Middlesex University Research Repository

An open access repository of

Middlesex University research

<http://eprints.mdx.ac.uk>

Jones, Simon ORCID logo ORCID: <https://orcid.org/0000-0002-6405-5698> (2016) Doing the right thing: computer ethics pedagogy revisited. Journal of Information, Communication and Ethics in Society, 14 (1) . pp. 33-48. ISSN 1477-996X [Article]
(doi:10.1108/JICES-07-2014-0033)

Final accepted version (with author's formatting)

This version is available at: <https://eprints.mdx.ac.uk/19166/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

Doing the Right Thing: Computer Ethics Pedagogy Revisited

Simon Jones

Department of Computer Science
Middlesex University
London, UK

It is now widely acknowledged that students of computing, as prospective professionals, need to be aware of the broader issues surrounding the design, deployment and use of information and communication technologies (ICTs). Courses in ethical, social and professional issues have consequently been integrated into computer science, and related degrees, for some time. A common framework for teaching and analysing computer ethics has emerged, and can be identified in various key texts which offer similar pedagogical approaches and address common topics (Kallman and Grillo, 1996; Forester and Morrison, 1998; Rosenberg, 2004; Stamatellos, 2007; Adams and McCrindle, 2008; Baase, 2009; Johnson, 2009; Quinn, 2005; Kizza 2010, Spinello, 2011; Tavani, 2010; Reynolds, 2012). This paper critically reviews key aspects of this framework from four main angles; ethical theories, professional codes of conduct, the law, and social issues. Some of the limitations of these approaches are considered. A revised framework is then outlined, and an alternative methodology proposed for *doing* ethics in a teaching context. This approach draws on the practical experience of teaching computer ethics to a culturally diverse body of computing undergraduates, both in the UK and at several international sites.

1. Reviewing the Standard Framework

Ethical theory occupies a central place in computer ethics. The main strands of Western moral philosophy are cited in most textbooks in the field. These include social contract and rights-based theories, virtue ethics, duty-based deontological theories and teleological approaches, such as consequentialism and utilitarianism. Many aspects of these ethical theories continue to have pedagogical value. The importance of substantiating claims through an appeal to rational argument is a fundamental requirement of any ethical evaluation. Kant's notion of moral obligation, of doing "good things" for intrinsic motivations because they are good in themselves, touches on a vital altruistic principle. Similarly, Kant's categorical imperative that an action is valid only if everyone else believes that one should act this way, touches on core principles of fairness and of not making an exception for oneself, as encapsulated in the question "what if everyone did it?" Social contract theory is useful in arbitrating disputes where different cultural standards or moralities clash, whereas virtue ethics' focus on the character traits of individuals as moral agents is particularly relevant to professional behaviour.

Each of these ethical theories, however, has its own particular limitations, and these are acknowledged in much of the existing literature. Contract-based theories tend to be minimalist in demanding only that citizens do no harm, without any obligation to proactively do good (Tavani, 2010, p. 63). Kantian theory's universal imperatives tend towards a certain moral absolutism through their inability to accommodate exceptional circumstances where "wrong" actions may result in "good" consequences, unintended or otherwise (Quinn, 2005, p. 67). Consequentialist theories, conversely, make it possible to justify *any* action, however unethical, in the name of a desired goal. Utilitarian approaches, where the goodness of actions is measured in terms of their maximisation of net happiness for the greatest number, can result in surprisingly inhuman conclusions, particularly where the happiness of the majority supersedes that of the minority (Johnson, 2009, p. 40). A further problem, identified by some authors, is that the privileging of these specific strands of western ethical theory is implicitly exclusive of other, political and moral philosophies, particularly those that are non-western in origin (Duquenoy, 2007, p.8). This is a limitation for teaching ethics in a culturally diverse, globally

dispersed learning environment, and for addressing the ethical challenges of a 21st century digital society (Ess and Thorseth, 2010).

While each of the major ethical theories has its respective merits, arguably, no single theory, on its own, is adequate for tackling all the complex modern ethical dilemmas raised by computer technologies. A range of theories and approaches are needed in approaching dilemmas and issues that are invariably context-specific. In most scenarios, for example, it is rarely possible to focus solely on the duties and intentions behind actions, without regard to their consequences, and vice versa. In real world decision-making, IT professionals invariably have to balance diverse, often conflicting demands and responsibilities. These may include legal and contractual obligations, business priorities and targets, and pressures to cut costs and deliver within timescales. Professional conduct may be shaped by a complex mix of personal values and character traits, including opportunism, self-interest, altruism, and common sense ethics.

The *codes of conduct* of various IT professional bodies are widely acknowledged as an important source of guidelines and standards of ethical practice. They are founded on similar principles and contain similar mission statements and codes of practice. Professional codes of conduct, however, have a number of limitations. They only apply where practitioners are members of such bodies, a major problem given the relatively low membership levels of professional bodies in the field of computing. Unlike other professions, the lack of any legal requirements to belong to a professional body in order to practice, or of any serious sanctions if codes are breached, means they have no “teeth”. Some, like McBride (2012) have argued that codes of conduct tend to result in a rule-based view of ethics, resulting in the need for “experts” to interpret such rules for practitioners, and creating a climate of formal compliance. Professional codes also tend to be aimed at *individual* practitioners, rather than organizations or teams comprised of diverse, specialised roles, where most development work typically occurs.

The *law* is a key element of computer ethics, and includes legislation that has a direct bearing on ICT, such as data protection, more general legislation (on equality and discrimination for example), contract law and various other regulatory frameworks. While legal compliance is seen as a fundamental touchstone for professionals, legalistic approaches to ethical issues also have their limitations. An obvious problem is the diversity of legal systems, internationally, and the jurisdiction of some laws, resulting in a lack of uniformity in enforcement and applicability in key areas.

