Emerald Insight

Journal of Information, Communication and Ethics in Society

Computer ethics beyond mere compliance

Richard Volkman

Article information:

To cite this document:
Richard Volkman , (2015),"Computer ethics beyond mere compliance", Journal of Information, Communication and Ethics in Society, Vol. 13 Iss 3/4 pp. 176 - 189
Permanent link to this document:
http://dx.doi.org/10.1108/JICES-10-2014-0043

Downloaded on: 10 November 2016, At: 21:11 (PT)
References: this document contains references to 19 other documents.
To copy this document: permissions@emeraldinsight.com
The fulltext of this document has been downloaded 235 times since 2015*

Users who downloaded this article also downloaded:

(2015),"Does computing need to go beyond good and evil impacts?", Journal of Information, Communication and Ethics in Society, Vol. 13 Iss 3/4 pp. 190-204 http://dx.doi.org/10.1108/JICES-10-2014-0045

(2015),"Mentorship in computer ethics: ETHICOMP as a “community mentor” for doctoral and early career researchers", Journal of Information, Communication and Ethics in Society, Vol. 13 Iss 3/4 pp. 326-345 http://dx.doi.org/10.1108/JICES-10-2014-0052

Access to this document was granted through an Emerald subscription provided by emerald-srm:563821 []

For Authors

If you would like to write for this, or any other Emerald publication, then please use our Emerald for Authors service information about how to choose which publication to write for and submission guidelines are available for all. Please visit www.emeraldinsight.com/authors for more information.

About Emerald www.emeraldinsight.com

Emerald is a global publisher linking research and practice to the benefit of society. The company manages a portfolio of more than 290 journals and over 2,350 books and book series volumes, as well as providing an extensive range of online products and additional customer resources and services.

Emerald is both COUNTER 4 and TRANSFER compliant. The organization is a partner of the Committee on Publication Ethics (COPE) and also works with Portico and the LOCKSS initiative for digital archive preservation.

*Related content and download information correct at time of download.

The current issue and full text archive of this journal is available on Emerald Insight at:
www.emeraldinsight.com/1477-996X.htm

JICES
13,3/4

Received 7 October 2014
Revised 24 January 2015
Accepted 3 March 2015

Computer ethics beyond mere compliance

Richard Volkman
Department of Philosophy, Southern Connecticut State University, New Haven, Connecticut, USA

Abstract

Purpose – The paper aims to examine the nature of computer ethics as a field of study in light of 20 years of Ethicomp, arguing that computer ethics beyond mere compliance will have to be pluralistic and sensitive to the starting places of various audiences.

Design/methodology/approach – The essay offers a philosophical rather than empirical analysis, but the ideal of open inquiry is observed to be manifest in the practice of Ethicomp.

Findings – If computer ethics is to constitute a real engagement with industry and society that cultivates a genuine sensitivity to ethical concerns in the creation, development and implementation of technologies, a genuine sensitivity that stands in marked contrast to ethics as "mere compliance", then computer ethics will have to persist in issuing an open invitation to inquiry.

Originality/value – The celebration of 20 years of Ethicomp is an occasion to reflect on who we are and what we mean to be doing. Inclusive of previous accounts (e.g. Moor and Gotterbarn), while going beyond them, an inquiry-based conception of computer ethics makes room for all the various dimensions of computer ethics.

Keywords Philosophy, Computer ethics, Collaboration, Information ethics, Community evolution, Teaching innovations

Paper type Conceptual paper

1. Introduction

One reason to celebrate birthdays, anniversaries and regular holidays is to punctuate our histories by periods of reflection that reaffirm the importance of what we care about and the trajectory of the story it inscribes. Graduations, weddings and retirements are likewise opportunities to reexamine core values and projects, but our celebrations of change and transition more explicitly acknowledge the perils and opportunities of a future that may not resemble the past. All this is as true of the stories of institutions as of individuals. Terry Bynum and Simon Rogerson, having co-founded the Ethicomp conference series two decades ago, are now handing over stewardship of this institution to a new generation of students, scholars and practitioners. Our celebration of 20 years of Ethicomp thus invites us to ask: Who have we been, and what shall we become? This means we need to ask again, in light of everything we have learned over the past two decades: What is computer ethics? How does one *do* computer ethics? What are the hopes and fears that inform our path forward?

Such questions were plainly lurking behind Rogerson's (2014) provocative plenary address to kick-off Ethicomp 2014. Having long championed the relevance of professional practice in the science and industry of computing, it came as no surprise that Rogerson would reaffirm his view that computer ethics ought to make some tangible, practical difference in the way computing artifacts are designed, developed

