% of that content. In 2005 a study by Gulli and Signorini estimated the size of the web at 11.5 billion pages. At the same time Google claimed that its database/index covers 8.1 billion pages, which would give it a 70%

Table 1. Opaque versus transparent technology

<i>Opaque</i> technology is	<i>Transparent</i> technology is
Embedded/hidden	On the ‘surface’/conspicuous
Passive operation (limited user involvement, often automatic)	Active operation (fair user involvement, often manual)
Application flexibility (open ended)	Application stability (firm)
Obscure in its operation/outcome	Transparent in its operation/outcome
Mobile (soft-ware)	Located (hard-ware)

⁵ This is because they need to index every word on a page (excluding stop words) since every word could potentially be a keyword that would be entered by a user. Thus a page with 1000 words will generate a thousand records in Google’s index.

coverage. However there has been a heated debate over these claims. For example it has been claimed by the editor of searchenginewatch that “Google, has sometimes included what they call ‘partially-indexed’ pages or what would more fairly be called link-only pages. These are pages Google knows about solely by links pointing at them. Nothing on the pages themselves has been indexed.” Others have estimated that Google’s coverage is somewhere between 50 and 60% of the publicly available web.⁶ Nonetheless, the point I want to make is that search engines, now and for the foreseeable future, will have to choose which site (or pages) they include and which they exclude from their index. Furthermore, one could claim, without much exaggeration, that for a website to exist it has to be in the index of a search engine, probably more specifically it has to be in Google’s index.⁷ A key question therefore is to know how Google decides which sites/pages to include and which to exclude? More specifically, what are the criteria being used by the Googlebot crawler or spider to target sites (or pages) for indexing. We could ask them but they will obviously not tell us as this is the intellectual property upon which their business depends. Thus we cannot inspect the code to determine these criteria as such.

Becoming indexed

We have discerned something of the nature of the Google spider algorithms from a paper on efficient crawling by Cho et al. presented at the WWW7 conference in 1998 (when Page was still a PhD student). This paper, which discusses commonly used metrics for determining the ‘importance’ of a webpage by crawling spiders, provides key insights relevant to my central claims. Because of its significance, I discuss it here in some detail. Cho et al write:

“Given a webpage P , we can define the importance of the page, $I(P)$, in one of the following ways...:

1. *Backlink Count*. The value of $I(P)$ is the number of links to P that appear over the entire web. We use $IB(P)$ to refer to this importance metric. *Intuitively, a page P that is linked to by many pages is more important than one that is seldom referenced. On the web, $IB(P)$ is useful for ranking query results, giving end-users pages that are more likely to be of general interest.* Note that evaluating $IB(P)$ requires

⁶ The publicly available web does not include pages protected by passwords. It has been estimated that the size of the ‘deep web’, as it is sometimes called, is 500 billion pages. Obviously such claims will be based on quite a bit of informed speculation.

⁷ Sixty percent (60%) of all searches in the USA is done through Google according to searchenginewatch.

counting backlinks over the entire web. A crawler may estimate this value with $IB'(P)$, the number of links to P that have been seen so far.

2. *PageRank*. The $IB(P)$ metric treats all links equally. Thus, a link from the Yahoo! home page counts the same as a link from some individual’s home page. However, since the Yahoo! home page is more important (it has a much higher IB count), it would make sense to value that link more highly. The PageRank backlink metric, $IR(P)$, recursively defines the importance of a page to be the weighted sum of the backlinks to it. Such a metric has been found to be very useful in ranking results of user queries. We use $IR'(P)$ for the estimated value of $IR(P)$ when we have only a subset of pages available.
3. *Location Metric*. The $IL(P)$ importance of page P is a function of its location, not of its contents. If URL u leads to P , then $IL(P)$ is a function of u . For example, URLs ending with “.com” may be deemed more useful than URLs with other endings, or URL containing the string “home” may be more of interest than other URLs. Another location metric that is sometimes used considers URLs with fewer slashes more useful than those with more slashes. All these examples are local metrics since they can be evaluated simply by looking at the URLs.” (1998, p. add in, emphasis added)