While some topic areas, such as intellectual property, have an abundance of legislation, others, such as AI, have very little. This highlights the disparity between the rapid speed of technology innovation, and the relatively gradual pace of legislative debate and enactment. This means that the law frequently lags behind new and emerging technologies, resulting in areas of application where there are legal vacuums and regulatory frameworks are absent. The associated risks of such technologies can be obscured by the initial rush of adoption and immersion, making regulation difficult.

Laws, moreover, are not always politically or economically neutral. Some laws protect the rights of citizens, while others are weighted in favour of state interests and contain exemptions that enable them to be over-ridden. Other laws are weighted in favour of private, corporate interests. The legal framework of copyright and patent law, for example, can be viewed as protecting the global economic interests, and reinforcing the dominance of the global software and entertainment industries. The legal extension of copyright, and the strengthening of penalties against infringement, have been the result largely of persistent government lobbying by such industries and their representatives (Berry, 2008, p. 51). Loopholes in the law can also be exploited, such as environmental regulations on IT energy use and hardware disposal which can be circumvented by companies through various *greenwashing* strategies (Delmas and Burbano, 2011, p. 3).

More fundamentally, the relationship between the “law” and “ethics” is not always harmonious (Duquenoy et al, 2005, p. 9). Technology can be built that is legally compliant but is not necessarily

ethical (some forms of spyware, for example), while actions that are “illegal” can, in some circumstances, be ethically justified (leaking “classified” information in the public domain, for example). Ethical conduct, then, may mean going beyond the law, or in exceptional circumstances, breaking it.

In many computer ethics texts, ICTs are acknowledged as having a *social* impact and generating social issues (Rosenberg, 2004; Baase, 2009; Tavani, 2010). Social and cultural norms are seen as variables that can shape the design, development and use of ICTs. In some texts, however, the nature of this social impact remains relatively untheorised, and tends to be presented as a one-way causal chain of effects (Stamatellos, 2007; Kizza 2010). Research in the field of science and technology studies suggest a more complex, dialectical relationship between technology and society (Bijker et al 1987, MacKenzie and Wajcman 1999). As Johnson argues, ICTs are *shaped* by social forces, but they also have a configuring role, actively shaping human actions and practices (Johnson, 2009, p.12). A simple example of this in action are the “choice architectures” which shape users’ decisions through their structure and order (Thaler *et al*, 2007, p.1). Where systems are designed to privilege a particular option by default, people tend to choose that option more often, even though they may have the freedom to change or switch off certain technology features.

As ICTs become more ubiquitous and woven into all aspects of human societies, so too are their social implications becoming more widespread and profound. To fully understand these implications requires an inter-disciplinary approach, one that retains some of the key elements of classical ethical theories (that are rational, principled and value-based) but combines these with more modern perspectives on technology. The following sections propose a revised framework along these lines. The first task of such a framework is to explore the social and economic context in which ICTs are designed, deployed and used with a much richer understanding of the shaping forces involved. This paper argues that the methodology of established approaches needs to be re-ordered, and in the case of social issues, reversed. Arguably, the “social” should be examined first, methodologically, in order for the ethical and professional implications of ICTs to be fully grasped.

2. Revising the Framework

2.1 Putting technology in its place

An obvious, but useful starting point, is the notion that ICTs are never outside of culture or society, but embedded in social structures and practices. Technologies are always accompanied by knowledges and discourses. The way they are represented in language frames how they are thought about, and has an important bearing on their design and use. This can be seen, for example, in media and marketing discourses about consumer technologies, such as smart phones, which valorise the “new” in terms of speed, efficiency and convenience, and which cultivate the “need” to be connected and always contactable, and the desire to possess the newest devices.

Ethical analysis can benefit from this kind of deconstruction work, and the unpacking of discourses around technology. Fostering a critical awareness of the importance of meaning and how it is reproduced through particular representations has an important place in computer ethics pedagogy. This might involve critically examining viewpoints that reflect the agendas of particular stakeholders, or media moral panics, such as those around computer games and violence. It could involve questioning the unproblematic use of terms such as “piracy” in discussions about intellectual property (David, 2010, p. 98).

Part of this pedagogic work also involves showing that technologies do not just appear out of thin air. They are designed, developed and implemented for a purpose, to achieve particular goals. Technology development is shaped by a complex set of social, political and economic forces. A cursory look at the history of computing reveals certain recurring drivers. From code-breaking technologies and missile ballistics, through semiconductor electronics to surveillance and encryption,

technology development has consistently been driven by the exigencies of intelligence, security and military needs. These needs continue to shape the development pattern of emerging technologies, such as bio-computing and nanotechnology.

Technology projects have been funded and shaped by the political and economic agendas of governments, not only through the sponsoring of military technology, but also to conduct myriad bureaucratic functions, from controlling and managing populations, to law enforcement, border control and counter-terrorism. Industrial and economic drivers have also been crucially important. Commercial applications have evolved from spin-offs of military technologies. They have been developed to solve particular business problems and to exploit opportunities. Technology innovations have also emerged from university-based academic research, both privately and publically-funded, while others have been initiated by creative entrepreneurs, both inside and outside academia. ICTs therefore need to be considered as commodities in their own right, whose success is determined by their histories of adoption and their relationships with rival products and services. Some technologies can and do succeed on their own technical merits, but rarely without financial backing and funding, or distribution and marketing resources.

