Journal of Information, Communication and Ethics in Society

Does computing need to go beyond good and evil impacts?

Randy Connolly Alan Fedoruk

Does computing need to go beyond good and evil impacts?

Randy Connolly and Alan Fedoruk
Department of Computer Science and Information Systems, Mount Royal University, Calgary, Canada

Received 10 October 2014
Revised 18 January 2015
Accepted 3 March 2015

Abstract

Purpose – This paper aims to demonstrate that computing social issues courses are often being taught by articulating the social impacts of different computer technologies and then applying moral theories to those impacts. It then argues that that approach has a number of serious drawbacks.

Design/methodology/approach – A bibliometric analysis of ETHICOMP papers is carried out. Papers from early in the history of ETHICOMP are compared to recent years, so as to determine if papers are more or less focused on social scientific examinations of issues or on ethical evaluations of impacts of technology. The literature is examined to argue the drawbacks of the impact approach.

Findings – Over time, ETHICOMP papers have moved away from social scientific examinations of computing to more philosophic and ethical evaluations of perceived impacts of computing. The impact approach has a number of drawbacks. First, it is based on a technological deterministic style of social explanation that has been in disrepute in the academic social sciences for decades. Second, it uses an algorithmic approach to ethics that simplifies the social complexity and uncertainty that is the reality of socio-technological change.

Research limitations/implications – The methodology used in this paper is limited in several ways. The bibliometric analysis only examined five years of ETHICOMP papers, while the literature review focused on published computing education research. It is possible that neither of these forms of evidence reflects actual common teaching practice.

Practical implications – It is hoped that the arguments in this paper will convince teaching practitioners to modify the way they are teaching computing social issues courses: that is, the authors hope to convince educators to add more focus on the social context of computing.

Originality/value – The use of bibliometric analysis in this area is unique. The paper's argument is perhaps unusual as well.

Keywords Computer ethics, Socio-technical systems, Education policy

Paper type Research paper

1. Introduction

One of the many breakthroughs in the teaching of computing over the past two decades has been the relatively widespread recognition of the importance of social and ethical issues (SEI) in the education of computing professionals. A 1991 computing curricula report argued that students "need to understand the basic cultural, social, legal, and ethical issues inherent in the discipline of computing" (Tucker, 1991). The Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE), the two main professional organizations for the different computing disciplines, have promulgated numerous curricula recommendations that have maintained this claim. The current recommendation for computer science assigns 11 core hours to Social Issues and Professional Practice: four of those are to ethics, one to the social context of computing and the remaining hours cover some key issue areas

Journal of Information, Communication and Ethics in Society
Vol. 13 No. 3/4, 2015
pp. 190-204

1477-996X
DOI 10.1108/JICES-10-2014-0045

such as privacy, intellectual property and security (ACM/IEEE-CS Joint Task Force on Computing Curricula, 2013). Other computing disciplines such as Information Technology (Lunt *et al.*, 2008), Computer Engineering (ACM/IEEE-CS Joint Task Force on Computing Curricula, 2004a), Software Engineering (ACM/IEEE-CS Joint Task Force on Computing Curricula, 2004b) and Information Systems (Topi *et al.*, 2010) also have social issues and ethics proscribed in their ACM curricula reports. Accreditation bodies like the Accreditation Board for Engineering and Technology, the British Computing Society, the Australian Computer Society and the Canadian Information Processing Society all recommend the inclusion of an ethics course in any accredited computer science program.

Happily, it appears that these curricula recommendations have been widely accepted. A recent survey of American universities found 88 per cent of computer science programs included ethics in their computer science curricula (Spradling *et al.*, 2008). Another survey that also included non-American universities found that a full 95 per cent of their respondents included ethics content in their computer science programs (Goldweber *et al.*, 2013). When one examines the past decade of papers at the principal ACM computing education conferences (SIGCSE, ITiCSE and ICER), while there have by no means been an overwhelming number of papers about the teaching of social issues and ethics (only 21 or about 1 per cent of the total), there certainly have been a number of interesting attempts at articulating the best ways of teaching these issues and integrating them into the computer science curriculum (Moskal *et al.*, 2002; Califf and Goodwin, 2005; Sanders, 2005; Flieschman, 2006; Martin and Kuhn, 2006; Purewal *et al.*, 2007; Canosa and Lucas, 2008).

Outside of the ACM computing education conferences, ETHICOMP has been at the forefront of this effort. Since its inception in 1995, ETHICOMP has broadened the disciplinary scope of the conversation, bringing in scholars from philosophy, communications and other areas. Indeed, it should be no surprise that the field of computer ethics, which can be seen as a branch of applied ethics (Brey, 2000b; Marturano, 2002), would be of great interest to philosophers. As a consequence, ETHICOMP, with its explicit interest in ethical reflection about computing, has perhaps over time tended to attract and interest philosophers more than computing scientists. As such, it might be expected that this interest would be reflected in the papers presented at ETHICOMP.