Journal of Information, Communication and Ethics in Society
Vol. 13 No. 3/4, 2015
pp. 176-189
© Emerald Group Publishing Limited
1477-996X
DOI 10.1108/JICES-10-2014-0043

and implemented. The more provocative part of the address came in its tone of dire warning against the many and various temptations that might lead the computer ethics community away from these practical goals. Rogerson cautioned against a computer ethics that might become narrowly academic, detached from real and pressing practical concerns. There is ethical work enough in dealing with what our arts and sciences have brought into the real world without getting distracted by imagined futures and the many worlds of science fiction. Rogerson went on to warn that, unless we remain diligent against it, computer ethics might come to be driven by personal and professional ambitions that prioritize funding, grants and establishing academic or industry fiefdoms over influencing the ethical culture of computing as a profession. Actually, improving the ethical practice of computing should set the agenda of computer ethics, rather than having the agenda of computer ethics set according to whatever promises the most funding or prestige. Ultimately, Rogerson advised that computer ethics should constitute a real engagement with industry that cultivates a genuine sensitivity to ethical concerns in the creation, development and implementation of technologies, a genuine sensitivity that stands in marked contrast to ethics as “mere compliance”. Although there was a great deal more to his talk, it was this contrast between *genuine sensitivity* and *mere compliance* that especially captured my attention. If such a contrast is to animate the practice of computer ethics going forward, as I agree it should, then we need to get clear about the implications of this distinction for our conception of computer ethics.

Computer ethics is a diverse and interdisciplinary field engaging a variety of audiences regarding a variety of issues from a variety of perspectives and levels of abstraction, so it comes as no surprise that there are many ways to engage the broad project of cultivating genuine sensitivity. Despite this great diversity, I suspect that nearly all of these approaches can agree that there is something amiss in an ethical project that settles for mere compliance. I will explore a few of the reasons why mere compliance falls so far short of our aspirations in computer ethics, and I will argue that the aspiration to cultivate genuine sensitivity requires a special concern for the importance of what our various audiences already care about. To cultivate ethical sensitivity, we must bring our audience either to care about something they did not care about before or to realize the full implications of something they already care about. Either way, because caring about something is not generally under one’s immediate voluntary control, the relationship will involve inquiry-based teaching and mentoring rather than content-based legislating and enforcement. In short, ethics is not simply a matter of policy or learning the right set of truths by rote. Ethics is all about asking questions and reflecting on the answers for oneself. I will argue for a conception of computer ethics as an invitation to reflection on the ethical significance of computing as a practice, a profession and an institution. Such a conception of computer ethics is not exclusionary, as it is consistent with a wide diversity of other commitments, projects and goals. The claim is simply that, whatever else computer ethics is and does, whatever other theoretical and practical commitments inform the practices of particular computer ethicists, it will have to pay special attention to the intrinsic motivations and concerns of its audience and appeal to those interests and values if it is to cultivate genuine sensitivity rather than mere compliance. As I will show, some ways of doing or conceiving computer ethics will be more suited to this than others, at least before certain audiences. It behooves the computer ethicist to appreciate exactly where her efforts interface with the practice of computing so she can practice computer ethics in a manner that will cultivate genuine sensitivity and not mere compliance within her specific audience. For many of us, that interface will come in the classroom as we invite aspiring computer professionals and end users to join us in appreciating the ethical significance of computing as a practice, a profession and an institution. For others, the interface between computer ethics and the professions and industries associated with computing will involve direct contact with engineers, managers and other professionals in the field. In either case, the principle job of the computer ethicist will be as a teacher and a mentor inviting others to a field of inquiry.

2. The ethical shortcomings of mere compliance

We describe someone as "compliant" when he operates in obedience to some external authority and not from his own inner motivations or interests. I may "comply" with your wishes, but I do not "comply" with my own. The merely compliant person has abdicated his own authority to another, leaving no genuine will of his own. On certain accounts, including especially those influenced by Kant, an action done from obedience or interest in extrinsic considerations can have no moral worth; on such views, mere compliance may lead to actions that accord with duty, but only actions that issue from one's own autonomous will, as a will legislating moral law can count as fully moral acts. Such views will have obvious reasons to shrink from an ethics of mere compliance, but even consequentialist and other views that do not make motives internal to the moral worth of action have reasons to encourage ethical motives, as having the right motives plausibly leads to better consequences. The point can be illustrated in a simple thought experiment. In which circumstance should one feel more secure in one's privacy:

- officials care nothing for privacy but are anxious to be in compliance with the letter of the law to avoid punishment; or
- officials care about privacy exactly as citizens care about it and with the same affective appreciation?

Going beyond compliance does not change or challenge any particular theory of right action, but it makes a difference for whether, how and to what extent we should expect agents to live up to spirit of whatever ethical aspirations they should ultimately have. Thus, even on a consequentialist account, concern for ethics beyond mere compliance is desirable if instilling motives beyond mere compliance will tend to lead to better consequences overall. Compliance is "sufficient" for right action, defined in narrowly utilitarian terms, but whether one will be more or less likely to behave in accord with the demands of right action plausibly depends on whether one also cares about what one is doing. That a concern for motives matters for the sophisticated consequentialist is *not* a controversial or new claim (Parfit, 1984; Railton, 1988). For consequentialists, motive may not matter in itself (i.e. motives may not be intrinsically valuable), but our motives have consequences (i.e. they have an instrumental value), and these extended consequences do matter.

