Ethics of Computing

Penny DUQUENOY

Middlesex University, London.

p.duquenoy@mdx.ac.uk

"The Introduction of words like "ethics" and "ought" into conversations about science seems almost always to engender a tension ..." (J. Weizenbaum) [13]

Key words: Computer ethics, professional responsibility, ethical design, software applications, artificial intelligence.

INTRODUCTION

Prior to the launch of new technologies we hear of the benefits these latest developments will bring. As the latest innovations become integrated with, and form the infrastructure of, everyday living, we begin to experience some less positive aspects. Clearly, computer technology brings both benefits and disadvantages – the degree to which either has an impact on individuals and their ability to live a “good life”¹ is the degree to which ethics is relevant to computing.

A number of ethical issues have been discussed and debated over the last twenty years or so under the broad category of “computer ethics”. This rather loose term has been criticised for placing ethics – a uniquely human characteristic – onto computers. Despite early predictions in Artificial Intelligence we are still far from being able to ascribe a moral viewpoint to mechanical devices. If, however, we talk about the ethics of computing we are talking about the use of computers – thus placing the moral perspective (and consequently the moral responsibility) firmly in the hands of computing professionals and the users.

It is issues of moral responsibility and moral choice that are at the heart of this paper.

Many computer professionals take the view that they are simply the providers of “tools”, and what people choose to do with these tools is not their responsibility. Whilst acknowledging that it would be naïve, and grossly unfair, to put all the moral wrong-doings of society at the door of computer developers, they cannot abdicate all responsibility. They, after all, are the ones who determine through their construction of artefacts the potential and limitation of actions (see [8] for example).

This paper begins by drawing on early discussions concerning ethical aspects of computer technology [9, 13, 14] and proceeds to use them to highlight different problem areas: technical, application, and environment. We then look at some of the major issues under discussion since these early writings, to show the scope of the problems, as well as their changing nature, as different technologies are introduced. This background sets the scene for the discussions of moral responsibility, and moral choices referred to above. Looking towards the future, we then discuss some of the technologies that are on the horizon, and notice some correspondence to concerns raised more than 20 years ago. Finally, we bring together the moral points raised in this work to provide a set of questions that could provide the basis of moral consideration when designing for the future.

BACKGROUND

As early as the mid-1950's Norbert Wiener [14] warned of the dangers implicit in machines that "acted" faster than we could react, and that had a complexity beyond our understanding. He foresaw the human loss of control in situations governed by computers - i.e. the loss of any timely intervention in an adverse situation, coupled with the inability to understand the cause of the problem. The focus here is on the technical characteristics (speed and complexity) of these machines, and their practical consequences. Clearly, if we do not understand what causes a problem then we can neither resolve it, nor predict other outcomes. If we are unable to predict future behaviour or outcomes we are in effect “out of control” of any situation determined by such complex devices. Similarly, if events are happening faster than we can respond to them, then those events are also beyond our control.

Some twenty years later, and prompted by the upsurge of interest in Artificial Intelligence, Joseph Weizenbaum's concern was with the envisioned practical application of computers [13]. That is, the view that computer technology, by virtue of its logical operation could mimic the rationality of human beings. Not only that, but also because they were not

prone to 'human error', they could be relied upon to perform tasks more efficiently than humans. His views on the moral limitations of computer applications are presented below.

These two pieces of work are important not only because they were the first early milestones in the area of the ethics of computing and tremendously influential to the ensuing debates. In the context of this article they are useful, in that they identify the two most fundamental aspects of computer technology in relation to ethics. In the first instance Weiner is concerned with the *technical* characteristics of computers, in the second Weizenbaum addresses the problem of their *application* in a social context.

The growing realisation that ethics was not only relevant but a vital consideration of computing gained impetus in the mid-1980's. Attention turned once more to technical characteristics, and their unique properties prompting a debate lasting 10 years on whether computers raise "special ethical issues" [9]. Central to this debate was the notion that the digital environment of computer technology does not easily map onto its analogue counterpart. In other words, the representation of material in binary form presents difficulties in a world that has traditionally operated with an analogue model. This latter observation sets a different context – the operational environment.

The above three articles identify the features of computers that, individually and in combination, provide the foundation for many of the difficulties we now face: their particular technical attributes, their application, and a digital domain that challenges previous mental models. We will return to these three categorisations later when discussing moral responsibility, and the moral assessment of computing.

THE ISSUES RAISED

To assess the impact of this technology, let us now look at the issues that have been raised by the different dimensions discussed above.