This paper will begin with a brief bibliometric analysis of papers presented at ETHICOMP in the early years of the conference (1995, 1996 and 1998) and at more recent years (2010, 2011). This analysis will then buttress a more extended argument about the need for a reorientation in the approach taken in the teaching of SEI within computing disciplines. The core of this argument is that over time, the papers at ETHICOMP (as well as at the other computing education conferences) have consolidated around a specific analytic approach, namely, the articulation of the impacts of information and computing technology (ICT) and the ethical evaluation of those impacts.

While the focus on computing ethics might seem to be an exceptional admirable one (and from a philosopher's point of view, a necessary one), this paper will argue that an over-focus on ethical evaluation has a number of drawbacks. First, it sometimes devolves into a technological deterministic style of social explanation that has been in disrepute in the academic social sciences for decades. Second (and most important), in the classroom, it can result in an algorithmic approach to ethics that simplifies the social complexity and the uncertainty that is the reality of socio-technological change.

This paper will then conclude that the way computing faculties teach this course needs to move a bit away from the preoccupation with the ethical evaluation of ICT impacts and instead (re)emphasize the social context aspects. In particular, the teaching of this material needs to integrate the decades-old insights of researchers in sociology, communications and the history of technology which emphasizes the complex interaction and co-construction between the social environment and any given technology and the resulting uncertainty of technological change. Doing so would make the SEI course more focused on explicitly understanding the social and psychological contexts of computing and less focused on its ethical evaluation.

2. ETHICOMP: then and now

This paper makes broad claims about how the SEI of computing are taught. What is the evidence for our claims? We provide two types of evidence. One of our evidence types are the published accounts of teaching such a course within computing programs (covered in Section 3); our sample here includes the major computing education conferences and related publications. The second evidence type is a bibliometric analysis of the extended abstracts of papers presented at ETHICOMP (covered in this section).

It should be stressed that both types of evidences are only proxies meant to provide some insight into how the SEI of computing are taught within computing programs[1]. As proxies, they are potentially limited depending on how close the proxy models actual teaching practice. Nonetheless, we do think there is some value in bibliometric analysis. While it is possible that ETHICOMP papers might not be representative of the way the majority of ICT SEI courses are taught within computing programs, using textual and qualitative analysis techniques can nonetheless provide us with unique evidence for larger shifts in thinking about techniques and practice.

At the beginning of this project, we intended to evaluate each of the conference papers across a subset of years. Unfortunately, the full papers presented at ETHICOMP are not warehoused in any publicly accessible electronic archive. However, the extended abstracts are available and that is what we used in our evaluation. A research assistant gathered the title, year and abstract for each paper in a spreadsheet. The year was then hidden and the row order was randomized (this was done to minimize experimenter's bias). We then read each extended abstract and assigned a value between 1 and 5 (with 1 = little focus and 5 = great deal of focus) to each of the six categories shown in Table I.

Category	Description
Ethics focus	The extent to which the paper engaged in ethical and/or philosophic discourse
Social focus	Did the paper explicitly make use of social scientific approaches to study social, political or psychological issues with a technology?
Impact focus	How focused was the paper on ethically evaluating possible outcomes of a technology
Professional focus	Did the paper concern itself with professional codes of conduct?
Emerging focus	Was the paper explicitly concerned with addressing the problems of evaluating emerging technologies?
Privacy focus	Was the paper concerned with privacy issues?

Table I. Evaluation categories

Generally speaking, most years of ETHICOMP had 60-80 papers. As we were interested primarily in any thematic changes between the early and later years of the conference, we performed our categorization on just a subset of years. Our sample of early papers included 1995, 1996 and 1998 ($n = 125$)[2]; our sample of more recent papers included 2010 and 2011 ($n = 159$). Once gathered, the data are analyzed using SPSS software.

For the analysis, the data were divided into two groups (years 1995-1998 and years 2010-2011). While there were differences in means between the two groups, it was not clear whether these changes were statistically significant for the sample size. As such, an independent samples *t*-test was conducted for each of the measures to investigate differences between these two periods.

While our primary interest was in the Ethics Focus, Social Focus and Impact Focus measures, the reader may be interested in knowing our results for the other categories. For the Professional Focus category, there was a significant difference (a decline) in the scores for the earlier period (M = 1.70, SD = 1.18) and the later period (M = 1.37, SD = 0.87), $t(221) = -3.01$, $p = 0.03$. This result may reflect the fact that in the early years of ETHICOMP, there was more of an interest in articulating professional codes of conduct for ICT programs; once those codes were articulated, there has been apparently less of a need to readdress this issue. As well, there was a significant difference (an increase) in the privacy and the emerging technology category means between the groups. Privacy has clearly become one of the most important and relevant SEI topics in recent years, and our results reflect that interest. The increase in emerging technology issues in the later years is perhaps also reflected in Brey's (2012) anticipatory ethics.

For the Ethics Focus category, differences for the 1995-1998 period (M = 2.64, SD = 1.25) and the 2010-2011 period (M = 2.60, SD = 1.24) were not significant, $t(282) = 0.286$, $p = 0.775$. This was the expected result, as the relative ethical and philosophical content of ETHICOMP papers should have remained relatively stable from year-to-year.