Extrinsic considerations, including rewards and punishments, only motivate insofar as the reward or punishment is a plausible consequence of compliance or noncompliance. Thus, compliance is only a reliable means to good consequences insofar as noncompliance might be found out. Not only will merely compliant actors be tempted to do the wrong thing when no one is looking, but the regulatory and other authorities exacting mere compliance will have reason to count as ethically significant only those features of an act that can be straightforwardly found out. Along this path, ethics risks being conflated with law, both in practice and in theory (Ladd, 1980). This is inevitable insofar as enforcement is the hallmark of law and an ethics of mere compliance depends on enforcement. Law and ethics are not the same thing, and ethical and legal injunctions cannot even be made coextensive insofar as ethics aspires to an excellence law and its cousins in command and control regulation cannot hope to attain. It is easy to enforce compliance with a rule of the form, "You are obliged to do x", but it is very difficult to enforce a rule that says, "You are obliged to do x well". It is quite impossible to enforce a rule that says, "You must do your best to do x as well as you can". There is no reliable way to detect noncompliance with such a rule, so there is no way to exact compliance. Thus, mere compliance is liable to lead to an ethics of "just good enough to satisfy the surveilling authorities", unless the external considerations that exact compliance can be devised as incremental rewards for excellent performance, incentivizing agents to make their best efforts and to improve upon their past best efforts.

There are mechanisms that mean to associate incremental rewards with performance in a decentralized manner, avoiding much of the fragility of detection suffered by regimes of oversight, and these may go some long way to improving performance along various dimensions, including ethical dimensions. Market institutions, for example, do a reasonable job of aligning incentives with performance such that one can expect greater rewards for doing better than some bare minimum. It is generally a good thing if success in the marketplace incentivizes ethical conduct, and we have good reason to structure the legal and regulatory regimes that define market institutions such that ethical conduct is rewarded in the market. This can go a long way towards encouraging ethical compliance. As things stand now, vendors with a reputation for shoddy work or a disregard for clients' interests will certainly find themselves at a disadvantage in a competitive environment; however, as is true of compliance by way of enforcing rules, these incentives to compliance only operate insofar as ethical failures and successes are readily detected and communicated, and even then market incentives only extend as far as they confer some competitive advantage. Although markets dynamically access, process and distribute vastly more information than command and control regulatory strategies, the problems of detection and enforcement are only diminished and not eliminated in market solutions. Markets enforce compliance only so far as desirable outcomes are represented in consumer demand expressed in purchasing choices between available competing products. Along any given dimension of value to consumers, earning immediate market rewards typically requires merely being no worse than others at the same price point. If I have few competitors or if all my competitors are ethically incompetent, it will not take much to outshine them and enjoy the market rewards that come with a good reputation. Of course, markets are dynamic information processing systems that encourage innovation to better satisfy consumer demand, even to the point of discovering and educating consumers to values of which they were formerly unaware, so there is market pressure to outdo competitors along whatever dimensions consumers might come to care about, and this can and often does include various dimensions of ethical performance. Market incentives to progressively better ethical conduct will be especially strong if there is a prevailing culture of ethical sensitivity. If customers are sensitive to these sorts of concerns (either on their own or as a result of marketing from ethically sensitive competitors), such that a clear competitive advantage is to be had from enjoying a good reputation, then markets can quickly produce a "race to the top" as firms invest in building and maintaining their reputations, eager to win customers for their high levels of ethical performance.

This might be considered the best case scenario for getting good outcomes from ethics as mere compliance, but even here the limits established by the boundaries of ready detectability are pressing. Especially in the professions, characterized in general by the need for special expertise that the general public does not possess, we cannot count on customers to reliably detect ethical failures. Insofar as clients are at the mercy of information about ethical performance from spokespersons of firms and other interested parties, they are susceptible to all manner of spin and hype as firms manage their reputations with savvy public relations instead of building a good reputation in the old-fashioned way of earning it. If it is possible to dupe the public on such matters, there will be an incentive to do so. These considerations loom all the larger insofar as the industry or profession is dynamic, innovative and revolutionary. In that case, as new firms, new providers and new products enter and leave the market at dizzying speeds, old-fashioned reputation may be difficult to earn, and when earned, it may be impossible to capitalize on it. Because no one lacking appropriate professional expertise, including especially the buying public, can be expected to follow all the nuances and subtleties of a rapidly changing industry and profession, consumers will have to rely on these professionals themselves to make these judgments, and this means there is little the market can do to discipline a profession that does not commit itself to the genuine pursuit of ethical excellence. Computing needs professional ethics at least because it must stand in a fiduciary relationship to the public (and perhaps for other reasons as well. Cf. Gotterbarn, 1997). As with other professions, but to an even greater degree in a relatively new and quickly evolving field like computing, market incentives to ethical conduct depend on the prevailing ethical sensitivity of professionals or clients or both, and this ethical sensitivity is not reliably bootstrapped by the market itself, at least not in any short term.