The topics covered and issues addressed between 1985 and 1995² are naturally indicative of the chronology of developments during that time. For instance, in 1985, liabilities in relation to defective programmes, and related issues of codes of conduct and professional ethics were major areas of concern. Also on the agenda were privacy, security, as well as power and democracy (as in [4] for example).

As personal computers became more widely available from the 1990's onwards, we see the increasing use of computers in the workplace. Consequently the topics under discussion reflect this move: quality of personal life, quality of work life, impact on employment and third world,

legal issues and computer crime. Around this time we also have the introduction of floppy disks and networks – all allowing file transfers – which bring hacking and viruses, and other computer crimes onto the scene. In addition, the topics of artificial intelligence and expert systems make their first appearance.

In latter years (since 1995) the phenomenal upsurge in computer use and inter-connection that has been enabled by the Internet and, in terms of public access, the World Wide Web has broadened the field further. The issues covered reflect the hazards of "interconnectivity": junk email; email monitoring and other aspects of surveillance; intellectual property (now including publishing issues relating to web pages, trademarks and logos); issues of anonymity and pseudo-nymity (including misrepresentation); easy access to illegal and harmful material (in particular pornographic material and its availability to young children); to name but a few.

Furthermore, we see issues discussed in human rights terminology such as freedom of speech, technology and democracy, and equality of access. We also see the appearance of items that are of global concern, specifically: Internet governance and regulation, free speech and content control, encryption, etc.

The issues of privacy and security, on the agenda from the early years, - remain a major concern but gain a change of emphasis. For example, discussions on these subjects in later books (since the Internet), along with discussing personal data, also emphasise monitoring, tracking (i.e. cookies) and surveillance. Thus, whilst there has always been a concern for issues falling under the banner of human rights, the details of how such rights are threatened change as the technology changes and allows a more diverse range of human action. The focus on the impact of the Internet is particularly apparent since 2000 [1, 3, 6, 8, 10, 11].

MORAL RESPONSIBILITY

Looking at the range of issues above it would be understandable to take the view that "these things happen", and that such problems are simply the trade-off we make when we embrace new technologies. However, whilst all of them are the result of developments and decisions made by computer scientists, as noted previously not all the blame should be placed at this door. However, there are some areas where responsibilities for outcomes most obviously fall to the experts. These are the directly technical aspects, such as defective programmes and failed systems, and the problems of hacking and viruses, for example. We have seen from the survey above that the profession has been addressing these issues with professional codes of

conduct, and developing technical security measures to combat unauthorised access.

It could also be argued that the privacy issues fall within the technical domain – if general-use programmes (applications) are produced that are lax in their protection of others (email for example) – then the shortfall should surely be the responsibility of the developers. However, other issues of application – such as how and where AI and expert systems are used for example – are the decisions of the users. Taking a wider perspective, how Internet applications are used - whether to exchange ideas or pornography - is also clearly the users' choice.

One of the dominating issues throughout the period is that of Intellectual Property. This is often used as a prime example of the difficulties of mapping digital to analogue [9] – in other words, a problem of understanding and adapting to the new environment as categorised above. It is difficult to ascribe responsibility in this case to either technical experts or users. The resolution of many of these problems has fallen to legislators (in the case of Intellectual Property this is a natural outcome, as it is a legal construct), but who also in their professional capacity have to understand and clearly define the difficult areas. Technical design can allow or disallow access to intellectual property (as in Digital Rights Management applications), but unless some balance is achieved regarding access and cost, users will seek ways to overcome the technical constraints.

MORAL CHOICES

Taking responsibility implies free choice. Society recognises that where individuals do not have a choice in their actions they should not be held responsible for them. Ethical action is also about choice – choosing good over bad, right over wrong, whatever we might determine such things to be. If we are asking technical experts, users and legislators to take moral responsibility for their actions or decisions, we have to assume choices are available. In the following paragraphs we return again to the three articles introduced above, and find that in each of the domains (technical, application and environment) each of the authors offer some interesting ideas for deliberation in this regard.

It is clear that, when designing and developing new technologies, choices are continually being made – how to improve performance, reduce costs, do something that has not been done before, etc. etc. These are all familiar and uncontroversial goals, but each one depends on previous work. In other words, new development does not happen in isolation. This is the gist of Weiner's warning - that even though scientists may have every good

intention as to the outcome of their work, that work is part of a larger picture. Each part contributes to continuing development and, eventually, to the whole. In other words, we should be always conscious of where our developments may lead us. Naturally it is beyond our capabilities to envisage all eventualities, but Weiner advocates a “continual scanning and re-evaluation” as the development proceeds, and he presses us to always “exert the full strength of our imagination to examine where the full use of our new modalities may lead us”.³