For the Social Focus category, there was a significant difference in the scores for the earlier period (M = 2.61, SD = 1.08) and the later period (M = 2.17, SD = 1.08), $t(282) = 3.44$, $p = 0.001$. Similarly, a comparison of the Impact Focus scores for the earlier period (M = 1.95, SD = 0.96) and the later period (M = 2.66, SD = 1.15) also revealed a significant difference, $t(280.9) = -5.671$, $p = 0.00$. These results together strongly suggest that over the years, it became less and less common for ETHICOMP papers to make use of social scientific approaches to the study of ICT, and increasingly more common to instead ethically evaluate possible outcomes of a technology.

What does this shift in ETHICOMP papers really mean? Does it indicate a shift in how SEI of computing were taught? We believe that it likely does indicate that shift, but it should be stressed that this is purely a supposition. It is also possible that ETHICOMP authors are quite unrepresentative. To assess that possibility, we also examined how the teaching of this content has been reported within the computing education research literature.

3. Impacts and their influence

There has always been an interest within the ICT SEI community in understanding and articulating the possible effects of computer technology. Back in 1991, the ACM/IEEE report mentioned earlier, for instance, argued that computing "practitioners must be able to articulate the impact of introducing a given product line to a given environment" (Tucker, 1991). The National Science Foundation-funded ImpactCS project (Martin and Weltz, 1999) of the mid-1990s explicitly continued this focus on social impacts; the articulation of the main SEI content and the analytic framework it recommended for this content became an important resource for computing faculty in constructing their SEI courses. More recently, an important project called "Computing for the Social Good" (Goldweber, 2013) has received high-profile attention and uses a similar framework.

Computing ethics has also been similarly focused on examining the possible effects of computing technology. Moor (1985) argued that ethical issues emerge when transformative technologies are put to use and new practices reveal gaps in current policies. And as we have evidently entered an era of dramatically changing technology, we must redouble our efforts at improving the scope of computer ethics (Moor, 2005).

One response to Moore's call for improved computer ethics has been the creation of evaluative frameworks. Many of these frameworks have been technology-driven in that an inventory of possible impacts is constructed for the evaluated technological practice and those impacts are evaluated using the various normative theories to determine their ethical standing. Wright (2011) details a typical approach. Some frameworks are specifically aimed at applying ethics early in the process (Brey, 2012), and others apply ethical guidelines from other fields such as business (Bose, 2012). Attention is sometimes paid to make the approach more interdisciplinary with philosophers, computer scientists and social scientists all contributing (Brey, 2000b). In some cases, there is recognition that the focus should be moved from the technology and toward people, the users of the technology (Huff and Barnard, 2009; Stoodley *et al.*, 2010). These technology-driven, assessing impacts-based frameworks are often adopted in the classroom. For example, Lincke and Hudspeth (2013) describe a teaching module that identifies an issue (energy used by server farms), identifies a number of impacts and then assesses those impacts.

Our interest in this paper is not, however, in the evaluation of these frameworks or of computer ethics in general. Rather, our interest is in on how these topics are typically taught to computing students. It is our belief that the way that this material is taught in computing classrooms is often quite less sophisticated than the frameworks proposed by leading philosophers in this area.

In many textbooks and published reports of these courses (Moskal *et al.*, 2002; Quinn, 2006a, 2006b; Baase, 2007; Canosa and Lucas, 2008), one can see a very particular teaching approach. First, provide the students with at least two substantive moral theories, generally utilitarianism and a Kantian deontology. Depending on the time available, other moral theories (perhaps Aristotle's virtue ethics or Rawls's distributive justice) might also be covered. These different moral theories are then used to guide the students through the evaluation of the impacts caused by computer technology in a paradigmatic impact area such as privacy, intellectual property, security and access to information[3].

The appeal of this approach for computing faculty is not hard to see. It is attractive because it is so *algorithmic*. Most computing faculty achieved their position through their knowledge and research in traditional computing topics, and, as a consequence, the SEI course is often not in a computing professor's primary knowledge area. As one computing professor has noted, the "many gray areas of computer ethics are often frightening [...] to professors who are worried about how to answer things of which they themselves are unsure" (Grodzinsky, 2004). Not surprisingly then, a clear-cut methodology for teaching this course has in fact been argued (Moskal *et al.*, 2002;

Quinn, 2006a, 2006b) to be the way to make this course less imposing for computing faculty; this algorithmic approach is perhaps especially appealing because, as one survey found, the vast majority (84 per cent) of all computer SEI courses within computing programs are taught internally by computing faculty (Spradling *et al.*, 2008).

4. What is wrong with impacts?

Perhaps, the first step in recognizing the shortcomings of the ethical impacts approach is to realize the central flaw in the "articulate social impacts" step in the computing SEI course algorithm. This flaw is predicated on what seems an obvious and common-sense belief, namely, the belief that technology is simply a tool available for us to achieve our ends. This belief encourages us to examine computer technology in a means-ends manner; that is, the SEI researcher identifies and observes what consequence the means is going to have on the social environment. It generally assumes that the means are by and large clear and unproblematic and will always work in the same way for all people at all times. While the impacts approach sees ICT as a tool, it also sees it as a very special type of tool that can have large-scale impacts on society and/or the people using it. That is, while ICT, like a hammer, is just a tool, its special general-purpose nature means it has far-reaching effects outside its tool domain, akin to a hammer that changes the weather or weakens the dollar every time it strikes a nail. This approach to technology is generally called technological determinism by those who study the history, philosophy or sociology of technology (Pinch and Bijker, 1987; Nye, 2007). In this approach, technological change is treated as very much the independent variable in societal change. According to this view – see Bijker and Law (1992) and Smith and Marx (1994) for a summary of the debates – technological inventions, especially key ones like the printing press, the steam engine, the computer, the Internet and social networking, have transformed the world, and thus, new technologies need to be subjected to analysis to understand the wide-ranging transformation they have had on us and the world.