The “Backlink” metric uses the backlink (or inlink) count as its ‘importance’ heuristic. The value of the backlink count is the number of links to the page that appear over the entire Web – for example the number of links over the entire Web that refers to <http://www.ibm.com>. The assumption here is that “a page that is linked to by many [other] pages is more important than one that is seldom referenced.” Obviously, this is a very reasonable heuristic. We know from academic research that it is wise to look at the ‘canonical’ works that is referred to – or cited in academic language – by many other authors. We know also, however, that not all topics necessarily have canons. Furthermore, whereas in some less mainstream fields, a small number of citations may make a particular work a canon, in other fields, it takes a vast number of citations to reach canonical status. Thus, the Backlink heuristic would tend to crawl and gather sites/pages and links for areas with a large potentially interested population (for example, topics/fields such as ‘shareware computer games’ since an even relatively unimportant site in this large field will be seen as more ‘important’ – have relatively more backlinks or inlinks– than an actually important site in a small field (such as ‘the local community services information’ page) which would have relatively less backlinks or inlinks because it has a smaller

potentially interested population. The essential point is that the large areas of interest determine the measure, or threshold, of ‘importance’ – through sheer volume of backlinks – in ways that would tend to push out the equally important small fields of interest.

With the “PageRank” metric this problem is exacerbated. Instead of treating all links equally, this heuristic gives prominence to back links from other ‘important’ pages – pages with high backlink counts. Thus, “ since [a link from] the Yahoo! home page is more important (it has a much higher IB [backlink] count), it would make sense to value that link more highly.” In the analogy to academic papers, a metric like this would imply that a particular paper is even more important if referred to by others whom are already seen as important – by other canons. More simply, you are important if others who are already seen as important indicate that you are important. The problem with the Backlink and PageRank metrics is that they assume that backlinks are a reliable indication of importance or relevance. In those cases where authors of pages create links to other pages they see as valuable this assumption may be true. There are, however, many organizations that actively cultivate backlinks by inducing webpage creators to add a link to their page through incentives such as discounts on products, free software utilities, access to exclusive information, and so forth. Obviously not all web pages creators have equal access to the resources and expertise to push their backlink counts up.

The “Location Metric” uses location information from the URL to determine ‘next steps’ in the crawl. “For example, URLs ending with “.com” may be deemed more useful than URLs with other endings, or URL containing the string “home” may be more of interest than other URLs.” Even though the authors do not indicate what they see as more important, one can assume that these decision are made when crawl heuristics are set for a particular spider. It may therefore be of great significance ‘where you are located’ as to how important you are seen to be. With the URL as the basis of decision making many things can aid you in catching the attention of the crawling spider, such as having the right domain name, being located in the root directory, and so forth. From this discussion on crawling metrics we can conclude that pages with many backlinks, especially backlinks from other pages with a high backlink counts, which are located at locations seen as ‘useful’ or ‘important’ to the crawling spider will become targets for harvesting. This is obviously a positive feedback loop which means that because a site is already in the index people will tend to find it and link to it which will make it even more difficult for those already left out to become included.

Getting indexed is only the first challenge. Once you are in the index the next challenge is to get onto the first (or second) page when results are returned to the user since users often limit their choices to these results.⁸ It is therefore important for website owners to know how Google makes decisions about the ranking of search results. I do not have the space here to deal with this matter in any detail. Nonetheless, I might just mention, as indicated above, that backlink count also plays an important role in ranking with the same implications as above. Furthermore, the criteria for ranking are also enclosed in proprietary software and subject to periodic changes, which could have disastrous consequences for a small online retailer who might suddenly disappear out of the first two pages for no apparent reason. There is a lot more I can say about ranking but space limitations preclude it – refer to Introna and Nissenbaum (2000). To summarise: search engines, through their undisclosed algorithms, constitute the conditions that make some websites/pages attractive or visible and others not. Users unwittingly contribute to this constitution by implicitly accepting this ‘as the way it works’ and thereby reinforcing it through their search and linking behaviour. Through this nexus of co-constitutive relations a particular world wide web is unwittingly being constructed that (in)cludes the interest of some and not others. Increasingly we have a web for the majority at the expense of the minority. Ironically, it was the vision of Tim Berners-Lee that the web would exactly be the place where even the smallest individuals, groups and organisations could become visible, interact and realise their potential. Through the logic of search engine algorithms the opposite is happening. Let us now turn to plagiarism detection systems.