This means that what ICTs look like, how they work and what they are used for, are inextricably bound up with the social and historical context *in* which and *for* which they are designed and developed. Who, or what, is funding the research and development of ICTs are powerful influences on the direction of their development. Facebook's privacy settings, for example, are set to open, by default, in order to maximise the sharing of personal data. Facebook has stealthily made increasing quantities of personal data openly available whilst simultaneously rendering these settings difficult to change and complicated to navigate for most users. Its encouragement of personal data-sharing is a calculated business decision to monetise users' profile data, by selling access to that data to advertisers and marketers. The code that enables this data-sharing is written and maintained by programmers and engineers who take their directions from administrators and executives, who in turn respond to investors and major shareholders.

Kuhn's (1962) groundbreaking work showed how scientific research aims are underpinned and informed by prevailing social priorities, assumptions and knowledge paradigms. The notion of value-free research and design in the field of computer science was similarly dispatched by formative ethicists like Nissenbaum (1998, p. 38). It is now widely accepted that ICTs are infused with cultural values. Those values get baked into computer systems, and systems become *infused* with them (Johnson, 2009, p. 18). They enter into all stages of the software development process, from user requirements analysis, through interface design, to testing and evaluation. They percolate down into the most detailed levels of fine-grained code through objects that name, represent and model real-world phenomena, or through rules and assumptions embedded in algorithms (Fuller, 2008).

Brey (2010) argues that the standard methodology of computer ethics tends to miss these embedded values, by focussing on the more transparent moral controversies that are related to the *uses* of technology. By doing so, it fails to identify morally *opaque* features which are either hidden from view or perceived to be morally neutral. These technology features are often experienced by users as fixed and unchangeable. They are, however, a function of their design. Engineers, technical administrators and programmers make decisions about the architecture of systems. Technologies are designed with intended purposes, and weighted in favour of particular uses. They allow certain actions while preventing or constraining others. These decisions involve judgements about what is important, and what is not, and, as such, involve exercises of power which have an inherently ethical dimension. Technologies, for example, can be designed to enable sharing of intellectual property, or to prevent its unauthorized copying and distribution, to protect privacy or infringe upon it. Technologies can be designed to control what can be accessed, downloaded or expressed online, or it can be designed to circumvent such measures.

Probing and exploring these embedded values, and revealing their articulations in technology is a vital part of ethical analysis. Such *disclosive* work, as Brey (2010, p. 51) calls it, is a necessary prerequisite to the normative evaluation of issues and proposing of solutions. This work is important too, because it is precisely these technology features, and their applications and uses, which tend to be at the root of many ethical issues.

2.2 Sourcing ethical principles

This paper shares the prevailing definitions of ethics in the field of computing, as *normative* (concerned with how we *should* or *ought to* act in particular situations) and *applied* (concerned with the principles and values that inform conduct in particular contexts). Where it differs from existing approaches, however, is by shifting ethical analysis away from an appeal to traditional, formalist theories in the field, towards alternative sources of ethical principles. Habermas's conception of *post-conventional* moral reasoning provides one such alternative. In contrast to pre-conventional or conventional stages of moral judgement (which are either punishment/reward driven, egocentric, individualistic or socially conformist in perspective) post-conventional moral consciousness is characterised by recognition and respect of difference, transcendence of personal interest, and acknowledgement that ethical and legal points of view may conflict. Post-conventional moral reasoning, crucially, is also informed by ethical principles (Habermas, 2001, p. 160).

This conception is grounded in fundamental presuppositions about humans as inherently communicative, autonomous social beings with a universally shared capacity for reason, and a profound desire to understand others, and the world around them. Habermas, in this sense, is part of a long tradition within political and moral philosophy of attempting to formulate an ethical foundation for human societies, a tradition which can be traced from ancient and classical beliefs, through to the modern era—from Aristotle, through Marx, to formative computer ethicists such as Wiener. These strands of thought are founded on a shared self-evident truth about all human societies—that humans work, live and form bonds with one another in communities, and tend to live and act by the values of those communities. As such, certain core values tend to recur across those societies, values that include, amongst others, non-harm, honesty, integrity, fairness, empathy, generosity, reciprocity and tolerance.

While the interpretation and expression of these values may vary between cultures, these core values can be shown to thread their way through most moral systems. For example, there are few such systems that do not affirm the principle of helping others and doing good (beneficence) or contain some form of injunction against needless harm to others (non-maleficence). The rule of reciprocity, of treating others as one would want to be treated oneself, is echoed in the scriptures of most of the world's religious belief systems and in many other non-religious and non-Western philosophies (Wilson, 1991).

These core values, and their accompanying human rights, are also enshrined in the constitutional and legal frameworks of many liberal democracies. They are affirmed too in the founding charters and universal declarations of international organisations such as the United Nations (1948, 1966) and the European Union (2000). Such principles include, amongst others, the right to freedom, autonomy and self-determination; justice, equality and non-discrimination; privacy; life, liberty and security of person and property; and freedom of thought, conscience, religion and expression. More recent frameworks that are specific to the *digital* society include declarations of information ethics by UNESCO and by the European Commission's *Digital Agenda for Europe* (Kinderlerer *et al*, 2012).

To these frameworks can be added the codes of ethics expressed from within the computing field itself, in the form of professional codes of conduct and ongoing academic discussions about the underpinning principles of computer ethics. The translation and application of these generic principles within the field of computing has been a source of debate. A key focus of this debate has been whether computer technologies, as a result of their unique properties, have generated new ethical issues, requiring new values, or whether such issues are simply new twists on age-old dilemmas

which predate the advent of computer technology. Arguably, there is truth in both of these positions. Many ethical issues do indeed touch on classic ethical principles which have been transposed to the digital domain. Equally, other ethical issues have assumed a greater significance as a result of the specific nature and implications of ICTs. As Moor (2004, p. 28) argues, the unique properties and applications of these technologies have brought particular issues, and values, to the fore. This *specificity* needs to be taken into account when attempting to catalogue the major ethical issues around ICTs, and the underlying values that are stake. Such a list will inevitably involve a mix of both general ethical values along with those that are specific to the domain of ICT. What follows is a brief overview of some of these major issues.