It is understandable why computer professionals find technological determinism attractive. After all, they are the people who are helping to invent some of these new technologies; it thus seems to feed their/our clear desire to be socially relevant (Purewal *et al.*, 2007) and to believe that computer geeks are actually the drivers of social change, and not politicians, business people or celebrities. This view is so widespread among computer professionals that, for instance, the authors' students and fellow computing department members find it difficult to believe that most current historians and sociologists of technology firmly reject technological determinism as being theoretically inconsistent as well as empirically under-supported. As one recent historian has noted, sweeping accounts "about machines that shape society remain popular, but they clash with the research of most professional historians of technology" (Nye, 2007). The academic field of science, technology and society (STS) studies that began in the 1960s has time and time again found that when examined carefully, most technologies rarely have had the effect that was expected – see, for instance, Pacey (1993) and Pool (1997) – and that the reason for this phenomenon is that "new technologies are shaped by social conditions, prices, traditions, popular attitudes, interest groups, class differences, and government policy" (Nye, 2007). Notice the direction of agency in this quote: it is technology that is being shaped or impacted by society, not the reverse.

Most technological deterministic impact prognosticators do their work by looking at the functional capabilities of a given technology and then imagining the impact of those functions. For instance, Internet search engines clearly make it easier to find knowledge; what then will be the impact of increased knowledge? Household technologies make it quicker to do housework; what will be the social impact of all that spare time? Antilock disc brakes make it less likely to skid and get into accidents; what will be the social impact of fewer accidents? In all these cases – and practically any other set of prognostications and impact evaluations that begin from an unquestioned belief that the functional capabilities of a technology (i.e. the means) do what is promised (i.e. achieve their ends) – the expected social impacts ended up being wildly wrong because the prognosticators believed in a naïve technological determinism. For instance, in his overview of the innumerable failed prognostications about the social impact of the Internet, Curran *et al.* (2012) sourly note that the cause of their faulty predictions was that they assessed "the impact of the internet not on the basis of evidence but on the basis of inference from internet technology".

Another relevant example is the introduction of household technology: it did not end up creating, in the words of Ruth Schwartz Cowan, less work for mother, but, in fact, more work because of a series of social changes that could not have been predicted if one limited one's analysis just to the functional capabilities of the household technologies. As Cowan (1983) demonstrated, household technologies created *more* housework due to changing expectations of what constitutes cleanliness (e.g. clothes changed daily instead of weekly), new unexpected technologies enabled by the technology (e.g. wall-to-wall carpets were unknown before vacuums) and the gradual displacement of household work done by external agents (e.g. laundry services, maids, nannies) to housewives partly as a consequence of household technologies and partly due to exogenous changes in the social and economic realm. Similarly, efficient Internet search engines have not resulted in people with too much knowledge; instead, unpredicted changes in how people interact with words (scanning replacing reading) and even possibly cognitive decline due to the brain's plasticity have arguably resulted in the exact opposite consequence (Carr, 2010). And as is readily apparent to anyone who actually drives an automobile, the introduction of anti-lock disc brakes has not reduced accidents at all, partly because drivers tend to drive faster and tailgate more closely due to the improved braking technology and also partly because of increases in the intensity of traffic due to ongoing changes in urban geography (Vanderbilt, 2008).

The first step then that we should take in our SEI course is to communicate how rarely technologies achieve their exact promise and, indeed, how many appear to do the opposite. This so-called Revenge Effect is well-documented (Tenner, 1997), and yet, we were unable to find it discussed at all in existing computer ethics textbooks or in published accounts of this course. As well, an equally important step we need to make in the teaching of the SEI course is to reject naïve technological determinism and help students understand the complex agency issues in the relationship between technology and society. One way to achieve this goal would be by beginning the SEI course with examples and readings in how certain vital technologies had little impact on some societies, or on how certain technologies were strongly modified and differently adapted in different cultures and countries. This more historically nuanced (and significantly more empirically accurate) approach is what is generally called social constructivism (Feenberg, 2000). In this approach, one looks at how technologies are researched, invented, financed, developed, adopted, marketed and propagated within the very complex system generally referred to as society. If one carefully examines a given technology within the social system in which it is embedded, it becomes extremely difficult to maintain a belief in technological determinism. Instead, one sees technologies much more strongly "impacted" by society rather than vice versa. For instance, one recent overview of the research about the social consequences of the Internet summarizes: "the weight of evidence points to one firm conclusion: society exerts, in general, a greater influence on the internet than the other way around" (Curran *et al.*, 2012). For these reasons, it is essential that our SEI course should integrate the historically and sociologically grounded insights of the STS research community. In other words, our computing SEI course should look more like a historical sociology course and a little less like a philosophic ethics course.