Plagiarism detection systems

Plagiarism detection systems (PDS) are increasingly used by universities as a response to the belief that plagiarism is on the increase (Lathrop 2000). The market leader, Turnitin, claims that their system is used by 5000 institutions in 80 countries worldwide (covering 12 million students and educators) and that 50,000 papers get submitted to their system every day. They also claim that their crawler Turnitinbot has downloaded over 9.5 billion Internet pages to

⁸ A study of travel agents using computerized airline reservations systems, which showed an overwhelming likelihood that they would select a flight from the first screenful of search results, is suggestive of what we might expect among Web users at large (Friedman and Nissenbaum, *Bias in Computer Systems. ACM Transactions on Information Systems*, 14(3): 330–347, 1996).

their detection database and that it updates itself at a rate of 60 million pages per day (Turnitin website). More recently academic publishers have also turned to Turnitin to help them protect themselves from publishing plagiarised material, which is obviously very damaging to their reputation (and profits one might add). This seemingly success story of Turnitin needs unpacking, which cannot be done here in full. Nevertheless, one of the most powerful arguments often put forward for adopting it (beyond resource constraints) is that it ‘levels the playing field’, indeed, that it is more fair than the hit and miss approach where individual teachers have to spot cases of plagiarism. The argument is made that teacher-based monitoring tends to pick up weak students or non-native speakers because of the obvious shift in sophistication when a piece of plagiarised text is found imbedded in an assessment document such as an essay or dissertation. But is it levelling the playing field or does it rather reconstitute a playing field that is even more uneven? I would argue that it is the latter. Moreover, that this is a much more serious issue since many of the important conditions (affordances) are now embedded in proprietary systems which are not open for scrutiny – an invisible micro-politics one might say. On a more profound level the increased pervasiveness of Turnitin – not only in academic contexts – might emerge as a site where the very notion of authorship, what it means, how it is attributed, and for what purposes may become contested. Nonetheless, let us consider some of the more mundane constitutive conditions in which Turnitin functions as a mechanism for detecting ‘plagiarism’.

If it is true that Turnitin covers almost all (if not all) of the web then anybody taking something from the web has an equal chance of being detected and that would most certainly be fair, a level playing field. However, what if Turnitin does not cover the entire web? In such a case the likelihood of somebody being detected would depend on whether they happen to take something from a place that Turnitin did (or did not cover). If Turnitin’s claim that they cover 9.5 billion pages is true and the estimate that the web consists of 11.5 billion pages is correct (which would give them 83.6% coverage) then one could argue that there is a relatively high probability that a student will be detected if they take something from the web. However these figures are misleading because a lot of the content that Turnitin needs to cover is in fact behind passwords (i.e. in the deep web), such as academic journals for example. In a small scale experiment we selected 103 fragments from a number of likely sources where students may take material from – in the publicly available as well as the deep web – and submitted it to Turnitin. Turnitin was only

able to detect 47 of these, a detection rate of 45.6%. There are a variety of reasons why web content may not be in Turnitin’s database. Some of the main reasons are:

- That the content is behind passwords, i.e. the deep web.
- That it is in image format (many papers in journals before 1999/2000 are).
- That it is new content, i.e. has not yet been picked up by the Turnitin.bot (we estimate the refresh rate of Turnitin to be approx 4–6 months).
- That webmasters used the robot exclusion standard to exclude the Turnitin.bot – there is evidence that this is increasingly happening because the Turnitin.bot is particularly aggressive in its requests.

