Moor Food for Thought: Five Key Issues in Computer Ethics
Erica L. Neely

A. Introduction

Almost thirty years ago, James Moor (1985) wrote "What is Computer Ethics?", arguing for the special status of computer ethics as a developing field. Displaying an uncanny prescience, he argues that computers are transforming our social institutions; as they continue to do so, they will raise many ethical questions. The potential scope of these questions is almost unlimited, he believes, due to the essential malleability of computers: since they have the potential to be used in a multitude of ways, they have the potential to raise a multitude of ethical dilemmas.

In this paper I will first briefly touch on some of the specific issues Moor raises. I will then turn my attention to the five overarching areas I see as most important in computer ethics today: privacy, identity, trust, responsibility, and access. Within each of these broad areas I will consider a number of related philosophical questions which have either arisen or become more pressing with advances in technology. My goal is not to answer all of these questions but rather to direct attention to issues which urgently require further philosophical attention.

B. Brief Overview of Moor

In his article, "What Is Computer Ethics?", Moor (1985) observes that the Computer Revolution has two stages: introduction and permeation. During the introduction phase computers were created; during the permeation stage the technology is integrated into our societal institutions. This transforms those institutions in such a way as to raise philosophical questions about the nature of the institution. For instance, the introduction of purely electronic forms of currency, such as Bitcoins, raises questions about the nature of money.

Moor mentions a number of examples in his paper, displaying an uncanny prescience for predicting many of the live issues we face today. For instance, he notes that the use of computers in tallying and/or recording votes raises questions about what a fair election is, since in the United States the results from eastern states can be displayed before the polls have closed in western states. Similarly, we have seen questions concerning voting machines and how to ensure they are operating fairly and without skewing results.[1]

Another interesting question Moor raises is how the nature of work will be transformed when computers or robots are able to perform many of our current tasks. While we have not yet reached the point where most jobs are simply a matter of supervising machines while they work, there are industries where this is becoming more standard. Commercial pilots, for instance, rely heavily on instrumentation to assist them in their flights; for much of the flight they are simply present to ensure nothing goes wrong. In many businesses the ability to telecommute is causing us to ask questions about the nature of a work environment, such as whether we lose anything vital when employees are not physically in the same location. Similar questions are being raised about the nature of education and whether it requires face-to-face interaction. The introduction

[1] As Kusters, Truderung, and Vogt (2012) discuss, there are ways of circumventing the verifiability checks we have tried to place on electronic voting machines; a perfect solution does not yet exist.

of Massively Open Online Courses (MOOCs) is causing the academy to reflect upon what is truly required for education and how to adapt new technologies to meet these requirements.

One of the reasons Moor sees computers as revolutionary is because of their malleability; they have the ability to be used in a wide variety of contexts and for a wide variety of tasks. That trait raises its own ethical issues. For instance, consider our new ability to create weaponry via 3D printers. This raises questions about how to control something which can be created at home from parts which are not themselves dangerous; typical controls on manufacturing depend on a sort of centralization which is no longer always required. While we have faced this issue in the past with the manufacture of homemade explosives or drugs, our new technologies greatly magnify this problem.

Lastly, Moor's concern over what he terms "invisible abuse" is timely as well. Computers can operate invisibly, i.e., in such a way that people are generally unaware of their presence. We are increasingly aware of ways in which computer technology has been used to invade people's privacy. Hackers can monitor your key strokes to steal your passwords. Governments can deploy surveillance drones to watch their own citizens. There is a huge potential for invisible abuse, and we need to consider how to address these sorts of illicit behaviors.

Having briefly touched upon specific examples in Moor's paper which are live today, let us turn to a broader consideration of computer ethics. In the following sections I will discuss five issues I see as key to the present and future of computer ethics. Four of these issues – privacy, identity, trust, and responsibility – fit fairly tightly together; our ability to deceive people concerning our identity, for instance, makes it difficult to hold a person responsible for his actions. The last issue, access, is a meta-issue concerning the societal ramifications of becoming increasingly technologically dependent. Together these issues encompass the critical areas for philosophical thought in this field.

C. Privacy and Anonymity

The first key issue in computer ethics is one we have become increasingly of aware lately: privacy. Our advances in technology have made it easier to invade people's privacy than ever before. Some of these invasions may seem relatively benign. At this point, most job candidates are aware that their prospective employers will Google them to see what appears; people see themselves as having a responsibility to police their online appearance. Yet, of course, this is restrictive – why should an employer have the right to seek out, say, photos taken by others during a party? If the event is not job-related, then it is of questionable relevance to any employment decision; why should we think we need to suppress these representations of our lives?