Privacy is a prime example of an issue that pre-dates the digital era, but one that has been transformed by computer technologies to become *the* predominant ethical concern of the times. Privacy is a core value that crosses over into virtually every major topic area in computer ethics. It is a multi-faceted category, encapsulating elements of confidentiality, anonymity and solitude, but also a fluid concept, whose value is context-specific (Nissenbaum, 2004, p.155; Tavani, 2010, p.138). Initially precipitated by advancements in data management technologies, and their potential to collect, exchange and analyse unprecedented amounts of information, privacy concerns have been heightened by various developments in the scale and complexity of these processes. These include the harvesting of online digital footprints from disparate sources, the aggregation and cross-correlative mining of this data for commercial or state intelligence, and new paradigms of data collection and processing around the convergence of distributed, wireless and geo-spatial technologies. Computer networks are merging into a common platform in which smart objects interact and automatically exchange data sensed about their environment. The miniaturization of processing combined with machine-sensed data from various signals and sensors is extending computing into a wider range of environments, breaking down the boundaries between public and private domains. As computers lose their stationarity and become ubiquitous in these environments, there is the potential for new kinds of continuous, remote monitoring and ever more fine-grained modes of surveillance.

The collection and processing of data on the human body has the potential to further threaten privacy and to violate bodily and psychological integrity through new kinds of intrusive surveillance. These include technologies that monitor physiological characteristics and emotions or which extract data from the body. They include various biometric technologies, and those that interface with, or are embedded in, the human body. As ICT becomes more and more apart of us, biologically, the human body and the way in which technologies work on it, will likely be at the heart of many future ethical issues in computing.

While concerns persist about the threat to individual privacy posed by the state, there has been a greater focus on those posed by private corporations, due to the exponential increase in data exchanged between organisations. While entry to privately-owned networks is becoming increasingly essential in order to work, communicate and access services, the terms and conditions of this entry involve the giving up of personal data and consenting to varying degrees of monitoring.

The status of privacy rights and expectations in an online world has become a source of ongoing debate about whether we have entered a "post-privacy era" of disclosure and transparency, or whether such claims are premature. For most commentators, however, privacy remains a fundamental right whose distribution is asymmetrical, one where citizens are increasingly transparent to monitoring, while organisations, businesses and governments claim a shield of privacy (Andrejevic, 2007, p. 7).

At the heart of many contemporary privacy issues is the collection and flow of personal information, in particular, the ability to collect increasingly detailed information about who people are, what they do, and their past and present locations. Privacy, as a result, is closely interwoven with the issue of *data protection*. Underpinning this issue is the right to determine when, how and what information is disclosed to, and used by, others. Data protection is an area where a clear body of identifiable principles exists in legislative form in many countries. Such legislation specifies the rights of data

subjects and the legal obligations of organizations regarding the collection and processing of personal information. It is the surreptitious nature of these processes, however, that has been an area of particular concern. Where these processes occur without the subject's knowledge or agreement, the principle of *informed consent* has become a central issue. More recent areas of concern have been around data retention and the "right to be forgotten," as well as changing categories of personal data, such as IP addresses, RFID and geo-location data, and new kinds of *sensitive* data, such as genetic and biometric data. The use of personal data for behaviour profiling and risk assessment is another contentious area, whether used for the prediction of crime or social unrest, acts of terrorism, or consumption patterns for advertising purposes. Where such profiles are used to make decisions about people, the consequences of inaccurate information can lead to discriminatory treatment, such as being detained or considered a potential security threat, or being denied a loan or a job. This can result in people being deprived of specific rights and opportunities, but can also have the effect of a more general form of *social sorting* which reinforces existing social and economic divisions and power relations (Lyon, 2007, p. 63).

Intellectual property has been a perennial issue in computer ethics, one brought to the fore by the digitalisation of information and the whole panoply of technologies which have facilitated the copying, distribution, sharing and locating of intellectual products. Patent, copyright and trademark issues continue to abound in the digital realm. Corporations have become increasingly aggressive in their attempts to increase their holdings and revenue streams by "enclosing" the internet and deploying various digital rights management systems to prevent or deter these processes (Benkler, 2006, p. 26; Lessig, 2006, p. 116). The ethicality of these attempts, and of the dominant legal and regulatory framework around intellectual property continue to be highly contested areas. Proprietary software has been challenged by various alternative models of ownership, by "gift" economies and by commons frameworks which fall under the broad umbrella of free and open source software (Berry, 2008, p. 9). At the heart of these issues lies the question of how to balance the rights and interests of creators, owners and users, while preserving the free flow and fair use of information and knowledge.