If these results are to some extent generalizable (we are not claiming at this stage) then a student taking something from the web has less than 50% chance of being detected, which is quite low. My problem is not that some are caught and some get away, as it were. I am rather more concerned with the mostly implicit assumption of teachers when they interpret the Turnitin results. They often believe that those that are not detected by Turnitin are innocent and those that are detected are guilty. I would suggest that both of these assumptions are wrong or could be wrong. The first is partly wrong because of the partial coverage of Turnitin as suggested by our experiments. The second one is often wrong for more subtle and complex reasons, related to the operation of the algorithm and its interaction with writing practices, which I now want to turn to.

I must first say that plagiarism detection software – contrary to what it name suggests – detects *copies not plagiarism*. How does it detect copies? A simple approach would be to compare a document character by character. However, this approach has a number of problems: (a) it is very time-consuming and resource intensive; (b) it is not sensitive to white spaces, formatting and sequencing changes; and (c) it cannot detect part copies from multiple sources. To deal with these problems a number of algorithms have been developed. Unfortunately many of these (such as Turnitin) are now proprietary software and therefore not available for analysis and scrutiny – black boxes in which it has become almost impossible to maintain the reversibility of foldings. However, I have studied the logic of certain published algorithms, such as winnowing (Schleimer et al. 2003), as well as doing some preliminary experimental research of the way the Turnitin algorithm seems to behave. From these we are able to draw some important conclusions, which I will discuss below.

All detection algorithms operate on the basis of creating a digital fingerprint of a document which it then uses to compare documents against each other. The fingerprint is a small and compact representation of the content of the document that can serve as a basis for determining correspondence between two documents (or parts of it). In simple terms the algorithm first removes all white spaces as well as formatting details from the document to create one long string of characters. This often results in a 70% reduction of the size of the document. Further processing is done to make sure that sequences of consecutive groups of characters are retained and converted through a hash function⁹ to produce unique numerical representations for each group of characters. The algorithm then takes a statistical sample from this set of unique numerical strings in such a way as to ensure that it always covers a certain amount of consecutive characters (or words in our human terms) and stores this as the document's fingerprint.¹⁰ A fingerprint can be as small as 0.54% of the size of the original document.

From this very limited description of the algorithm it is clear that the detection algorithm is dependent on certain characteristics of the copied text to remain intact for detection to be possible. In some cases a small amount of change in the right way (or place) will make a copy undetectable and in other cases a large amount of changing will still make it possible to detect. One of the key requirements for detection is that a *sufficiently long string of consecutive characters* from the original is retained in the copied version. The location, within the fragment, of the consecutive string is also important. For example in experiments we did with Turnitin it became clear that if one would change one word in a sentence at the right place – often between the 6th to 14th word in the sentence – then Turnitin did not recognise it even if all the rest of the sentence remained exactly the same. Indeed we were also able to submit a fragment of 300 words where we changed approximately every 7th to 10th word and remain undetected. In contrast a small fragment of 26 consecutive unchanged words were detected by Turnitin. Given this behaviour of the algorithm it is

possible for a student to incorporate large amounts of copied material by intentionally or unintentionally changing words in the right places in the text submitted and remain undetected. Now my concern here is not to suggest ways that students might cheat. My concern is rather the way this behaviour of the algorithm might constitute an uneven playing field, especially for non-native speakers.

We know that non-native speakers learn to write by using fragments as 'patches' to imitate the vocabulary and structure of expressions as part of their transition to become competent in academic writing (Howard 1993, 1995; Shi 2004; Leki and Carson 1997). This is true not only for non-native speakers, it is also true for native-speaking academics when paraphrasing a difficult-to-understand text—even material within their own discipline. Roig (2001), in a fascinating study, provided college professors in psychology (all members of the American Psychological Society) with two different texts to paraphrase: the first was a difficult text from a peer-reviewed psychology journal article and the second was an easy-to-read text from an introduction-level psychology textbook. Twenty-six percent (26%) of the professors appropriated text – strings of five words in length or more *without quotation marks* – from the original text, whereas only three percent (3%) appropriated text from the piece that was easier to read. If psychology professors – and most probably native speaking students – feel the need to 'stay close' to the text when confronted with difficult material, we can see why, students who understand the importance of 'speaking' like the teachers and the people they read, do the same when it comes to doing their assessments. We also know that it is possible to use phrases and fragments from a text to say something completely different than that which the original author has said. Nevertheless, this is not my concern here; rather, my claim is that non-native speakers will tend to use larger fragments of consecutive words, for fear of losing the meaning, than native speakers. Furthermore, native speakers will tend to have the vocabulary and linguistic skills to make changes to the fragments without a loss of meaning – especially in the middle of sentences where it really matters from a detection point of view. Thus, it is my claim that non-native speakers who appropriate fragments as part of their writing practices will be disproportionately detected as opposed to native speakers. I will furthermore claim that the non-native speaker might have more legitimate reasons to have fragments in the text as opposed to the native speaker – although that may be a controversial point of view. I do not want to go too far in my analysis here as it is still based on preliminary work.