Surveillance in general occurs at many levels. Google, for instance, tracks which web sites you visit and tailors its advertising to your profile. In some sense, of course, a web page is a public place, since anyone can access it, hence we should not expect the same degree of privacy as in a private place. Yet we often think that we are accessing them anonymously – no one knows that we are looking up particular medical symptoms from our home computer because we are

concerned about our health. However, in general, this belief in anonymity is illusory; we are being tracked for corporate purposes.[2]

On a larger scale, both the United States' National Security Agency (NSA) and Great Britain's Government Communication Headquarters (GCHQ) have admitted to tapping into the emails and cell phone calls sent by ordinary citizens; while they claim these are for larger security purposes, their actions constitute a huge invasion of privacy.[3] The public place argument does not work here, since these are communications between two private individuals. Furthermore, such large-scale invasions would not be feasible without the aid of computers to record and help filter the data collected; this unethical behavior is enabled by our technological advances.

Lastly, the increasing use of unmanned drones for surveillance, particularly outside of a recognized conflict, is also problematic. Since activities in public places can be viewed by passersby, they are clearly not private. However, in general we have a sense of anonymity while performing those actions. Yes, someone might witness us entering a grocery store and someone else might witness us leaving, but it is unlikely to be the same person. Moreover, unless we are being followed, no single person is witnessing us at all times. Our actions, most of the time, are not being given special attention – they are part of the background scenery for other people, observed and quickly forgotten.

However, surveillance drones can observe the background scenery and, if recording, later bring up our actions for special attention. They can be that single observer, correlating all of our movements to form a cohesive picture of our actions. Furthermore, we are not necessarily able to ensure that we are not being watched. In ordinary life, we tend to think that a quick look to see if anyone is around and watching is sufficient before performing an action we do not wish to have observed, such as adjusting an errant article of clothing. With the advent of small surveillance drones, we may frequently be observed even if we are not aware of it; this may make us less willing to take even small, harmless actions which we do not wish observed.

One of the key issues in computer ethics, therefore, is how to conceive of privacy and anonymity in the modern world. The fact that we can observe people with more ease and efficiency than before does not imply that doing so is ethical; we must consider what circumstances justify an invasion of privacy. Furthermore, we can ask whether there is a difference between corporate and governmental tracking – are Google and the government equally justified in tracking which websites you visit? And, lastly, to what extent have we tacitly consented to those invasions by entering public spaces, whether real or virtual? These are all key questions which must be addressed.

D. Identity

Developments in computers and other similar technologies also raise pressing questions about identity, particularly with the advent of the internet and virtual communities. One question

[2] There are ways of blocking Google Analytics and AdSense from your browser, but people who are unaware they are being tracked are unlikely to investigate ways of blocking them.

[3] Gellman and Poitras (2013) and MacAskill et al. (2013) discuss the actions of the NSA and GCHQ's surveillance programs.

concerns how to understand identity of digital artifacts. We can create a file on one computer and copy it to another. Given the relative ease of altering date and time stamps, we can essentially create an identical copy of the original file. As such, the question of originality becomes difficult: how do we distinguish which one is the original? Is there a reason to distinguish them?

Historically we have prized originality: we value an original painting or an original manuscript more than subsequent copies, even if those copies are excellent facsimiles. However, we have techniques which allow us to distinguish those copies, enabling us to preserve the distinction between original and copy. If there is no way to distinguish between the two, is there any reason to have a preference for the original? Does it even make sense to speak of the original of a digital manuscript or artwork?[4] As more creative activities move into a virtual setting, these will become pressing issues.

Another aspect of identity pertains to who an individual is online. Many people have profiles on a variety of web sites, each of which represents some aspect of themselves. Hence a person may use Facebook to connect with his friends, LinkedIn to present himself to potential business contacts, and Match.com to try to find a potential romantic interest. These sites each present a picture of the individual, yet it is likely that they display slightly different facets of him; he probably does not post photos from his family reunion to LinkedIn, nor does he upload his résumé to Facebook. Similarly, the description of himself which he uses to attract a potential date is likely to be different than the description which will lure a potential employer. Yet all of these connect to him in some way.