Free expression is another classic, pre-digital issue which has been inflected by ICTs and which continues to raise fundamental questions about what content, if any, should be regulated or censored, how, by whom, and at what point, technically? Some have argued that the Internet's disparate communication platforms and Web 2.0 technologies have opened up a more participatory communication models of *commons-based peer production* (Benkler, 2006, p. 8) thereby transforming consumers into active *producers* (Bruns, 2008, p. 2). In so doing, they have facilitated a global explosion of computer-mediated mass self-communication, and a borderless, free flow of information, from subversive expressions of political dissent to the micro-banalities of social media (Castells, 2009, p. 135). Others like Lessig (2006, p. 4), however, have warned that the Internet's technical infrastructure can, and has been, configured to restrict what can be accessed or expressed, and to enable greater regulatory control and surveillance. This has created the possibility of sealed, bordered internets within specific nation states (Spinello, 2011, p.48). The precise technical means to accomplish this have varied, but they include the centralised control of primary internet traffic, various blocking and filtering mechanisms, layers of digital identification and authentication, and geolocation software (Zittrain and Palfrey, 2007). The use of these measures by states for political censorship, security and intelligence purposes has been a key concern. However, private actors have also become powerful information gatekeepers with an increasingly pivotal role in the distribution and construction of knowledge. Search engines, for example, can affect the diversity of content available, through their selective filtering and ranking of search results (Vaidhyanathan, 2011, p. 21). This touches on the related issue of *net neutrality*, the principle that networks should not discriminate between traffic, either in terms of content, origin or destination, nor attempt to mandate digital identity and traceability.

These principles also connect with issues around *information and communication ethics*, raising questions about information validity, and associated issues of trust, confidentiality, and the public right to know. Online communication and connectivity also has potential implications for identity and

self-expression in spheres such as entertainment and work. The relationship between online and offline identities, and between virtual and “real world” behaviour is a key issue in areas such as social networking and games (Nabeth, 2009, p. 8; Papacharissi, 2010, p. 304). The potential for deception and the disinhibiting aspect of digital environments where online social identities can be manipulated or concealed, also raises issues of authenticity and trust.

Security issues in computer ethics have traditionally been framed in terms of the threats and risks entailed by ICTs, especially those posed by various forms of *cybercrime*. While many such threats might appear ethically uncontentious, there remain issues to explore around the circumstances in which some actions (such as hacking) might be ethically justifiable. Equally pertinent, in terms of ethical and professional issues, is how organisations *respond* to these threats, raising questions about what should be secured, and why, and the obligations to design systems with appropriate security features. Where the social and economic infrastructures of societies are increasingly dependent on computer systems, the *reliability* of such systems continues to be an issue. Where computer systems have unintended or unpredictable, and harmful, consequences, whether as a result of their unique properties or design errors, assigning responsibility also remains a key issue.

The principle of *transparency* stems from another key property of computer technologies, their *invisibility*. Moor’s notion that computer operations, most of the time, to most people, are hidden from view, and beyond their comprehension, remains as pertinent as ever (Moor, 1985, p. 272). This is particular so, given the complexity and scale of many programs, and the opacity of their source code, functions and algorithms. These “invisible” properties are open to various forms of abuse by technical specialists and insiders, for example, the fraudulent financial transactions that have been enabled by the digitalisation of stock-market trading. Numerous other examples of the invisibility factor exist, in the covert surveillance of communications data, in cookie technologies and behavioral advertising, in the algorithms behind search engine technology, or the prescriptive functions of “closed” proprietary programmes whose source code is concealed due to pressure from copyright owners and patent holders.

Transparency issues have arguably been made more pressing by intelligent systems which are specifically *designed* to be invisible, where users may not notice that they are interacting with technology which is embedded in the environment or working autonomously in the background. Ambient intelligent applications, for example, which respond to the presence of people and take actions on their behalf, according to stored assumptions and rules, also raise issues of *autonomy* and *self-determination*. Where devices interact with one another and make autonomous decisions independently of human intervention, be they smart homes or autonomous military vehicles, issues are raised about the delegation of human control, judgement and responsibility.

Equality and *fairness* encompass a broad range of issues related to the design, distribution and access to computer technologies. Technology can perpetuate and amplify existing inequalities in a number of ways, around divisions of class, age, gender, race or disability, for example through data profiling, as discussed above, or by verifying identities and determining eligibilities. Equal access to ICTs and the fair sharing of their benefits continue to be major issues given the persistence of various disparities in the distribution of access to computer resources and connectivity. Those who lack these resources or connectivity are potentially disadvantaged by a corresponding lack of access to job opportunities, education, information and various other commercial and public services. Access depends not only on having the financial means to acquire the necessary technical resources but also the necessary *skills* and *knowledges* to use such resources. This can result in knowledge divides between those who have the expertise to use technology to their benefit, and those who don’t. Where access also depends on having particular physical or sensory abilities, this also raises questions about inclusiveness of design for groups whose abilities are impaired. With digital inclusion increasingly becoming a key aspect of *social* inclusion, internet access may constitute a new kind of “right to be connected”.

The *environmental impact* and *sustainability* of ICTs has moved to centre stage as a consequence partly of their growing integration into the social and economic infrastructures of many societies. This environmental impact can be traced through each stage of the ICT lifecycle. The manufacturing of hardware components uses finite raw materials whose extraction is often carried out under inhuman conditions and results in toxic pollution (British Computer Society, 2011, p.13). The growing demand for data communications, processing and storage capacity have made the energy consumption of ICTs a central issue, particularly where that energy is derived from depleting resources, and where its carbon emissions may be contributing to global warming (Tomlinson, 2012, p.2). Environmental threats are also posed at the disposal end of the computer lifecycle, where discarded electronic devices containing toxic substances end up in landfills, and where much of the recycling of materials is unregulated and results in toxic waste. These issues are exacerbated by the high turnover, or *churn*, of hardware at the consumer end of the lifecycle where devices quickly lose their value, and become obsolete. This is reinforced by a pervasive upgrade culture in the technology industry, and touches on wider issues around planned obsolescence and waste in consumer culture (Slade, 2006, p.5). However, while computer technology is part of the problem, through its negative environmental impact, it is also potentially part of the *solution*, through various "green IT" strategies (Murugesan, 2008, p.26). These include the manufacturing of devices with fewer toxic substances, and in a more modular fashion, making it easier to upgrade individual components. They also include smart energy consumption and power management through sensor networks in power grids and buildings; more efficient use of computer resources such as data centres and servers, through consolidation and virtualisation; avoiding data duplication and redundant storage; and the use of cloud computing models.