⁹ A more technical definition of hash function is "A hash function is a function that converts an input from a (typically) large domain [input values] into an output in a (typically) smaller range (the *hash value*, often a subset of the integers) (from http://www.en.wikipedia.org/wiki/Hash_function).

¹⁰ Refer to Schleimer, Wilkerson, Aiken, Winnowing: Local Algorithms for Document Fingerprinting. Proceedings of the ACM SIGMOD International Conference on Management of Data, June 2003, 76–85, 2003 for a more detailed discussion.

In summary: my suggestion is that the large-scale use of Turnitin may be creating a set of constitutive conditions in which some students are being constituted as ‘plagiarists’ and others not in an unfair uneven playing field. Most importantly, and quite ironically, most of the teaching staff that use Turnitin are not aware of this and are contributing to this unfair constitution of plagiarists in a sincere effort to be fair. Let us now turn to some general conclusions and observations about the morality of technology and disclosive ethics.

Some comments on the ethics of (dis)closure

From our discussion of search engines and plagiarism detection systems above it is clear that the constitutive conditions that constitutes some sites/pages as ‘attractive’ (and others not) and some students as ‘plagiarists’ (and others not) are *not* simply properties of software objects, but they are also not properties of the humans either. Indeed there is a fundamental simultaneity (or hybridity) of agency at play in this nexus of relationships – in which sites become attractive and students become plagiarists – which makes it is very difficult to locate intentionality and properties in any definitive way. We cannot say that the designers of Google had an intentional strategy of hegemonisation to make the ‘big guys’ push out the ‘small guys’. Nor can we say that the designers of Turnitin intended to discriminate against non-native speakers. The material agency of their code is but one element in the nexus of constitutive relationships. There are a multitude of other intentions and agencies at work that shapes the hegemonisation of the ethico-political site in ways that transcend (even pervert) the intentions and agencies of individual actors (or actants in Latour’s language). Likewise, we cannot simply say that the software objects are neutral means and it is the people who use them that are at fault, or that they simply use them in an inappropriate ways. Of course some of that is true, however, the software objects do embody certain (im)possibilities, (dis)functions, affordances/prohibitions that condition the way they are taken up as part of ongoing social practices (in searching and detecting). The politics and ethics of technology is diffused and multiple. That is why we have a moral obligation to *disclose* them on an ongoing basis, to maintain the reversibility of foldings, in Latour’s words. How and when – in concrete terms – must this disclosure happen? Let me briefly mention some ideas.

(a) *Disclosive ethical archaeology*

Technological sites needs to be subject to ongoing disclosive scrutiny through a process of disclosive

archaeology as was done with search engines, plagiarism detection systems above—and others such as ATMs (Introna and Whittaker, 2006), facial recognition systems (Introna and Wood 2004; Brey 2004) and virtual reality computer games (Brey 1999), to name but a few. When I use the term archaeology here I am thinking of Foucault’s work – i.e. the conditions that rendered an event or situation possible. As he explains:

... it is rather an enquiry whose aim is to rediscover on what basis knowledge and theory [technology in our case] became possible; within what space of order knowledge [technology] is constituted... Such an enterprise is not so much a history, in the traditional meaning of the word, as an “archaeology” (Foucault 1994, xxi–xxii)

The purpose of disclosive archaeology is not to focus on material agency or human agency as such but rather to make visible the ongoing conditions of possibility that they co-constitute in the world where they function as such. It must trace the contingent simultaneity of intentions, decisions, affordances, interpretations, uses, codes, programmes, etc. to reveal the nexus that co-constitutes the ethico-political site of technology. I would suggest that when funding is given for research and development, in for example nano technology, a simultaneous grant must be given for ethicists of technology to scrutinise and make visible the choices being made by the scientists, engineers and designers. Disclosive ethical archaeology should be an ongoing practice that is inherently part of the ongoing coming into being of the technological site.