Is he the same person in all of these venues? Does it depend on how differently he presents himself? After all, much has been made of the ability to present oneself as someone radically different online.[5] Suppose he loves baking but is afraid to admit that; he could start a blog where he claims to be female. Is the person represented still him? Similarly, he could play an online game and create a persona for his avatar which he uses in all his interactions. Is this avatar him? Is it an aspect of him, like a fragment of his personality? Answers have varied, ranging from arguments that we have multiple identities, such as Sherry Turkle's view (2004) of our identity as distributed, to John Suler's view (2002) that we simply highlight different aspects of our personality in different settings. I tend to agree with Suler, but the questions are difficult.

Identity questions interact with other philosophical concerns. For instance, we may again raise privacy issues by asking whether all of these presentations should be linked to a single individual – should we be able to track one's online presences to a single person? One reason we frequently care about identity is that it enables us to trust people more easily; if you might not be who you claim, we are more wary. Similarly, we want to ensure that you are held responsible for your online actions; this raises questions of responsibility. I shall consider each of these issues in turn.

E. Trust

[4] Clearly we can distinguish between an original file and later revisions if the content has changed; the question is whether the distinction makes sense for two identical files.

[5] Note that this ability may be somewhat overstated, as I have argued elsewhere. (Neely 2013)

There are two facets to the issue of trust that are currently pressing. The first concerns trust in online settings. As mentioned in the last section, deception frequently seems easier in a virtual setting; how then are we to trust people we only interact with online? Questions of trust are not unique to online interactions, of course; we can be deceived in face-to-face settings as well. However, some of the identity categories we can typically discern with relative ease in face-to-face interactions – for instance, gender, age, perhaps nationality, and so forth – are not as easy to judge in online interactions. While there are positive aspects to this inability to categorize people, the possibility of deception raises barriers to trusting others online.[6]

If the stakes are relatively low in a particular online interaction, this possibility may not be of significant concern. For instance, if someone on a mailing list is falsely presenting themselves as a woman, often it will not matter.[7] However, the question of trust becomes much more worrying in the context of how to trust information presented online. This relates directly to identity issues since trust in information often stems from trust in the authors of that information; if we believe that you are an expert in a field, we are much more likely to trust what you say than if we have no such belief. The rise of the internet has made it significantly easier to connect with people from all over the globe. It has also simplified the mechanism for sharing information and conjecture with a wide audience. However, these technological advances have concurrently broadened the possibility for testimonial deception; as such, we face challenges about how to choose sources online, particularly if the authorship of those sources is unclear. As more of our information moves to a virtual setting, this will become of key concern.

The second aspect concerning trust deals with technology on a more general level. In our everyday lives, we rely on machines and computer programs to accomplish a variety of tasks; these range from mundane activities such as paying a credit card bill online or making a cellular phone call to specialized tasks such as performing robotic surgeries or tallying votes in an election. Most of us are not competent to assess how the machine or computer program works, yet we rely on them in an increasing variety of venues. We are simply trusting that they will work correctly without unanticipated negative consequences.

The stakes of this trust are high, since frequently technology has the potential to be of great aid in performing many tasks. We can miniaturize robots and surgical instruments so that they fit in much smaller incisions, for instance, leading to a shorter recovery time in patients. Similarly, we can avoid human error in tallying a large quantity of data; assuming the program is written correctly, the machine will produce the correct answer and not succumb to tiredness or boredom. Yet the potential for mistake or deception exists as well; machines can be programmed incorrectly, whether in error or due to malice. If I am not competent to assess the reliability of the technology I am using, then it difficult for me to detect when my trust is misplaced.

[6] One positive aspect of this blindness to common identity categories is that it may allow people to interact without being stereotyped. Stendal et al. (2011) discuss this with regard to disability, but their points may be extended to other categories as well.

[7] Assuming that their being a woman is not somehow relevant to interacting with them; a man posing as a female survivor of breast cancer, for instance, might be problematic.

In general, we seem to justify trust by relying on the ability of experts to provide checks on each other. Thus the idea is that if a particular company manufactures fraudulent voting machines or ineffective robots, eventually someone will notice the results and investigate the causes. This will result in a loss of trust on the part of the general public. However, this leaves us open to cases involving large-scale collusion, and it only truly addresses trust in some kind of regulated setting.[8] Online, there is much more freedom to deceive and much less regulation; it is not clear that this system of checks and balances will work effectively. This ties into our next ethical issue, namely, responsibility.