3. Putting Ethics into Practice: Some Conclusions

The ethical issues, and revised approaches, outlined above have some important methodological and practical implications for computer ethics pedagogy. The nature of these issues, and the principles that they invoke, will undoubtedly evolve, in tandem with the emergence of new technologies and new domains of application. Most ethical dilemmas in the computing field are context-specific. They are therefore subject to change, as existing norms and practices change, and new ones are generated in response to new situations and dilemmas. Ethical issues therefore need to be continuously reviewed, curricula need to be updated, and codes of practice revised, by teachers, researchers and professionals.

Rather than providing a prescriptive decision-making schema, this framework offers a set of key questions that are pursued through a series of discreet, but inter-related stages or *moments*. The *first stage* of any ethical analysis is to identify a particular ethical dilemma, invariably in the form of a specific case where ICTs raise potential ethical issues. The *second stage* involves an analysis of the specific technologies involved, and the *social* context of their design, deployment and use. This entails some description of key facts and stakeholders, and an analysis of relevant economic and political drivers. This disclosive moment also analyses how values are embedded in the distinctive features or properties of technologies, and how these might generate ethical issues. This kind of analysis is an exercise in critical thinking in its fullest sense (Brookfield, 2005, p. 229). It implies an ability to distinguish opinion from fact, to recognise biases and fallacies, and to substantiate claims with evidence. It may also involve an unpacking of taken-for-granted beliefs, common-sense assumptions and dominant ideologies to reveal how technologies are socially and historically constructed, and how they may be skewed in support of particular interests and agendas.

A critical understanding of this "big picture" context is a necessary condition for the *ethical* moment, the *third stage*, where the values and principles that are at stake in any specific case are explicitly identified. The *legal* moment (the *fourth stage*) involves the consideration of any applicable laws or regulatory frameworks. Legal issues might arise from questions of jurisdiction, political and economic interests, or the ethicality of the law. In cases where there is no legislation, there is an opportunity to speculate about what such legislation might look like, or whether a particular issue can, or should be resolved by legal or regulatory means. The *fifth stage* of the framework is the *practitioner* moment

where the implications of the three preceding stages are followed through for *professional* practice. This involves considering the applicable codes of conduct of relevant professional bodies, while noting the caveats and limitations discussed above. The scope of analysis can be extended here to include a broader range of practitioners, such as managers and users of ICT, as well as more specific areas of practice, from games developers to CCTV operators. Existing codes of practice in these different domains can be explored, along with relevant debates around ethics and professional responsibility. The policies and codes of *organizations* such as businesses or government departments can also be considered, in the form of mission statements, for example. The *final stage* is to assess and *evaluate* potential solutions and practical courses of action. Different options and alternatives should be weighed up, judgments made, and solutions proposed. These should be justified and rationalized by appealing explicitly to the aforementioned ethical principles, laws and codes of conduct.

This approach enables students to demonstrate that they have thought through an issue and arrived at a balanced conclusion by considering different arguments and perspectives. The criteria of rationality should remain at the heart of this evaluative process. In this sense, the framework is informed by similar criteria and conditions to those proposed by Habermas in his procedural approach to “discourse ethics” (Habermas, 2001b). Any solutions should be informed by relevant facts and justified with sound reasoning, rather than personal preferences, gut instincts or opinions. While there may be no single “right” answer to an ethical dilemma, some actions or decisions are more defensible, rational and ethically acceptable than others.

Take together, these different moments of analysis—social, ethical, legal and professional—combine to form a multi-dimensional analytical framework. This is, necessarily, an inter-disciplinary exercise. The value of inter-disciplinary approaches is that different theories and methods, each with their own explanatory strengths, can be brought to bear on multi-faceted phenomena and objects of study. Connections can be made, and insights generated, across disciplinary boundaries. The value of inter-disciplinarity can be illustrated by looking at any of major issues in computer ethics. A topical question such as social networking services and “privacy”, for example, cannot be adequately grasped within the confines of any one single discipline. *Computer science* can shed light on the technical configuration of social networking platforms and the computational tools used to collect, analyse and share personal data. *Economics* can tell us about the business models and commercial drivers that shape some of these technologies, including the monetization of profile data. *Interaction design* may help us understand interface design, and the choice architectures implicit in privacy settings and options. *Communication* and *cultural studies* might help through textual analyses of privacy policies, terms and conditions, or by deconstructing wider discourses about “privacy”. Finally, to understand the meanings of online privacy for different social networking users, and their relationship to digital identities, we might look to *psychology* and the *social sciences* more generally

This framework offers a potential methodology for exploring the implications of *any* computer technologies, existing or emerging. Pedagogically, it provides students with a means of *doing* ethics for themselves, allowing them the flexibility to pursue their own topics of interest, in any domain where ethical issues arise. This is particularly pertinent when teaching students from different cultural backgrounds, and at disparate international sites and campuses. The ability to identify ethical issues in topic areas of high interest, and to tease out the underlying values at stake in such issues, is a crucial part in getting students to take ownership of such values.

While the development of post-conventional modes of moral reasoning is key to increasing levels of ethical awareness amongst adult learners, it is a capacity that has to be consciously fostered and learned (Brookfield, 2005, p. 257). This an urgent issue in the light of research which has found relatively low levels of moral judgement skills amongst computing students in higher education, regardless of level, and despite exposure to computer ethics in the curriculum (Holland, 2011, 121). The challenge is how to engage students through relevant topics, and through values and principles

that resonate with their own experiences, while at the same time encouraging critical approaches that look beneath the surface and question taken-for-granted assumptions.