Moving on to the second domain, the application of computer technology, Weizenbaum is very clear about his choices. He names three areas where computer applications should not be pursued. In the first category are “ones whose very contemplation ought to give rise to feelings of disgust in every civilized person” and “all projects that propose to substitute a computer system for a human function that involves interpersonal respect, understanding and love in the same category”[13]. We must remember that these comments are set against a background of new research into Artificial Intelligence. The first quotation refers particularly to connecting animals to computers (specifically visual and brain systems) the second is a response to suggestions that a programme he created to demonstrate computer “conversation”⁴ could be used to replace psychotherapists. Concerning the latter suggestion he states “...there are some human functions for which computers *ought* not to be substituted. It has nothing to do with what computers can or cannot be made to do. Respect, understanding, and love are not technical problems.” [13].

Finally, Weizenbaum warns against anything which “can be seen to have irreversible and not entirely foreseeable side effects”, especially when there is “no pressing human need for such a thing”. He illustrates his point using the example of speech recognition, pointing out that although promoted as an efficient method for physicians to record notes and take actions more efficiently “such listening machines, could they be made, will make monitoring of voice communication very much easier than it now is.” With uncanny foresight, he continues: “Perhaps the only reason that there is very little government surveillance of telephone conversations in many countries of the world is that such surveillance takes so much manpower ... speech-recognizing machines could delete all “uninteresting” conversations and present transcripts of only the remaining ones to their masters”. [13].

The choices are not so explicitly laid out in James Moor’s paper. As a philosopher, his mission is to identify the “revolutionary” aspects of computers rather than pursue an opinion. However, these aspects – which according to Moor are their invisibility, logical malleability, and social impact – give grounds for discussion. The invisibility factor has a similar consequence to Weiner’s warnings about loss of control – when processes

are out of sight we are likely to either ignore them, or be unaware of them. In both cases these are usually classed as the benefits, if not the purpose, of computers – that is, to take the cognitive load off the user. Designers can choose to enhance or reveal “invisibility” – dialogue boxes for example reveal occurrences in programmes, often warning that something is wrong. Dialogue boxes characteristically offer the user choices, but users are often unaware of their choices (for instance, in rejecting “cookies”). The charge to designers and developers is to at least be aware of the inherent dangers of “invisibility”, and to incorporate choice for the user in the design where necessary – particularly where safety is an issue.

What about “logical malleability”? Moor’s explanation goes as follows: “Computers are logically malleable in that they can be shaped and moulded to do any activity that can be characterized in terms of inputs, outputs, and connecting logical operations ... The logic of computers can be massaged and shaped in endless ways through changes in hardware and software” and consequently “the limits of computers are largely the limits of our own creativity.”[9] The scope for choice here is clear – we can shape and mould computers to create an environment of our choosing. And so we arrive at his third revolutionary aspect – social impact. There is no doubt about the social impact of computer technology, we need look no further than the list of issues mentioned earlier to see the evidence of the range and scale of impact. However, if further evidence should be required we have only to remind ourselves of the almost global panic as we approached the year 2000, and the cost of the Y2K bug!

THE FUTURE

We have seen the way discussions in the field have been progressing and how the priorities for consideration have changed since the overwhelming rise of the Internet. It seems likely then that there will be more changes in the future. Having said that, some of the current issues - such as Internet governance, and security and privacy, are far from any resolution and will continue to confront us for a long time yet.

The security loopholes of the Internet are almost impossible to address [6], and whilst the emphasis is currently focussed on individuals as intruders into our computer systems (hackers), it is also possible that governments can use these loopholes as well. The debates on privacy are likely to increase as surveillance and monitoring become easier, and governments continue to feel threatened by secure encryption (which may be used against the interests of national security, and law and order).⁵

Regulating, or governing, the Internet is not an easy task - for any regulation to be effective a global approach will be needed. It is likely that the "Net" will fragment into sectors - one level reverting to the original government and scientific communications medium (this is already planned, and referred to as the "Grid"), another level supporting eCommerce, and yet another level providing the public communications space that we are now using. It is possible that we may see a move away from government (democratic) regulation, towards "self-regulation" - which is in reality decided by large corporations. We can see already that regulation is implemented via the technology itself - for example, access to certain web sites can be restricted; encryption protects intellectual property on video and audio content⁶ - and there is no reason to suppose this approach will lessen. It is more likely that these technical means will be developed and used in the interests of government and corporate policy. In other words, developers and designers will set tomorrow's scene.