(b) *Transparent design*

With transparent design I mean two things. First, transparent design means the opening up of the design and development (and implementation) activity to multiple stakeholders for ongoing scrutiny and debate. Examples of this might be participative design as well as participative technology assessment. However, it is important that these participative activities be understood within the co-constitutive horizon of the participation itself and within the relations of power that makes all such activities of participation possible and meaningful. Second, and very importantly, designing technology in such a way that it is relatively transparent in its operation – i.e. that it is possible for ordinary informed users to understand the (un)intentions, (im)possibilities, (dis)functions, affordances/prohibitions of the artefacts that constitute their way of being. The free software movement and the open source movement

could be seen as examples of transparent design. It is always possible for an informed user to ‘read’ the code, and even amend it if they so wish. Increasingly more and more of the artefacts that we buy are blackboxed to the degree that they are not accessible even to expert users. This problem becomes even more pronounced in systems based on neural technology where even the experts might be surprised by the behaviour of their own artefacts.

(c) *Engaging and multistable design*

I think it is important, as Verbeek (2005) and Borgmann (1984) suggested, that we build technologies that engage us so that they do not become ‘blackboxed’ and subsumed into the ‘background’ in such a way as to disappear from our attention and scrutiny. The development of the camera may be a useful example. The creation of a good photograph is the result of the interaction between many variables such as the type of film (ISO), shutter speed, aperture setting, composition, etc. The ability to vary these elements through various settings on the camera allows the photographer to develop an intimate relationship with a particular camera to such a degree that they insist on using that particular camera (not just the model and the make) when they take photographs. They care for their camera as a unique instance that is the necessary condition for them to be constituted as a good photographer. On the opposite spectrum we have the disposable camera in which all elements are preset and blackboxed. It leaves very little room for engagement. The camera requires only that you aim and shoot. It was designed to be disposed of and this is what you do once it is full. It constitutes you as ignorant and you constitute it as disposable (Introna 2003). This may be a convenient way of ‘getting the picture’ but through this delegation to the camera it proceeds to enrol the person into a whole set of intentions and programmes that are mostly unsaid such as having to get it developed in a certain way, having to buy different cameras for different situations, etc. More engaging artefacts have higher levels of multistability. In other words these technologies afford multiple interpretations – and ways of doing – to the users. This means that users can interpret and use these artefacts in multiple ways that encourage an active engagement with the artefact, as is the case with a manual camera. One could say that the higher the level of engagement the less likely it will be for the user to become unwittingly enrolled into political programmes not of their choosing. Unfortunately this

engagement is often sacrificed or exchanged, unsuspectingly, for the benefit of ‘convenience’.

(d) *Materialising morality*

If technology is indeed society made durable (as suggested by Latour 1991) and argued above then we can make some of our values more durable by explicitly building some of these values into the technological site (Latour 1992; Achterhuis 1995; Verbeek 2006). For example, instead of cautioning users to be careful in online transactions we can make them safer. Obviously, there are many arguments against such approach (such as issues of freedom and autonomy of the user). However, the fact is that we are already explicitly and implicitly building values into technological sites and this is now undisclosed. Seemingly ‘technical’ choices are being made without scrutiny that may have a profound affect on our future possibilities. In moralising technology we must ensure that this technology is transparent and open to ongoing scrutiny. Nevertheless, delegating to artefacts some of our moral responsibility might be good if it happens in the context of disclosure, transparency and engagement suggested above.

These suggestions are not complete, unproblematic or uncontroversial. Nevertheless, they do go some way in taking the ethics of technology seriously. They are steps towards maintaining the reversibility of foldings, which is our very important moral concern when it comes to our encounter with technology—as Latour suggested.