Habermas suggests there is a universal core of *moral intuition* in all societies, at all times [Habermas, 1992, p. 201]. Arguably, most rational actors in society, including students and IT professionals, possess some kind of *moral compass*. People constantly make moral decisions and judgments in their everyday lives. These judgements may be contradictory, or formed from a combination of emotion, reasoning, and post-hoc rationalisations of behaviour. Intuition alone, however, is an untenable approach to moral deliberation, one which offers no guarantee of moving beyond pre-conventional/conventional modes of reasoning. What it does offer is a starting point for pedagogic practice in computer ethics, and a set of foundational moral beliefs and raw materials with which to work in an educational context.

The approach outlined in this paper aims at generating a more dynamic, rounded understanding of ethical issues, and a greater degree of reflexivity on the part of prospective practitioners. As such, this framework is of potential relevance to professionals as well as students. Understanding the broader social contexts in which computer technologies are designed, implemented and used is a prerequisite for professional social responsibility. Ethical considerations can and should be embedded in projects from the outset, from the planning and design stages, right through the development life cycle, to implementation and evaluation. Key actors, users and stakeholders can, and should, be involved in the design of systems and devices from the earliest stages. Such approaches, arguably, will enable better quality judgments and more informed decisions regarding the ethical design, development and deployment of ICTs. It is important, in this sense, to offer paradigms of value-sensitive design that are human centred, environmentally sustainable and socially desirable. It is important also to offer models of good practice, and examples of technologies which positively *affirm* values such as equal access, inclusive design, fairness, transparency, autonomy and freedom. It is important, finally, to provide a vision of what “good” looks like, and how to “do the right thing”.

References

- Adams, A. and McCrindle, R. (2008), *Pandora's Box: social and professional issues of the information age*, John Wiley, Chichester.
- Andrejevic, M. (2007), *iSpy: surveillance and power in the interactive era*, University Press of Kansas, Lawrence, Kan.
- Baase, S. (2009), *A Gift of Fire: social, legal, and ethical issues for computing and the internet*, Pearson, Upper Saddle River, NJ.
- Bijker, W.E., Hughes, T.P. and Pinch, T.J. (Eds.) (1987), *The social construction of technological systems: new directions in the sociology and history of technology*, MIT Press, Cambridge, Mass
- Benkler, Y. (2006), *The Wealth of Networks*, Yale University Press, New Haven, Conn.
- Berry, D. (2008), *Copy, Rip, Burn: the politics of copyleft and open source*, Pluto Press, London.
- Brey, P. (2010), “Values in technology and disclosive computer ethics”, in Floridi, L. (Ed.) *Cambridge Handbook of Information and Computer Ethics*, Cambridge, Cambridge University Press, pp. 41-58.
- British Computer Society (2011) “Green Impact of Hardware”, *ITNOW*, Vol. 53, No. 3, pp.10-17.
- Brookfield, S. (2005) *The power of critical theory for adult learning and teaching*, Open University Press/McGraw-Hill, Maidenhead.
- Bruns, A. (2008), *Blogs, Wikipedia, Second Life, and Beyond: From Production to Produsage*, Peter Lang, New York.
- Castells, Manuel (2009), *Communication Power*, Oxford University Press, Oxford
- David, M. (2010), *Peer to Peer and the Music Industry: the criminalization of sharing*, Sage, Los Angeles.
- Duquenoy, P., Jones, S., Rahanu, H. and Diaper, D. (2005) *Social, Legal and Professional Issues of Computing*, Middlesex University Press, Enfield.

Simon Jones (2016) "Doing the right thing: computer ethics pedagogy revisited", *Journal of Information, Communication and Ethics in Society*, Vol. 14 Iss 1 pp. 33 - 48
<http://dx.doi.org/10.1108/JICES-07-2014-0033>

Duquenoy, P (2007) *Ethics in the Environment of the Information Society*, Conference on the 'Ethical Dimensions of the Information Society', Strasbourg, UNESCO/Council of Europe, available at: <http://portal.unesco.org/ci/en/files/25453/11909027561Duquenoy-Penny.pdf/Duquenoy-Penny.pdf> (accessed 24 February 2015)

Ess, C., and Thorseth, M. (2010), "Global Information and Computer Ethics", in Floridi, L. (Ed.) *Cambridge Handbook of Information and Computer Ethics*, Cambridge, Cambridge University Press, pp. 163-180.

European Union (2000) *Charter of Fundamental Rights of the European Union*, available at: http://ec.europa.eu/justice/fundamental-rights/charter/index_en.htm (accessed 25 July 2014).

Delmas, M.A. and Burbano, V.C. (2011), "The Drivers of Greenwashing," *California Management Review*, available at SSRN: <http://ssrn.com/abstract=1966721> (accessed 28 July 2014)

Forester, T. and Morrison, P. (1998), *Computer Ethics: cautionary tales and ethical dilemmas in computing*, MIT Press, Cambridge, Mass.

Fuller, M. (Ed.) (2008), *Software studies: a lexicon*, MIT Press, Cambridge, Mass.

Habermas, J. (1992) *Autonomy and Solidarity: Interviews with Jurgen Habermas*, 2nd edition. Edited by P. Dews. Verso, London/New York.

Habermas, J. (2001a) "Moral Consciousness and Communicative Action", in *Moral Consciousness and Communicative Action (Studies in Contemporary German Social Thought)*. Translated by C. Lenhardt and S.W.Nicholsen, MIT Press, Cambridge, Mass. (Originally published in 1983), pp. 116-194.