The past has shown us that developments in computer technology are a result of choices made in many areas: development, infrastructure, government policies and take-up by the population. The technological drive is not pre-determined, and there is no "inevitable" future. The future, both beneficial and otherwise, will be formed from the technology that exists at the moment, and choices that will continue to be made in the areas mentioned above. Some choices have already been made, in the sense of research initiatives promoted by governments and other bodies encouraging research in particular areas - current key words are "ambient intelligence", "ubiquitous computing", and the "semantic web"⁷. It seems that future technologies are likely to be increasingly "intelligent" and everywhere! (Even in our clothes, and in our bodies.)

It would be foolhardy in these times of extraordinarily rapid change to offer predictions for the future. Past experience shows technologies put to very different uses than those originally envisaged by the developers (the Internet is a prime example, as is text messaging and mobile phones). It is possible however to consider the consequences of emerging technologies on certain basic human values such as free will (characterised by the ability to make choices), and respect for human life and human dignity (suggested by Weizenbaum).

Work in the field of Artificial Intelligence (AI) has been steadily progressing since at least the 1950's following Alan Turing's celebrated insight relating computation with intelligence. Questions raised even now can be projected towards any future work in this area. So-called "expert" systems and intelligent agents put decision-making in the control of computer technology. Evaluative judgements are inherently human attributes and, as we have noted, provide the basis for ethical action. What

ethical and philosophical questions are raised when they are put into the domain of intelligent agents? Will we simply be exchanging human fallibility for machine fallibility? If so, have we gained or lost? With judgement and decision-making comes responsibility and accountability (a basic premise of the judicial system) - how are we to accommodate these attributes into the domain of Artificial Intelligence? And finally, if any notion of ethics is dependant on free will, and the freedom to choose between actions, what is the ethical status of "intelligent agents" that exhibit free will and free choice? There are particular areas of concern where intelligent systems are used and making decisions and judgements on our behalf - for example in medical diagnosis⁸. Perhaps we should give consideration to Weizenbaum's question "is it there a pressing need for such a thing"? [13]

Computer simulations have proved immensely helpful in training - for instance aircraft pilots (and almost certainly in military defense). Simulation techniques are the backdrop for Virtual Reality (VR) - a technological representation of the physical world which includes human representation. Thought must be given to what is represented, and how it is represented. Representations in a virtual world could have an impact on personal identity (impersonation and misrepresentation), or on human dignity (violence or degrading behaviour) [12]. We will need to ask whether the interactive nature of virtual reality surpasses boundaries which might have previously been considered acceptable (e.g. in film making) when directed at a passive audience. Does it make a *moral* difference whether we watch, or we participate? Intuitively it does.

Experiments have been carried out in the area of computer implants⁹ and computer chips are now used in animals. Further research is likely to investigate the potential of implanted technology in humans for medical assessments and monitoring. What are the implications of being "always connected"? If tagging is seen as acceptable (used to track offenders instead of being in prison), why not implanted tags? After all, animals are tagged for the purposes of tracking and tracing, as well as for records of medical status. Aside from any health implications, these issues will ensure that the privacy debate will remain lively and controversial.

In all of these examples we are reminded of Weizenbaum's concerns: connecting animals to computers; simulating human functions involving respect, understanding and love; and surveillance. Finally, if we are to take information itself as a value - justified by a right to knowledge - then we must accept that everyone should be entitled to equal access. Whilst many governments are committed to promoting equal access to communication technologies - what about equal access to the information they provide? There are risks of preventing access to information through

software, for example "digital rights management software" - introduced as a response to fears regarding entitlement to Intellectual Property. This software is designed to specifically restrict access, and is supported legally by United States government through their introduction of the Digital Millennium Copyright Act 1998 which criminalizes any attempt to technically side-step Digital Rights Management technologies. In this case the technology has government backing – but in other situations corporations and organisations can, through software, regulate use. [8]

CONCLUSIONS

We have seen how computer technology not only changes interaction, but also facilitates different ways of interacting, and opens up potential activities hitherto unrealisable. When taking the ethics perspective it is easy to be dismayed at the possibilities of adverse intent. We should not forget that when all is said and done, computer technology does provide real benefits in a great many areas – a prime example being global communications. We do recognise that the benefits of these technologies have been immense. However, the arguments for the benefits of new technologies are well supported by the companies who develop and supply them, and the media. There should also be a balancing point of view – and this too is gaining momentum. There are already a number of initiatives in the field of business ethics, and computer ethics is gaining ground in this respect¹⁰. Raising public and organisational awareness is a first step.