F. Responsibility

Questions of identity and trust lead to questions concerning responsibility: we want to know who we are interacting with so that we can hold them accountable for their actions. Similarly, we are more likely to trust a person or technology if we believe some accountability exists. Technological developments cause issues of responsibility to arise in a number of settings. First, there are questions about responsibility for how technology is used. Historically, of course, this arose with the atomic bomb; in a generalized form the question is whether a person who works on developing a technology is then responsible for harms caused by it. A recent example of this has arisen with the invention of 3D printing: what responsibility does a creator have for developing an invention which permits us to print our own guns?

Similarly, we can ask who should be held responsible for autonomous machines.[9] If an unmanned combat drone kills a civilian in Afghanistan, who is responsible? It presumably cannot be the machine, since it seems to lack necessary criteria to be a moral agent.[10] Is the designer somehow at fault because her programming failed to distinguish between civilian and enemy targets? Is the person who deployed the drone responsible, since she chose to use the machine in a circumstance where such a mistake could be made? If the latter, is it the person who gave the order or the person who physically took the steps to launch the drone? These questions are increasingly important due to the amount of automation in our lives. While frequently the stakes will be lower than killing an innocent person, machine error is clearly possible; we must consider who is responsible for the harms caused by those machines.

Moving to a slightly different context, our technologies have also made it possible to cause harm in virtual settings. The classic example is the Lambda MOO case, in which a man used his online avatar to sexually assault the avatars of other players.[11] While this did not cause equivalent harm to sexually assaulting someone in the physical world, it did cause distress and harm to the players involved; therefore the player causing the harm needed to be held accountable. Determining an acceptable punishment for virtual harms is difficult, however, and something we need to address as a society.

[8] This could be regulation by law, by economic pressures, or anything similar.
[9] Where "autonomous" is used here to mean something like acting without being directly guided by humans.
[10] At a minimum, it seems that consciousness is required to be a moral agent, and these drones currently lack that. See (Himma 2009) for further discussion of this point.
[11] This case has been discussed repeatedly; the original commentary appears in Dibbell (1993).

There have been proposals for dealing with harm in virtual communities and suggestions (Alemi 2008, Johansson 2009) that virtual punishments may work best; being ostracized by a group you are connected with may work better than handing out a fine, say. This will not work for all cases, however. Sometimes the line between online and physical communities is blurred; for instance, cases of cyberbullying are extensions of offline bullying. A purely virtual punishment may not suffice in these cases. Furthermore, even some virtual harms deserve real world consequences, but our current institutions are ill-equipped to handle these crimes. For instance, when hackers in different countries band together to hack a particular website, how do we punish them?

The technology which has given us new opportunities has also given us new abilities to cause harm. Much has been remarked (Suler 2004) about the fact that people will do and say things online that they would not say in other contexts; this appears to be, in part, because of the perception that they escape consequences for their actions. As such, an effort to address questions of responsibility are key both for preventing crimes and for establishing a degree of civility which could enhance many users' experiences online.

G. Access

Questions of privacy, identity, trust, and responsibility form a strongly interconnected web of ethical problems. However, there is a meta-issue which also deserves attention, namely, access to technology. Moor discussed the two stages of the computer revolution, noting that we were now at the stage of permeation. While I agree he is correct, we are clearly far from ensuring that everyone has access to technology. Yet we have reached a sufficient level of saturation that it is easy for people in relatively privileged situations to take its availability for granted.

This is a mistake; there are many people who cannot access various technologies for a variety of reasons. Some of these reasons are due to internal characteristics of the user. For instance, a blind or deaf user will have difficulty attempting to access a technology if it depends on sight or speech. These barriers are frequently a matter of design, specifically, they are a result of the invisible values which Moor mentioned. Programmers and designers are unlikely to be deliberately constructing barriers to access, but they frequently assume a user with roughly similar capabilities as themselves; this makes it difficult for the technology to permeate user groups with different characteristics.