5. The importance of uncertainty

The reason why revenge effects occur is due to the fact that "socio-technological transformation is a highly complex process which involves many uncertainties" (Meijer *et al.*, 2006). While uncertainty is a key concept in fields like economics, management, medicine, environmental science and a variety of other disciplines, within moral philosophy in general, and computer ethics in particular, it is underappreciated (Walker, 2003). This is an important problem because the substantive moral theories (such as deontology and utilitarianism) that are the bedrock of the usual computer ethics course require relatively clear and unambiguous information about effects to make judgments (Floridi, 1999; Floridi and Sanders, 2002; Walker, 2003). Typical problems or dilemmas for which macro-ethical approaches are applied are most often done in a context of complete knowledge (if you do action X, then Y people will be harmed, but Z people will be benefited). This is appealing for computer scientists, who often work with problems modeled by idealized abstractions for which complete knowledge is possible. The ethics of technology, by contrast, should be recognized as residing in a context of at least partial uncertainty or ambiguity. Furthermore, the degree of uncertainty is greater for emerging technologies, and the more complex the technology, the more uncertain we are as to the developmental trajectory of a technology (Pool, 1997). Shirky (2010) has similarly noted that the more complex and open-ended the new technology, "the less completely anyone can extrapolate the future from the previous shape of technology".

There are many places where uncertainty intersects with the life cycle of a technology. We can visualize this as shown in Figure 1 (Sollie, 2007), which shows how a technology may be used/repurposed/deployed in a wide variety of different ways across time (or perhaps across different cultural contexts).

For the evaluating agent x (the professor or the student or the developer or the journalist) examining emerging technology y at time t_1, the agent must have knowledge of the development trajectory to morally evaluate it. Unfortunately, as we have already seen, we very often cannot know the actual development trajectory and if we hypothesize one based purely on its functional capabilities, we more often than not will be woefully wrong. Thus, in reality, we need to recognize that technology y has multiple trajectories (o_1, o_2, o_3, …) and that it might be more important to attempt to understand which trajectories are more likely (by untangling the complex web of history, interests and agents) than applying a prescriptive ethical judgment on a single predicted trajectory. Speaking in the context of social theory, Couldry (2012) has noted that there may even be six types of uncertainty facing anyone considering analyzing the impacts of technology in today's complex, highly differentiated societies, and as a consequence,

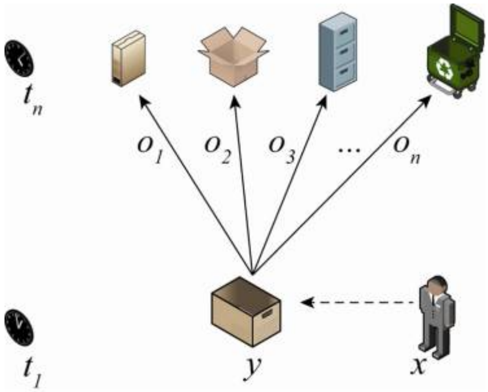

Figure 1.
Problems with evaluating technology

advocates a very cautious, sociologically focused approach to the consideration of technological impacts.

Brey (2012) grappled with this particular quandary in his paper on anticipatory technology ethics. He argues that there are two reasons why it is possible to do ethical evaluation on technology early in its life cycle. One of the reasons is that the emerging technology may represent generic ethical issues that can readily be philosophically interrogated independent of the technology; the other reason is that the philosopher can make use of forecasting and future studies. Based on the exceptional lack of accuracy by, for instance, early expert forecasting about the Internet in the 1990s or social networks in the mid 2000s (Curran *et al.*, 2012; Connolly, 2001), we are skeptical about the effectiveness of any ethics based on foresight analysis (Stahl, 2011) or future studies. As mentioned in the previous section, such futurist predictions are almost always predicated on assessing future outcomes based on the functional capabilities of technology, an endeavor which historical evidence shows is rarely accurate when viewed retrospectively.

6. Beyond moral evaluation

One could argue that while uncertainty may be a feature of many aspects of life, we still manage to make practical and ethical decisions. That is, just because the arc of technological development is uncertain, that does not suddenly excuse us from making moral judgments. While this may be true, the argument made in this paper is that moral evaluation (i.e. normative or prescriptive ethics) in the uncertain realm of socio-technological change should perhaps be only tentative at best. That is, we believe that the ethical discussion in the computing SEI course should be mainly descriptive ethics. In descriptive ethics, the focus is on discovering and articulating ethical perspectives or issues. As a consequence, it is the authors' belief that the approach we should be taking in the computing SEI course is to expose the students to the many levels of uncertainty in the domain of technological evaluation, thereby allowing the students to achieve a level of critical awareness that weaves some ethical insight into a richer understanding of the complex nature of socio-technological change.