Just because new technology introduces ethical challenges, it is important not to forget that how we respond to them, and how we can shape the future, is in our hands. (After all, someone chooses what programming code will be devised and written.) There has always been, and will always be, those who exploit situations for their own advantage and against others. Fundamental to any discussions on ethics is the principle that human beings have free will - that is, each person can make a free choice in regard to thinking and their actions. We should perhaps follow Wiener's advice and "exert the full strength of our imagination to examine where the full use of our new modalities may lead us".

Taking some of the key values drawn out in this paper (choice, dignity and respect, equality of access) can provide a basis for "examining new modalities". We could ask of future technologies:

- To what extent do they allow or prevent individual choice?
- To what extent do they raise or diminish human dignity?
- To what extent do they respect or impoverish person-hood?
- To what extent do they provide equal opportunity, and equal access?

And in all cases we should ask ourselves: What is the trade-off, and how far are we prepared to go?

REFERENCES

- [1] Baase, Sara (2003) *A Gift of Fire: Social, Legal, and Ethical Issues in Computing* (2nd Ed.). Prentice Hall.
- [2] Duquenoy, Penny. and Harold Thimbleby (1999) "Justice and Design" in *Human-Computer Interaction*, INTERACT'99, M. Angela Sasse and Chris Johnson (Eds.) IOS Press on behalf of the International Federation for Information Processing (IFIP).
- [3] Hamelink, Cees J. (2000) *The Ethics of Cyberspace*. Sage Publications.
- [4] Johnson, Deborah G. and Snapper, John W. (1985) *Ethical Issues in the use of Computers*, Wadsworth Publishing Company. (Out of print).
- [5] Johnson, Deborah G. (1985) *Computer Ethics* (2nd Ed.). Prentice-Hall Inc.
- [6] Kizza, Joseph Migga (2002) *Computer Network Security and Cyber Ethics*. McFarland.
- [7] Kling, Rob (Ed.) (1991) *Computerisation and Controversy: Value conflicts and Social Choices* (2nd Ed.), Academic Press.
- [8] Lessig, Lawrence (1997) *Code and other Laws of Cyberspace*. Basic Books.
- [9] Moor, James H., (1995) "What is Computer Ethics?" in *Metaphilosophy*, 16 (4)
- [10] Langford, Duncan (Ed.) (2000) *Internet Ethics*. McMillan Press.
- [11] Spinello, Richard (2000) *Cyberethics*. Jones and Bartlett.
- [12] Whitby, Blay (1996) *Reflections on Artificial Intelligence*. Intellect Books.
- [13] Weizenbaum, Joseph (1993) *Computer Power and Human Reason*. W.H. Freeman and Company, San Francisco. First published 1976.
- [14] Wiener, N. (1960) "Some Moral and Technical Consequences of Automation". *Science*, 131, pp.1355-8.

-
- ¹ Aristotle's Eudaimonia – where well-being and morality are intrinsically linked.
 - ² The initial analysis of topics covered in the computing ethics literature between 1985-1995 was carried out by Prof. Jacques Berleur, and the 1995 to current survey was a collaboration between Prof. Berleur and Penny Duquenoy.
 - ³ A similar approach was suggested in [2] whereby the idea of using John Rawls Theory of Justice, in particular designing from a "veil of ignorance" was used to encourage designers to imagine different possible perspectives.
 - ⁴ Called "Eliza" this programme responded to input questions with seemingly intelligent replies.
 - ⁵ There have been many discussions on this topic. For both sides of the argument see Dorothy Denning "Clipper Chip Will Reinforce Privacy" and Marc Rotenburg "Wiretapping Bill: Costly and Intrusive" in [6]. In the UK the

- Regulation of Investigatory Powers Act (RIP) (2000) provoked a storm of protest from civil liberties groups.
- ⁶ This technology is promoted as "Digital Rights Management" software (DRM).
- ⁷ For example, the Engineering and Physical Sciences Research Council (UK) proposes to further research into ubiquitous computing, and the "semantic" web; the European research funding agencies refer to "ambient intelligence". Tim Berners-Lee: " the semantic web will raise moral questions" (speaking in Oxford, UK, 2001).
- ⁸ For discussions on ethical aspects of Artificial Intelligence and Virtual Reality see [11]
- ⁹ For example, the research of Kevin Warwick, Professor of Cybernetics at University of Reading. (<http://www.kevinwarwick.com/>)
- ¹⁰ Ethics is gaining a higher public and organisational profile. To take just one example, the Royal Society of Arts in 1997 organised a Forum for Ethics in the Workplace. The discussions have instigated a number of projects addressing similar questions to the ones raised in this article. Who makes the decisions about innovation in industrial science? At what point in the R&D process are these decisions made? What are the criteria? Do ethical or social considerations play any part? *RSA Journal*, 1/6 2002 p.26.