Neither compliance as enforced rules nor compliance as enforced by market incentives is a reliable means to the best outcomes when detection of ethical lapses is costly or requires special expertise, and this is why professional ethics must start from a place already beyond mere compliance on consequentialist as well as deontological accounts of ethics. Mere compliance fares even worse when it is considered from the vantage of a typical virtue ethics. Virtue ethics regards the evaluation of character as more basic than the evaluation of actions or states of affairs. The central concern is not whether the person does the right thing or brings about a good result but whether he is an excellent person. The conflict between being an excellent person and being merely compliant is perhaps most obvious with respect to the virtue of integrity. It is commonplace that integrity means doing the right thing when no one is looking, and we have already seen that the merely compliant agent has no reason to do the right thing if no one is looking. The person of integrity is not concerned with whether he can get away with it, as integrity involves being steadfast in one's constitutive commitments such that one's life hangs together as the life of a single, coherent thing and not as a mere heap of whim and desire. The person of integrity stands for something and can be counted on to act accordingly as his actions flow from his commitments. Mere compliance to external authority thus tends to run against one's integrity. Indeed, defiance of authority one regards as illegitimate may be the clearest expression of integrity. Mere compliance with respect to small matters (e.g. traffic laws or small favors) need not always constitute a failure of integrity, and compliance to an accepted authority can even be an expression of one's integrity so long as one's commitment to the legitimacy of that authority is a constitutive element of one's character. However, this indicates a sense in which such compliance is no *mere* compliance; the affirmation of the authority's legitimacy must ultimately come from the agent himself and not in mere obedience to any external authority. Whatever authority the person of integrity obeys will ultimately derive its legitimacy over him from his own judgment, as one who does as he is told just because that is what he has been told fails to manifest integrity. He cannot be counted on to have any settled character beyond his unquestioned obedience, and to that extent he lacks any identity of his own. He is not one thing; he is simply whatever some authority tells him to be. He is not an excellent person. Indeed, he is barely a person at all. He is a mere utensil. Thus, an ethics that takes integrity seriously will have to reject an ethics of mere compliance.

3. What we care about

This discussion of integrity starts to bring out more clearly what seems to be the fatal flaw in any effort to advance ethics by way of the tools of mere compliance; any such project will have to profess an indifference to what Harry Frankfurt has called "the importance of what we care about" (Frankfurt, 1988). As we have already seen, ethics as mere compliance is liable to fall short of the full realization of our ethical ideals precisely because compliance means doing something about which *the agent does not really care*. If she did care about doing it, then she would do it because she cares about doing it and not as an act of mere compliance. Moreover, if she cared about doing it, she could be reasonably expected to do it as well as she is able *because* it is important *to her*. If she cared about doing it, she would not need extrinsic rewards to do it well, and her doing it well would not depend on our ability to detect her failures and shortcomings.

It should be noted that Frankfurt considers the importance of what we care about to be a branch of inquiry distinct from ethics on account of his narrowly defining ethics as only concerning our relations to *other people* and especially regarding matters of *right* and *wrong*. This is a grievous error on Frankfurt's part, as our discussion so far already indicates that what we care about is a matter of great ethical import on Kantian and consequentialist grounds alike. The case is even worse when we turn to consider virtue ethics, as narrowly limiting ethics to matters of moral obligation begs a number of important philosophical questions against virtue ethicists and their ilk. To be clear, Frankfurt's point is meant to be neutral as between theoretical frameworks, and should not be mistaken for a call to adopt or reject an "ethics of care" in any of its various forms; to describe some tradition as an "ethics of care" means the standards of right action or good character relate to caring for various moral patients, but interrogating standards of right and wrong or good and bad is not Frankfurt's concern. His concerns have more to do with moral psychology and motivation. He goes too far, however, in sharply distinguishing ethics from the importance of what we care about as utterly distinct fields of philosophical inquiry, especially if he means to be neutral with respect to various theoretical frameworks. Frankfurt's point in distinguishing ethics from caring seems to be that caring about something does not itself constitute the judgment that some action is morally obligatory, and this much is certainly true; in fact, the contents of one's moral beliefs and one's affective ties can and do often come apart in spectacular ways. However, as Frankfurt agrees when he discusses the case of a mother who cares nothing at all for her children, we do criticize and shame one another on the basis of what we care about or fail to care about. We might criticize such a parent even if she was unfailing with respect to the narrow requirements of moral obligation; she seems to have failed at being a good parent on account of her failure to achieve appropriate affective attitudes, even as she does her duty without fail. She may do what is morally obligatory, but she is not an excellent or flourishing person. This is exactly the sort of case that leads virtue ethicists to criticize conceptions of morality that regard affective relations as beyond the pale of ethics, and it is significant that Frankfurt (1988, p. 91) comes to sound much like a virtue ethicist himself:

> Especially with respect to those we love and with respect to our ideals, we are liable to be bound by necessities which have less to do with our adherence to the principles of morality than with integrity or consistency of a more personal kind. These necessities constrain us from betraying the things we care about most and with which, accordingly, we are most closely identified […] what they keep us from violating are not our duties or our obligations but ourselves.

Surely such considerations indicate that ethics is not narrowly about obligation. If we regard ethics as the study of what there is most reason to do or want, instead of being narrowly about obligation, then what we care about is a matter of great ethical import. In any case, what we care about makes all the difference for whether our actions are a matter of *mere compliance* with all the shortcomings that attend that approach. Ethics beyond mere compliance requires agents to care about ethics. This poses a significant practical puzzle about how to do ethics, for it may not be perfectly clear how we can bring our audiences to care about something they do not already care about.