As Brey (2000a) has noted, the main problem with contemporary computer ethics is that the analytic effort is limited to evaluating well-recognized morally controversial technological practices for which there is a policy or legal vacuum. While not unimportant certainly, the problem here, as discussed previously, is that it either depends on a level of epistemic certainty that may be unwarranted or it may be evaluating a technological system that is so firmly established that its momentum has moved it beyond the effective reach of moral prescription (Hughes, 1994). Furthermore, it tends to ignore the problems embedded in morally nontransparent computer practices. A promising alternative approach to the usual macro-ethical evaluation in SEI courses is to pursue what Brey (2000b) has referred to as disclosive ethics. In this approach, rather than applying moral theories to well-known impacts, the focus is on disclosing the assumptions, values and interests built into the design, implementation and use of technology. This type of activity could even be a cooperative venture "in which philosophers, computer scientists and social scientists collaborate to disclose embedded normativity in computer systems and practices". Such a course would then try to guide the students in the unpacking of the normative assumptions of real-world computer practice, which, in turn, requires clarifying the complex web of social interactions that play a role in constructing the different trajectories a technology may take.

This approach is not about providing bright-line tests for evaluating the moral rightness and wrongness of technological practice but about opening up the black box of technological practice for understanding and to interrogate the assumptions, the embedded power relations and the rhetorical and ideological contexts that we bring to those practices. While this may seem to some an abrogation of moral duty, it is perhaps the path that may end up being ultimately more socially responsible because it exposes our computing students to a more socially- nuanced understanding of their profession. Looking at the description of the core components in the SEI knowledge area in, for instance, the earlier CS-2008 report (ACM Interim Review Task Force, 2008), there was, in fact, a recognition that it is equally important to have students appreciate the social and historical context of computing as it is to ethically evaluate it. Similarly, in our analysis of ETHICOMP papers over time, we saw that in the earlier years of the conference, it was much more common for papers to make use of social science approaches. We (in computing that is) need to return our principal analytic focus in the SEI course back to this social context and focus the students' attention on different perspectives and on the complexities of socio-technological change.

There has already been a somewhat similar intellectual transition in the fields of communication that study the interactions between new technologies of communication and the broader society. This new research approach is generally referred to as "mediatization" (Livingstone, 2009; Finnemann, 2011). Moving forward from the media effects tradition within communications (one which is quite close in spirit to the impacts approach to technological evaluation), scholars using the new analytic approach of mediatization begin with the recognition that a technology works at four levels: as a technology, as a social institution, as a context for other technologies and as a cognitive construct in the experience of the user (Krotz, 2009). That is, mediatization theory attempts to capture how media technology (especially new digital media) shapes *and* is shaped by broader transformations in culture and society. Like with the social constructivist approach mentioned earlier, in mediatization, technology is not a pool cue driving history but is "entangled" in cultural and social change (Couldry and Hepp, 2013). As a consequence, the best approach for untangling is via careful and nuanced empirical research into all four working levels of technology. This mediatization research approach in communications could be a potential model for our ICT SEI education moving forward.

7. Limitations

While the previous section argued that the teaching of ICT SEI should be less philosophic and more sociologically oriented in its methodological approach, we should add several caveats. First, it should be stressed that technological innovation does affect individuals and social groups. Indeed, some of the most compelling recent research in this area has been around trying to better understand the psychological changes individuals appear to undergo in the face of new specific technologies. We are not advocating that teachers of SEI course should ignore these potential effects. It should be noted that such evidence-based studies generally make use of a social science repertoire of methods. While is important to also problematize such studies so that students understand their difficulties and limitations, we do think that the basic sociological and psychological orientation of such studies is what our students need more of in their ICT SEI course.

Second, it should be emphasized as well that this is not an easier way to teach an ICT SEI course. As noted by some practitioners (Moskal *et al.*, 2002; Quinn, 2006a, 2006b), the great benefit of the algorithmic ethics approach is that it provides a straightforward approach for teaching this material. It is even possible that it results in students who are more sensitive to the broader consequences of their own computing practices. Nonetheless, we believe that a less algorithmic approach would better achieve that laudable aim – though we of course have no concrete evidence for this claim other than the types of arguments we have made in this paper.

8. Conclusion

Computing has been immeasurably improved by the many dedicated scholars who have worked tirelessly to convince the rest of the field about the importance of SEI. While most computing programs now do include a dedicated SEI course in their curriculum, there is often some anxiety associated with teaching this course due to a perceived lack of the appropriate knowledge by computing faculty. As a consequence, there has been a convergence on a particular approach to teaching this course: articulate the impacts of different computer technologies and then apply a moral theory such as utilitarianism or deontology to those impacts. Evidence for the increasing prevalence of this approach was provided in this paper via the proxy measures such as that obtained via a bibliometric analysis of ETHICOMP papers. Our analysis of the conference papers showed a significant decrease in social scientific approaches over time that was mirrored by a significant increase in ethical impact analysis. This paper has also argued that this approach leans too heavily on a social theory, namely, technological determinism, which is both theoretically unsound and empirically under-supported. Furthermore, this paper also argued that the use of large moral theories may be inappropriate for domains in which there is a strong lack of epistemic certainty. As the field of technological change is indeed characterized by substantial uncertainty, this paper concluded that we need to transform the way we teach the computing SEI course so that it is more closely allied with the insights of the broader technology-and-society research communities within sociology, communications and history. Doing so would make the SEI course much more focused on explicitly understanding the social contexts of computing and less focused on its ethical evaluation.