Other barriers to access can come from features which are not intrinsic to the user. For instance, consider users from a low socio-economic class. Technologies usually have some cost – even inexpensive computers cost money – and not everyone is able to afford that cost. This raises questions about how to ensure equal access across a society; if we computerize vital services, we must ensure that all members of our society can access those services.[12] If we do not ensure equal access, then we create a hierarchy favoring the more affluent members of society; at the same time, the poorer classes are increasingly unable to function in this technologically-dependent society. Given the pressures to take advantage of our technological advances,

[12] Estonia, for instance, is having issues with this, since they have moved to a largely computerized system of government; this is problematic for those citizens who are more concerned with having enough food than whether they have a computer. (Nelson 2013)

determining how to avoid this kind of social injustice will become more urgent in the coming decades.

One final related concern is the issue of continued access. Humans have been preserving information in written formats for thousands of years. Many of these records are still accessible today, ranging from the Rosetta stone to illuminated manuscripts to a favorite childhood novel. However, it can be difficult to access computer data that is only ten or twenty years old. Our ability to collect and record information may be faster and more efficient than ever before, but we must deal with issues of obsolescence – how will we ensure that in the future we can continue to access this information? While it may not be terribly important if I cannot access emailed jokes from my university days, this issue increases in importance the more we move critical information online. We have an obligation to ensure access both to people today and to future generations.

H. Conclusion

Moor displayed an impressive degree of prescience in his paper on computer ethics. Many of the issues he used as illustrations, such as the nature of work and the potential for invisible abuse, present challenges today. In addition to these specific issues, I have argued that there are five broader philosophical issues deserving our concern. Specifically, I believe that areas concerning privacy, identity, trust, responsibility, and access will become increasingly important as permeation increases.

While Moor argued for the special status of computer ethics thirty years ago, I believe that the situation is even more urgent now. Our abilities rapidly outpace our social ideas of how to handle the ethical situations that arise; hence our confusion over the limits of privacy online or how to handle virtual crimes. The malleability of technology demands a malleability in approach and an ability to quickly adapt to a changing field. He who hesitates may not be lost, but he will certainly be overwhelmed by the whirlpool of ethical concerns.

References

Alemi, F. (2008) An Avatar's Day in Court: A Proposal for Obtaining Relief and Resolving Disputes in Virtual World Games. *UCLA Journal of Law & Technology*, 12: 1-54.

Dibbell, J. A. (1993) A Rape in Cyberspace: How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database into a Society. Retrieved from http://www.juliandibbell.com/texts/bungle_vv.html.

Gellman, B., & Poitras, L. (2013, June 7) U.S., British intelligence mining data from nine U.S. Internet companies in broad secret program. *The Washington Post*. Retrieved from http://www.washingtonpost.com

Himma, K. E. (2009) Artificial agency, consciousness, and the criteria for moral agency: what properties must an artificial agent have to be a moral agent? *Ethics and Information Technology*, 11: 19-29.

Johansson, M. (2009) Why Unreal Punishments in Response to Unreal Crimes Might Actually Be a Really Good Thing. *Ethics and Information Technology*, 11: 71-79.

Kusters, R., Truderung, T., & Vogt, A. (2012) Clash Attacks on the Verifiability of E-Voting Systems. *Proceedings of the 2012 IEEE Symposium on Security and Privacy*. doi: 10.1109/SP.2012.32

MacAskill, E., et al. (2013, June 21) GCHQ taps fibre-optic cables for secret access to world's communications. *The Guardian*. Retrieved from http://www.guardiannews.com

Moor, J. H. (1985) What Is Computer Ethics? *Metaphilosophy*, 16: 266-275.

Neely, E. (2013, July) Intertwining Identities: Why There is No Escaping Physical Identity in the Virtual World. Paper presented at the annual meeting of the International Association for Computing and Philosophy, College Park, MD.

Nelson, S.S. (2013, June 12) Tallinn: The Former Soviet City That Gave Birth To Skype. *NPR*. Retrieved from http://www.npr.org

Stendal, K., et al. (2011) Virtual worlds: A new opportunity for people with lifelong disability? *Journal of Intellectual & Developmental Disability*, 36: 80-83.

Suler, J. (2002) Identity Management in Cyberspace. *Journal of Applied Psychoanalytic Studies*, 4: 455-459.

Suler, J. (2004) The Online Disinhibition Effect. *CyberPsychology & Behavior*, 7: 321-326.

Turkle, S. (2004) Our Split Screens. In A. Feenberg & D. Barney (Eds.), *Community in the Digital Age: Philosophy and Practice*, (pp. 101-117). Oxford: Rowman & Littlefield.