On certain rationalistic conceptions of the human condition, in which matters of the heart are to be contrasted with what issues from the head, coming to care about something can seem an inscrutable and capricious affair; if we imagine our audiences might damn well choose to care about anything they like, that it is no use trying to understand or modulate the mysteries of love and affection, then we have every reason to despair that we cannot get audiences to care about ethics. Fortunately, this is an incorrect conception of how we come to care; as it turns out, affairs of the heart are no more capricious or inscrutable than our faculties for fixing belief. Despite his tragically narrow conception of ethics, Frankfurt's essay contributes a great deal to our understanding of what it would take to advance an ethics beyond mere compliance by articulating the relation between what we care about and our integrity. In light of this relation, what we care about cannot be a matter of whim and caprice. "The notion of caring implies a certain consistency or steadiness of behavior, and this presupposes some degree of persistence" (Frankfurt, 1988, p. 84). Because what is decided in one moment can be undecided or ignored in the next, Frankfurt contends that what we care about cannot be a matter of mere decision or whim. "The fact that a person decides to care about something cannot be tantamount to his caring about it. Nor is it a guarantee that he will care about it" (Frankfurt, 1988, p. 84). That what we care about is not simply a function of what we decide to care about has profound implications for how we conceive our affective faculties, and consequently, for how we do ethics if we agree with Rogerson that successfully doing ethics issues in some tangible, practical difference for behavior that is not a matter of mere compliance:

> The fact that someone cares about a certain thing is constituted by a complex set of cognitive, affective, and volitional dispositions and states [...] It certainly cannot be assumed that what a person cares about is generally under his immediate voluntary control (Frankfurt, 1988, p. 85).

You can choose whether to comply with the ethical commands of some purported authority, but you cannot simply choose whether you care about them. At first blush, the involuntary nature of affective reactions may seem to make them even less amenable to instruction or reflection, but this is exactly backwards. Indeed, this insight reveals above all that our affective reactions resemble perception and belief formation in that they come about as an involuntary response to stimuli. To care about something is to "see" in a rather immediate way that something is important. Far from making such reactions inscrutable and capricious, the parallel with epistemic states and dispositions suggest there are more or less reliable ways to produce such affective reactions. If they are a response to stimuli, then presenting the stimuli to normal, healthy individuals should produce the response. Getting an audience to care about something is not some mysterious conjuring trick; we do it in just the way we get an audience to believe or see for themselves that something is so. We just need to direct their attention to the right stimuli by asking the right questions. Significantly, if our method is to ask questions, we do not even need to know in advance or with much precision what the right stimuli are. Inquiry puts us on the path to encounters that foster sensitivity to such stimuli, whatever they turn out to be. On this score, the parallels between cognitive and affective reactions run deep.

It is generally accepted that belief and perception are no more fixed by mere choice than caring. That is, one does not simply choose what she sees or whether or not she regards some proposition as true; one is either convinced of it or not. Furthermore, as is also the case with caring about something, actually believing something cannot be a matter of compulsion or mere compliance. I can make you comply with my beliefs on the exam, checking off all the "right" answers to get a good grade, but I cannot make you actually believe it any more than I can make you actually care about something by insisting on mere compliance. It is an empirical question whether and to what extent one can cultivate belief or caring by compelling subjects to declare the prescribed attitudes over and over despite their having no such attitudes, but such a practice is at least ethically dubious, as it frankly means to indoctrinate and subvert the independent thinking of the subject while it pressures her to behave dishonestly, all the while risking backlash and resentment against the intended outcome. Fortunately, there is a more honest and still fairly effective way to convince an audience: appeal to the sorts of evidence they are already prepared to accept. Invite them to inquire with you into the evidence framed in terms of shared premises, and see where the inquiry leads. In the best case, the audience will come to see things as you see them through their own eyes (often figuratively, but sometimes literally, as in a laboratory demonstration), convinced as you were convinced on account of the evidence being so convincing. Of course, there is no guarantee that such an inquiry will convince every member of every audience. It may happen that all sorts of background beliefs color or distort what appear at first to be shared premises, and it may happen that shared premises will be hard to find after all, but if you are especially honest and open about all this, it promises to be a learning experience for everyone involved. This is, after all, more or less the method of science, and if you are ever distraught by the discovery that not everyone you mean to instruct sees things quite the way you see them, recall to yourself that "One repays a teacher badly if one always remains nothing but a pupil" (Nietzsche, 1883, p. 190).