Habermas, J. (2001b) "Discourse Ethics: Notes on a program of philosophical justification", in *Moral Consciousness and Communicative Action (Studies in Contemporary German Social Thought)*. Translated by C. Lenhardt and S.W.Nicholsen, MIT Press, Cambridge, Mass. (Originally published in 1983), pp. 43-115.

Holland, L. (2011), "Do higher education computing degree courses develop the level of moral judgement required from a profession?", *Journal of Information, Communication and Ethics in Society*, Vol. 9 No. 2, pp.116-126.

Johnson, D. (2009), *Computer Ethics: analyzing information technology*, Pearson, Upper Saddle River, N.J.

Kallman, E.A. and Grillo, J.P. (1996), *Ethical Decision Making and Information Technology: an introduction with cases*, McGraw-Hill, New York.

Kinderlerer, J., Dabrock, P., Haker, H., and Nys, H. (2012) "Ethics of Information and Communication Technologies". *European Commission, Opinion No. 26*, European Union, Brussels, available at: http://ec.europa.eu/bepa/european-group-ethics/docs/publications/ict_final_22_february-adopted.pdf (accessed 29 July 2014).

Kizza, J.M. (2010), *Ethical and Social Issues in the Information Age*, Springer Verlag, New York.

Kuhn, T.S. (1962), *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago.

Lessig, L. (2006), *Code: version 2.0*, Basic Books, New York.

Lyon, D. (2007), *Surveillance Studies: an overview*, Polity Press, Cambridge.

MacKenzie, D. and Wajcman, J. (1999) "Introductory Essay" in MacKenzie, D. and Wajcman, J. (Eds.) *The Social Shaping of Technology*, Open University Press, Buckingham, pp. 3-27.

McBride, N. (2012) "Have IT ethics become if-then-else statements?", *British Computer Society*, available at: <http://www.bcs.org/content/conWebDoc/38393> (accessed 25 July 2014).

Moor, J. (1985), "What is Computer Ethics?", *Metaphilosophy*, Vol 16, pp. 266-275.

Moor, J. (2004), "Reason, Relativity, and Responsibility in Computer Ethics", in Bynum, T.W. and Rogerson S. (Eds.) *Computer Ethics and Professional Responsibility*, Blackwell, Oxford, pp. 21-37.

Murugesan, S. (2008), "Harnessing Green IT: principles and practices", *IT Professional*, Vol. 10 No. 1, pp. 24-33.

Nabeth, T. (2009) "Virtual online social environments, real digital identities issues", in Jaquet-Chiffelle, D., Benoist, E., and Anrig, B., (Eds) *D2.6: Identity in a Networked World; use cases and scenarios*, Future of Identity in the Information Society (FIDIS), European Commission, pp. 8-9. Available at: <http://www.fidis.net/resources/fidis-deliverables/identity-of-identity/#c1752> (accessed 28 November 2014)

Nissenbaum, H. (1998), "Values in the Design of Computer Systems," *Computers and Society*, March, pp. 38-39.

Simon Jones (2016) "Doing the right thing: computer ethics pedagogy revisited", *Journal of Information, Communication and Ethics in Society*, Vol. 14 Iss 1 pp. 33 - 48
<http://dx.doi.org/10.1108/JICES-07-2014-0033>

Nissenbaum, H. (2004), "Privacy as Contextual Integrity", *Washington Law Review*, Vol. 79, No. 1, pp.119-157

Papacharissi, Z. A. (Ed.) (2011) *A Networked Self: identity, community and culture on social network sites*, Routledge, New York.

Quinn, M. J. (2005), *Ethics for the Information Age*, Addison-Wesley, London.

Reynolds, George (2012), *Ethics in information technology*, Cengage Course Technology, Boston, Mass.

Rosenberg, R. (2004), *The Social Impact of Computers*, Elsevier Press, San Diego, Calif.

Slade, G. (2006), *Made to Break: technology and obsolescence in America*, Harvard University Press, Cambridge, Mass.

Spinello, R.A. (2011), *Cyberethics: Morality and Law in Cyberspace*, Jones and Bartlett, Sudbury.

Stamatellos, G. (2007), *Computer Ethics: a global perspective*, Jones and Bartlett, Sudbury.

Tavani, H. T. (2010), *Ethics and Technology: Ethical issues in an age of information and communication technology*, John Wiley, Hoboken, NJ.

Thaler, R. H., Sunstein, C. R. and Balz, J. P. (2010) "Choice Architecture", *Social Science Research Network*, Rochester, NY, available at: <http://ssrn.com/abstract=1583509> (accessed 25 July 2014).

Tomlinson, B. (2012), *Greening Through IT: information technology for environmental sustainability*, MIT Press, Cambridge, Mass.

United Nations General Assembly (1948), *Universal Declaration of Human Rights*, Resolution 217 A (III), available at: <http://www.un.org/en/documents/udhr/index.shtml#a12> (accessed 25 July 2014).

United Nations General Assembly (1966), *International Covenant on Civil and Political Rights*,

United Nations, Treaty Series, Vol. 999, available at: <http://www.refworld.org/docid/3ae6b3aa0.html> (accessed 25 July 2014).

Vaidhyanathan, S. (2011), *The Googlization of Everything: (and why we should worry)*, University of California Press, Berkeley.

Wilson, A.W. (Ed.) (1991), "World Scripture: a comparative anthology of sacred texts, *International Religious Foundation*, available at: <http://www.unification.net/ws> (accessed 25 July 2014).

Zittrain, J. and Palfrey, J. (2007), *Access Denied: The Practice and Policy of Global Internet Filtering*, Oxford Internet Institute, Research Report No. 14, available at:

<http://www.oii.ox.ac.uk/research/publications/RR14.pdf> (accessed 25 July 2014).