Notes

1. Note that our focus is on the way that SEI content is typically taught within computing departments and not within philosophy or sociology departments.
2. There wasn't an ETHICOMP in 1997.
3. It should be stressed that not all ITC SEI textbooks follow such an explicitly algorithmic approach. Yet, as one textbook author told me, the publishers encourage this approach because they feel that this is the preference of those teaching the courses.

References

ACM Interim Review Task Force (2008), *Computer Science Curriculum 2008: An Interim Revision of CS 2001*, ACM Interim Review Task Force.

ACM/IEEE-CS Joint Task Force on Computing Curricula (2004a), *Computer Engineering 2004: Curriculum Guidelines for Undergraduate*, ACM Press and IEEE Computer Society Press.

ACM/IEEE-CS Joint Task Force on Computing Curricula (2004b), *Software Engineering 2004: Curriculum Guidelines for Undergraduate Degree Programs in Software Engineering*, ACM Press and IEEE Computer Society Press.

ACM/IEEE-CS Joint Task Force on Computing Curricula (2013), *Computer Science Curricula 2013*, ACM Press and IEEE Computer Society Press.

Baase, S. (2007), *A Gift of Fire: Social, Legal, and Ethical Issues for Computing and the Internet*, 3rd ed., Prentice Hall.

Bijker, W. and Law, J. (1992), *Shaping Technology/Building Society: Studies in Sociotechnical Change*, The MIT Press.

Bose, U. (2012), "An ethical framework in information systems decision making using normative theories of business ethics", *Ethics and Information Technology*, Vol. 14 No. 1.

Brey, P. (2000a), "Disclosive computer ethics", *Computers and Society*, Vol. 30 No. 4.

Brey, P. (2000b), "Method in computer ethics: towards a multi-level interdisciplinary approach", *Ethics and Information Technology*, Vol. 2 No. 2.

Brey, P. (2012), "Anticipating ethical issues in emerging IT", *Ethics and Information Technology*, Vol. 14 No. 4.

Califf, M. and Goodwin, M. (2005), "Effective incorporation of ethics into courses that focus on programming", *Proceedings of the 36th Technical Symposium on Special Interest Group on Computer Science Education, St. Louis, MO*.

Canosa, R.L. and Lucas, J. (2008), "Mock trials and role-playing in computer ethics courses", *Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education, Portland, OR*.

Carr, N. (2010), *The Shallows: What the Internet is Doing to Our Brains*, Norton.

Connolly, R. (2001), "The rise and persistence of the technological community ideal", in Werry, C. and Mowbray, M. (Eds), *Online Communities: Commerce, Community Action, and the Virtual University*, Prentice-Hall.

Couldry, N. (2012), *Media, Society, World: Social Theory and Digital Media Practice*, Polity Press.

Couldry, N. and Hepp, A. (2013), "Conceptualizing mediatization: contexts, traditions, arguments", *Communication Theory*, Vol. 23 No. 3.

Cowan, R. (1983), *Less Work for Mother*, Basic Books.

Curran, J., Fenton, N. and Freedman, D. (2012), *Misunderstanding the Internet*, Routledge.

Feenberg, A. (2000), "From essentialism to constructivism: philosophy of technology at the crossroads", in Higgs, E. (Ed.), *Technology and the Good Life?*, The University of Chicago Press.

Finnemann, N.O. (2011), "Mediatization theory and digital media", *Communications*, Vol. 36 No. 1.

Flieschman, W.M. (2006), "Meta-informatics and ethical issues in computing", *ACM SIGCSE Bulletin*, Vol. 38 No. 3, pp. 232-236.

Floridi, L. (1999), "Information ethics: on the philosophical foundation of computer ethics", *Ethics and Information Technology*, Vol. 1 No. 1.

Floridi, L. and Sanders, J.W. (2002), "Mapping the foundationalist debate in computer ethics", *Ethics and Information Technology*, Vol. 4 No. 1.

Goldweber, M. (2013), "A framework for enhancing the social good in computing education: a values approach", *ACM Inroads*, Vol. 4 No. 1, pp. 58-79.

Goldweber, M., Barr, J., Clear, T., Davoli, R., Mann, S., Patitsas, E. and Portnoff, S. (2013), "A framework for enhancing the social good in computing education: a values approach", *ACM Inroads*, Vol. 4 No. 1, pp. 58-79.

Grodzinsky, F. (2004), "Panel: responding to the challenges of teaching computer ethics", *SIGCSE'04: Proceedings of the 35th SIGCSE Technical Symposium on Computer Science Education*, *Norfolk, VA*.

Huff, C. and Barnard, L. (2009), "Good computing", *IEEE Information Technology and Society Magazine*, Vol. 28 No. 3, pp. 47-54.

Hughes, T.P. (1994), "Technological momentum", *Does Technology Drive History? The Dilemma of Technological Determinism*, The MIT Press.

Krotz, F. (2009), "Mediatization: a concept with which to grasp media and societal change", in Lundry, K. (Ed.), *Mediatization: Concept, Changes, Consequences*, Peter Lang.