The purposes of inquiry are not the purposes of drill or other closely guided instruction, and what inquiry means to teach is notoriously hard to measure. Inquiry is explicitly modeled on the practice of working scholars and scientists and is meant to improve students' depth of understanding, critical thinking and ability to find meaning rather than merely expanding the sum total of their useless "knowledge" (promptly and safely forgotten after the exam, for such are the wages of mere compliance). In short, the practice of inquiry is intended to foster the sort of independent thinking that discovers new knowledge rather than parroting what is already known. It is intended to cultivate a scientific sensitivity that we have every reason to suppose parallels the ethical sensitivity we aspire to cultivate when we do ethics. Whatever its merits in the classroom (it is not the only or best way to approach every purpose, after all), open-ended inquiry is what scholars do when they explore the questions that matter to them in their research, and it is ideally what scholars do when they share their results with the wider scholarly community in journals and at conferences. Open-ended inquiry is what we do at Ethicomp, inviting all to partake in the conversation without regard to adopting any particular creed or secret handshake. We start by asking questions, and we explore our various proposed answers together. If consensus emerges, we suspect it results from our having come into contact with some or other facts which provoke a like response in us. Where consensus is not forthcoming, we keep looking, eager to cultivate our sensitivity to whatever facts there are in the reflective process of looking for whatever facts there might be. This is what we do in our various conversations regarding the ethical implications of computing, whether in the classroom, with friends, family and colleagues or at conferences like Ethicomp. This is, in short, what I take computer ethics itself to be. Computer ethics is an open invitation to reflect on the ethical significance of computing as a practice, a profession and an institution. When we invite an audience to think through ethical issues with us, they are put in a position to discover for themselves that ethics matters, and this discovery encourages (but does not guarantee) ethical conduct that is not mere compliance. The hypothesis is that engaging in a sustained and earnest discussion of ethical issues is likely to foster a greater sensitivity toward the sorts of concerns in question, and there is empirical evidence to support it (Jagger and Volkman, 2014). This is the best way I know to cultivate genuine ethical sensitivity without encouraging mere compliance, and our confidence that it can work should correspond to our confidence that ethics really does matter, that there really are facts of this sort that provoke responses in creatures like us. Our audiences do not need to take our word for it when they can see it with their own eyes. When exposed to ethical issues, they will discover for themselves what matters after all.

4. What is computer ethics?

It must be emphasized that there is no necessary competition between a conception of computer ethics as an invitation to inquiry and reflection versus other accounts of the nature and scope of computer ethics, but it does advise these other accounts to acknowledge that each will have particular uses before particular audiences, and it behooves the ethicist to deploy the right framework for the right audience and the right occasion. An approach oriented toward inquiry over drill or direct instruction is not utterly unstructured or rootless; on the contrary, various theoretical perspectives are essential in conceptualizing and structuring discussion, provoking questions as well as proposing answers. Judiciously introduced and applied, theory enriches inquiry and discussion. However, the misuse of theory can effectively shut down inquiry, replacing it with mere drill and rote. In that case, the audience is no longer engaging their own rational and affective faculties and seeing for themselves what matters. They are being asked to comply with some cold theory.

Such a misuse of theory threatens to dominate in certain computer ethics classrooms, especially when these classes are led by computing faculty with their own keen (and professionally appropriate) interest in getting at the right answers rather than appreciating the wonder and aporia that have been the hallmark of doing philosophy since Socrates. Connolly and Fedoruk (2014) observe that when computer ethics is taught by computing faculty:

> […] 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.

The irony in all this is that, although the computing faculty mean to be applying the best philosophical theories, practically no philosopher would dream of teaching a course in applied ethics in such a ham-fisted way. Computer ethicists must do a better job of reaching a wider audience of computing faculty. We must bring them along to understand that, for all the reasons indicated above in defense of inquiry based learning, philosophical theories, especially in ethics, are not appropriately applied the way an engineer would apply a formula learned from physics. When we misapply philosophical theories as if they were established formulas or laws, we are not doing philosophy, and we certainly are not inviting students to do philosophy. We are not inviting them to see for themselves that ethics matters.

Having a particular theory or account of computer ethics can structure and promote inquiry instead of stifling it, but only if we remind ourselves that theories bring certain matters into focus while blurring other matters out of focus. In a policy discussion among computing students, it makes good sense to reflect on Moor's (1985) argument that the computer revolution "transforms social institutions" and that these transformations are apt to produce "conceptual muddles" leading to "policy vacuums" that require philosophy and conceptual analysis to resolve. This line of thought helps the audience to see why philosophy is relevant to solving certain policy problems. On the other hand, using Moor's formulation of computer ethics to constrain or exclude from consideration whatever is not a matter of policy shuts down worthwhile inquiry, and this is not a good use of theory (nor is it ever Moor's intention to shut down inquiry). For example, although online lurking in social media raises significant ethical questions that cannot be reduced to any question of policy, law or moral obligation, Moor would surely agree that it is worth considering the implications of lurking for character and flourishing (Volkman, 2011). Similarly, Gotterbarn's conception of computer ethics as a field of professional ethics can be illuminating for an audience of future computer professionals, but it is less useful for settling policy matters or before an audience of end-users or managers. Surely, Gotterbarn would heartily agree that computer ethics belongs in the contexts of Business Ethics, Medical Ethics, Social and Political Issues and general introductory courses in Moral Problems as well as in courses designed especially for future computer professionals, even as his work in computer ethics focuses on computing as a profession. It should be clear that all this is no criticism of Moor or Gotterbarn; on the contrary, I am explicitly endorsing each approach in its appropriate domain, and I introduce both in the classroom to structure and facilitate discussion. The point is simply that such approaches are best understood as aspects of the wider process of inquiry that I am advocating as a "big tent" conception of computer ethics.