References

- H. Achterhuis. De moralisering van de apparaten. *Socialisme en Democratie*, 52(1): 3–12, 1995.
- A. Berg and M. Lie. Feminism and Constructivism: Do Artifacts Have Gender?. *Science, Technology and Human Values*, 20(3): 332–351, 1995.
- W.E. Bijker, *Of Bicycles, Bakelites and Bulbs Toward a Theory of Sociotechnical Change*. MIT Press, Cambridge, MA, 1995.
- W.E. Bijker, T.P. Hughes and T.J. Pinch, *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. MIT Press, Cambridge, MA, 1987.
- A. Borgmann, *Technology and the Character of Contemporary Life*. Chicago University Press, Chicago, IL, 1984.
- G.C. Bowker and S.L. Star, *Sorting Things Out: Classification and its Consequences*. MIT Press, Cambridge, MA, 1999.
- P. Brey. Philosophy of Technology meets Social Constructivism, Techné. *Journal of the Society for Philosophy and Technology*, 2(3/4): 56–79, 1997.

- P. Brey. The Ethics of Representation and Action in Virtual Reality. *Ethics and Information Technology*, 1(1): 5–14, 1999.
- P. Brey. Ethical Aspects of Face Recognition Systems in Public Places. *Journal of Information, Communication and Ethics in Society*, 2(2): 97–109, 2004.
- M. Callon. Society in the Making: The Study of Technology as a Tool for Sociological Analysis. In W.E. Bijker, T.P. Hughes and T.J. Pinch, editors, *The Social Construction of Technological Systems*, pp. 83–103. MIT Press, Cambridge, MA, 1987.
- J.D. Caputo, *Against Ethics*. Indiana University Press, Indianapolis, 1993.
- J. Cho, H. Garcia-Molina, and L. Page. Efficient crawling through URL ordering, Paper presented at the Seventh International World Wide Web Conference, Brisbane, Australia, 14–18 April 1998.
- T. Conley. Folds and Folding. In C.J. Stivale, editor, *Gilles Deleuze: Key Concepts*, pp. 170–181. McGill-Queen's University Press, Montreal, 2005.
- S. Critchley, *The Ethics of Deconstruction: Derrida and Levinas*. 2 ed. Edinburgh University Press, Edinburgh, 1999.
- S. Critchley. Is There a Normative Deficit in the Theory of Hegemony?. In S. Critchley and O. Marchant, editors, *Laclau: A Critical Reader*, pp. 113–122. Routledge, London, 2004.
- M. Foucault, *The Order of Things: An Archaeology of the Human Sciences*. Routledge, London, 1994.
- B. Friedman and H. Nissenbaum. Bias in Computer Systems. *ACM Transactions on Information Systems*, 14(3): 330–347, 1996.
- S. Graham and D. Wood. Digitizing Surveillance: Categorization, Space and inequality. *Critical Social Policy*, 20(2): 227–248, 2003.
- M. Heidegger, *Being and Time*, trans John Macquarrie and Edward Robinson. New York, Harper and Row (1927/1962).
- M. Heidegger, *The Question Concerning Technology and Other Essays*. Harper Torchbooks, New York, 1977.
- R.M. Howard. Plagiarisms, Authorships, and the Academic Death Penalty. *College English*, 57(1): 788–805, 1995.
- R.M. Howard. A Plagiarism Pentimento. *Journal of Teaching Writing*, 11(2): 233–245, 1993.
- L.D. Introna. Oppression, Resistance and Information Technology: Some Thoughts on Design and Values, Design for Values. Ethical, Social and Political Dimensions of Information Technology workshop sponsored by the NSF DIMACS, Princeton University, USA, 27 February to 1 March, 1998.
- L.D. Introna. On the Ethics of (Object) Things. paper presented at the Critical Management Conference 3, Lancaster University, England, 17–19 July 2004. Available at <http://www.lums.lancs.ac.uk/publications/viewpdf/000235/> (Last accessed on 18/07/2006).
- L.D. Introna and F.M. Ilharco. The Meaning of Screens: Towards a Phenomenological Account of Screenness. *Human Studies*, 29(1): 57–76, 2006.