Lincke, S.J. and Hudspeth, C.S. (2013), "Energy efficiency: teaching for accreditation, ethics and technology", *Frontiers in Education Conference*, Oklahoma City, OK, pp. 484-490.

Livingstone, S. (2009), "On the mediation of everything: ICA presidential address 2008", *Journal of Communication*, Vol. 59.

Lunt, B.M., Ekstrom, J.J., Gorka, S., Hislop, G., Kamali, R., Lawson, E., LeBlanc, R., Miller, J. and Reichgelt, H. (2008), *Information Technology 2008: Curriculum Guidelines for Undergraduate Degree Programs in Information Technology*, ACM Press and IEEE Computer Society Press.

Martin, C.D. and Weltz, E. (1999), "From awareness to action: integrating ethics and social responsibility into the computer science curriculum", *Computers and Society*, Vol. 29 No. 2.

Martin, F. and Kuhn, S. (2006), "Computing in context: integrating an embedded computing project into a course on ethical and societal issues", *SIGCSE'06: Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education*, *Houston, Texas*.

Marturano, A. (2002), "The role of metaethics and the future of computer ethics", *Ethics and Information Technology*, Vol. 4.

Meijer, I., Hekkert, M., Faber, J. and Smits, R. (2006), "Perceived uncertainties regarding socio-technological transformations: towards a framework", *International Journal of Foresight and Innovation Policy*, Vol. 2 No. 2.

Moor, J.H. (1985), "What is computer ethics?", *Metaphilosophy*, Vol. 16 No. 4, pp. 266-275.

Moor, J.H. (2005), "Why we need better ethics for emerging technologies", *Ethics and Information Technology*, Vol. 7 No. 3, pp. 111-119.

Moskal, B., Miller, K. and King, L. (2002), "Grading essays in computer ethics: rubrics considered helpful", *SIGCSE'02: Proceedings of the 33rd SIGCSE Technical Symposium on Computer Science Education*, *Cincinnati, KY*.

Nye, D. (2007), *Technology Matters*, The MIT Press.

Pacey, A. (1993), *The Maze of Ingenuity: Ideas and Idealism in the Development of Technology*, The MIT Press.

Pinch, T. and Bijker, W. (1987), "The social construction of facts and artifacts: or how the sociology of science and the sociology of technology might benefit each other", *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*, The MIT Press.

Pool, R. (1997), *Beyond Engineering: How Society Shapes Technology*, Oxford University Press.

Purewal, T., Bennett, C. and Maier, F. (2007), "Embracing the social relevance: computing, ethics and the community", *SIGCSE'07: Proceedings of the 38th SIGCSE Technical Symposium on Computer Science Education*, *Covington, Kentucky*.

Quinn, M. (2006a), "Case-based analysis: a practical tool for teaching computer ethics", *SIGCSE'06: Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education*, *Houston, Texas*.

Quinn, M. (2006b), *Ethics for Information Age*, 2nd ed., Addison-Wesley, Houston, Texas.

Sanders, A.F. (2005), "A discussion format for computer ethics", *SIGCSE'05: Proceedings of the 36th SIGCSE Technical Symposium on Computer Science Education*, *St. Louis, MO*.

Shirky, C. (2010), *Cognitive Surplus*, Allen Lane.

Smith, M. and Marx, L. (1994), *Does Technology Drive History? The Dilemma of Technological Determinism*, The MIT Press.

Sollie, P. (2007), "Ethics, technology development and uncertainty: an outline for any future ethics of technology", *Journal of Information, Communication & Ethics in Society*, Vol. 5 No. 4.

Spradling, C.L., Soh, L. and Ansorge, C. (2008), "Ethics training and decision-making: do computer science programs need help?", *SIGCSE'08: Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education*, *Portland, OR*.

Stahl, B. (2011), "IT for a better future: how to integrate ethics, politics and innovation", *Journal of Information, Communication & Ethics in Society*, Vol. 9 No. 3.

Stoodley, I., Bruce, C. and Edwards, S. (2010), "Expanding ethical vistas of IT professionals", *Information Systems Frontiers*, Vol. 12, pp. 379-387.

Tenner, E. (1997), *Why Things Bite Back: Technology and the Revenge of Unintended Consequences*, Vintage.

Topi, H., Valacich, J., Wright, R., Kaiser, K., Nunamaker, J., Sipior, J. and de Vreede, G. (2010), *IS 2010: Curriculum Guidelines for Undergraduate Degree Programs in Information Systems*, ACM Press.

Tucker, A.B. (1991), "Computing curricula 1991", *Communications of the ACM*, Vol. 34 No. 6.

Vanderbilt, T. (2008), *Traffic: Why We Drive the Way We Do (and What It Says About Us)*, Knopf.

Walker, W. (2003), "Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support", *Integrated Assessment*, Vol. 4 No. 1.

Wright, D. (2011), "A framework for the ethical impact assessment of information technology", *Ethics in Information Technology*, Vol. 13, pp. 199-226.

Further reading

Connolly, R. (2011), "Beyond good and evil impacts: rethinking the social issues components in our computing curricula", *Proceedings of the 16th Annual Conference on Innovation and Technology in Computer Science Education, Darmstadt*.

Corresponding author

Randy Connolly can be contacted at: rconnolly@mtroyal.ca