Nowhere is it more apparent that the usefulness of a theory is relative to the audience before which it is used than in the case of such especially abstract and general theories as Floridi's (2013) *Information Ethics* and Bynum's (2006) *General Theory of Flourishing Ethics*. The proposals that information is *per se* intrinsically valuable or that entropy is *per se* a great natural evil are provocative philosophical theses, and these theories deserve to be examined alongside others subject to our inquiry, especially before a philosophical audience. However, they are not really designed or suited to moving an audience of end-users or managers or computer professionals, at least not directly. In the big tent view that computer ethics is an invitation to inquiry, such proposals are welcome a place at the table (of course there is a table in the tent – otherwise, where would we put our drinks?), but it will not be their place to move certain audiences to greater genuine ethical sensitivity, as they have to be articulated and defended in terms of a great many abstract premises and background understandings that do not already have a hold on the relevant audiences. For a philosopher casting about for a theory, in *that* domain of inquiry, such proposals are live options to be evaluated and considered, but they are liable to leave a computer professional or end-user in the audience cold. Inquiry must begin where the audience already is, and many of the audiences of computer ethics are sure to find Information Ethics or General Flourishing Ethics alien to the point of unintelligibility. This does not mean there are no insights to be had from entertaining these approaches in our investigations of policy questions or professional ethics or the aporia of lurking, but it does mean, as I have argued in much greater detail elsewhere (Volkman, 2011), that such theories are of limited applicability to Rogerson's (and my own) declared project for computer ethics that it should constitute a real engagement with industry and society at large that cultivates a genuine sensitivity to ethical concerns in the creation, development and implementation of technologies.

Floridi (2013, p. 112) has expressed a certain impatience with at least part of this wider project by declaring:

> The arguments offered in this book do not provide threatening answers regarding the consequences of an action, as one might to the teenager's question "why should I do it?", but seek to answer the more mature and difficult question, "what would be the right thing to do?"

Perhaps the philosophical project of devising theories of ethics can (at least in principle) come apart from considering the importance of what we care about in pursuing the practical project of cultivating genuine ethical sensitivity, although such a claim requires its own philosophical defense. Earlier in the same work, Floridi (2013, p. 68 emphasis original) makes the point more moderately, essentially conceding that the project of Information Ethics is a philosophical project that might be distinguished from other sorts of ethical investigations:

> When the primary aim of an ethical investigation is to understand what is right and wrong and, hence, what the best course of action would be, irrespective of a specific agent's behavior and motivation, it becomes possible to adopt a more *philosophical* viewpoint.

Such speculations about the ultimate nature of right and wrong certainly have their place, especially (but not only) in the philosophy classroom, but if we agree with Rogerson that ethics is something we do with a practical purpose in mind and not just as an idle academic enterprise, we should be suspicious of theories that purport to do ethics without concern for the problems of moral motivation. If the problem of moral motivation is never addressed, then ethics (at least as described by such theories) may very well seem as nothing but an idle academic enterprise to those audiences that do not already endorse the theory in question. Indeed, there is reason for suspicion of such theories on philosophical terms alone; it is hardly obvious that there is any such thing as the "best course of action […] irrespective of a specific agent's behavior and motivation". Williams (2001) offers compelling reasons to think there is something quite amiss in theories that purport to justify reasons for action external to the motivational set of the agent whose reasons they are purported to be, and this has led to a vast literature on the relation of moral motivation to moral theory. As Korsgaard (2001, p. 120) observes, "The force of the internalism requirement is psychological: what it does is not to refute ethical theories, but to make a psychological demand on them". As Churchland (2012) contends, whether a given theory can meet this psychological demand becomes both a more tractable and a more pressing question the more we know about how human brains and behaviors have evolved and operate, and theories do not all satisfy the psychological demand with equal plausibility or in equal consonance with the facts of psychology, biology and neuroscience. However all this plays out in the realm of evaluating philosophical theories (much controversy remains here), to accomplish the practical goals of computer ethics it is clear enough that we must not only make room in the tent to accommodate the internalist psychological demand but we must also make every effort to meet the internalism requirement directly by framing our invitation to inquiry with special sensitivity to our various audiences.

Other approaches to computer ethics also aspire to erecting a big tent, and the view defended here enjoys great affinities with some species of what Bynum (2006) calls *Human-Centered Flourishing Ethics*, although I would hesitate to describe my characterization of computer ethics as anything like a *theory*. I think it is more apt instead to see this conception of computer ethics as articulating a particular role for the computer ethicist. Whatever the audience we mean to reach, our success will depend on approaching the audience in the spirit of inquiry and as facilitators of discussion, helping them to think through ethical issues for themselves. This means the domain of computer ethics will have to include just about anything that can reach some audience where they already are. Such a broad conception of computer ethics goes a long way to mitigate the worry that, for example, science fiction cases might become a source of distraction for us. Anything that sets the stage for inquiry is a welcome part of the project, and certain works of science fiction are especially well suited to sparking the imaginations of certain audiences. As Gerdes (2014) contends, "the expressive power of art provides us with ways to understand ourselves and our being and actions in the world". It is our role as ethicists to poke and prod our various audiences to such understanding, and works of art and inspiration have as much if not more claim to this function as any theory. In any case, we cannot hope to specify in advance the bounds of inquiry or its many possible starting points. So long as we are fulfilling our role as educators and hosts of further inquiry, we serve the practical project of computer ethics.