- L.D. Introna and H. Nissenbaum. The Internet as a Democratic Medium: Why the Politics of Search Engines Matters. *The Information Society*, 16(3): 169–85, 2000.
- L.D. Introna and L. Whittaker. Power, Cash and Convenience: Translations in the Political Site of the ATM. *The Information Society*, 22(5): 325–340, 2006.
- L.D. Introna and D. Wood. Picturing Algorithmic Surveillance: The Politics of Facial Recognition Systems. *Surveillance and Society*, 2(2/3): 177–98, 2004.
- A. Lathrop, *Student Cheating and Plagiarism in the Internet Era: A Wake-up Call*. Englewood CO, Libraries Unlimited, 2000.
- B. Latour. Where are the Missing Masses? Sociology of a Few Mundane Artefacts. In W. Bijker and J. Law, editors, *Shaping Technology, Building Society: Studies in Sociotechnical Change*, pp. 225–258. MIT Press, Cambridge, Mass, 1992.
- B. Latour, *Pandora's Hope Essays on the Reality of Science Studies*. Harvard University Press, Cambridge, Mass, 1999.
- B. Latour. The Promise of Constructivism. In D. Ihde and E. Selinger, editors, *Chasing Technoscience: Matrix for Materiality*, pp. 27–46. Bloomington and Indianapolis, Indiana University Press, 2003.
- B. Latour. Technology is Society made Durable. In J. Law, editor, *A Sociology of Monsters: Essays on Power*, pp. 103–31. Technology and Domination, London, Routledge, 1991.
- B. Latour. Morality and Technology: The End of the Means. *Theory, Culture and Society*, 19(5/6): 247–60, 2002.
- J. Law, *The Sociology of Monsters: Essays on Power, Technology and Domination*, London, Routledge, 1991.
- S. Lawrence and C.L. Giles, Accessibility and Distribution of Information on the Web, 1999.
- I. Leki and J. Carson. Completely Different Worlds: EAP and the Writing Experiences of ESL Students in University Courses. *TESOL Quarterly*, 31(1): 39–69, 1997.
- D.A. Norman. *The Psychology of Every day Things*. Basic Books, New York, NY, 1988.
- B. Pfaffenberger. Technological Dramas. *Science, Technology, and Human Values*, 17 282–312, 1992.
- M. Roig. Plagiarism and Paraphrasing Criteria of College and University Professors. *Ethics and Behavior*, 11(3): 307–323, 2001.
- T.R. Schatzki. The Sites of Organizations. *Organization Studies*, 26(3): 465–484, 2005.
- T.R. Schatzki, *The Site of the Social. A Philosophical Account of the Constitution of Social Life and Change*. University Park Penn, Pennsylvania University Press, 2002.
- S. Schleimer, D. Wilkerson and A. Aiken (2003) Winnowing: Local Algorithms for Document Fingerprinting. Proceedings of the ACM SIGMOD International Conference on Management of Data, June 2003, 76–85.
- L. Shi. Textual Borrowing in Second-Language Writing. *Written Communication*, 21(2): 171–200, 2004.
- S. Sismondo. Some Social Constructions. *Social Studies of Science*, 23(3): 515–553, 1993.

- P.P. Verbeek, *What Things Do – Philosophical Reflections on Technology, Agency, and Design*. Penn State University Press, Penn State, 2005.
- P.P. Verbeek. Materializing Morality – Design Ethics and Technological Mediation. *Science, Technology and Human Values*, 31(3): 361–380, 2006.
- L. Winner. Upon Opening the Black Box and Finding it Empty: Social Constructivism and the Philosophy of Technology. *Science, Technology, and Human Values*, 18(3): 362–78, 1993.
- J. Woodward, C. Horn, J. Gatune and A. Thomas. Biometrics: A Look at Facial Recognition. Documented Briefing prepared for the Virginia State Crime Commission (2003). Available at <http://www.rand.org>.