The ACM/IEEE-CS Joint Task Force (2014) Software Engineering Code of Ethics and Professional Practice explicitly embodies this role when it declares:

> The Code is not a simple ethical algorithm that generates ethical decisions […] Ethical tensions can best be addressed by thoughtful consideration of fundamental principles, rather than blind reliance on detailed regulations […] As this Code expresses the consensus of the profession on ethical issues, it is a means to educate both the public and aspiring professionals about the ethical obligations of all software engineers.

No one could understand this and come away thinking that living up to the code is about mere compliance.

The role of the computer ethicist, in my view, is not simply to legislate or testify on behalf this or that vision of the true, the beautiful and the good, but to serve as a mentor and a gadfly, bringing others to see for themselves what matters and why. That this is at the heart of what the Ethicomp series has always been about is captured in Flick's (2014) thesis that:

> ETHICOMP is a "community" mentor – a group of people who nebulously are able to offer temporary (individually) but on-going support that could be considered mentor-like throughout a career progression. The atmosphere and nomenclature within the conference of the community as "family" fosters this embracing nature of relationships formed there between academics of all levels, and those who "come through the ranks" are usually happy to contribute when they are able to.

I would hasten to add that this mentoring relation extends not only to academics of all levels but to our various other audiences as well. We mentor professionals, aspiring professionals, managers, end-users, policymakers and each other by offering an open invitation to inquiry about the broad ethical implications of the practices, professions and institutions of computing and the design, development and implementation of computer artifacts. In this manner, we cultivate a genuine sensitivity to ethics beyond mere compliance. In my view, that is something worthy of celebration.

References

ACM/IEEE-CS Joint Task Force (2014), "Software engineering code of ethics and professional practice", available at: www.acm.org/about/se-code (accessed 29 July 2014).

Bynum, T. (2006), "Flourishing ethics", *Ethics and Information Technology*, Vol. 8 No. 4, pp. 157-173.

Churchland, P. (2012), *Braintrust: What Neuroscience Tells us about Morality*, Princeton University Press, Princeton, NJ.

Connolly, R. and Fedoruk, A. (2014), "Why computing needs to go beyond good and evil impacts", *Proceedings of Ethicomp 2014*, *Paris*.

Flick, C. (2014), "What's Stopping Me? The importance of Ethicomp as a "community mentor" for doctoral and early career researchers", *Proceedings of Ethicomp 2014*, *Paris*.

Floridi, L. (2013), *The Ethics of Information*, Oxford University Press, Oxford.

Frankfurt, H. (1988), *The Importance of What We Care About*, Cambridge University Press, Cambridge.

Gerdes, A. (2014), "IT-ethical issues in sci-fi film within the lifespan of Ethicomp", *Proceedings of Ethicomp 2014*, *Paris*.

Gotterbarn, D. (1997), "Software engineering: a new professionalism", in *The Responsible Software Engineer*, Springer, London, pp. 21-31.

Jagger, S. and Volkman, R. (2014), "Helping students to see for themselves that ethics matters", *The International Journal of Management Education*, Vol. 12 No. 2, pp. 177-185.

Korsgaard, K. (2001), "Skepticism about practical reason", *Varieties of Practical Reasoning*, Massachusetts Institute of Technology Press, Cambridge, MA, pp. 103-126.

Ladd, J. (1980), "The quest for a code of professional ethics: an intellectual and moral confusion", In *Ethical Issues in the Use of Computers*, Johnson, D. and Snapper J. (1985), Wadsworth, Belmont, CA, 8-13.

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

Nietzsche, F. (1883), "Thus spoke Zarathustra", in Kaufmann, W. (1982), *The Portable Nietzsche*, Penguin Books, New York, NY.

Parfit, D. (1984), *Reasons and Persons*, Clarendon Press, Oxford.

Railton, P. (1988), "Alienation, Consequentialism, and the demands of morality", in Scheffler, S. (Ed.), *Consequentialism and Its Critics*, Oxford University Press, pp. 93-133.

Rogerson, S. (2014), "Ethicomp 2014 keynote address", Public Speech, Ethicomp 2014, 25 June, Paris.

Volkman, R. (2011), "Lurking: finding one's self while remaining hidden", *Proceedings of Ethicomp*, *Sheffield*.

Williams, B. (2001), "Internal and external reasons, with postscript", *Varieties of Practical Reasoning*, Massachusetts Institute of Technology Press, pp. 77-98.

Corresponding author
Richard Volkman can be contacted at: volkmanr1@southernct.edu

For instructions on how to order reprints of this article, please visit our website:
www.emeraldgrouppublishing.com/licensing/reprints.htm
Or contact us for further details: **permissions@emeraldinsight.com**