
Dynamic Technology Challenges Static Codes of Ethics

A Case Study

Brinkman, Bo

University of Miami, Computer Science and Software Engineering.

Flick, Catherine

De Montfort University, Faculty of Technology.

Gotterbarn, Don

East Tennessee State University, Department of Computer and Information Science.

Miller, Keith

University of Missouri, Computer Science.

Vazansky, Kate

Bemidji State University, Department of Mathematics and Computer Science.

Wolf, Marty J.

Bemidji State University, Department of Mathematics and Computer Science.

Corresponding Author: Don Gotterbarn, don@gotterbarn.com

Abstract

We describe the process of changing and the changes being suggested for the ACM Code of Ethics and Professional Conduct. In addition to addressing the technical and ethical basis for the proposed changes, we identify suggestions that commenters made in response to the first draft. We invite feedback on the proposed changes and on the suggestions that commenters made.

Keywords: ACM, Code of Ethics, Professional Conduct

Introduction

A hallmark of the best professional organizations is their commitment to serving the community in an ethical fashion, an understanding of which is frequently expressed in a code of ethics. Codes describe how members of a particular professional organization ought to relate to society. Sometimes codes have an almost contractual nature, stating the professional's obligations to society and to the profession, where failing in either of these may result in loss of membership in the profession. Codes of professional societies can also place an obligation on that organization to promote and follow the standards of the code.

As society and the profession change, the code describing that relationship also needs to change.

1972

The Association for Computing Machinery (ACM) has developed several statements of ethical guidance. Its first code was the Code of Conduct adopted in 1972. It was organized as 4 professional ideals (Ethical Considerations) with mandatory rules (Disciplinary Rules) and was applicable only to those ACM members with voting rights. This Code of Conduct was published relatively early in the development of computing; the programming language ‘C’ and the microprocessor had just been invented.

1992

Twenty years later, after the development of parallel computing, the PC, and 20 Meg hard drives, the ACM significantly revised its 1972 Code of Conduct. In 1992, the ACM developed a Code of Ethics and Professional Conduct¹. Like the 1972 Code, the 1992 version was based on a set of fundamental ethical considerations found in many professional codes of ethics, including don’t cause harm, maintain integrity, competence, self-improvement, and impartiality. Both computing and our understanding of its role changed between 1972 and 1992. In addition to making some statements about conduct, the primary function of the 1992 Code was to aid practitioners in decision making as new situations arose. This code was stated in terms of aspirations rather than the fixed behaviors of the 1972 Code. Rapid technology change could quickly make rigid rules irrelevant. The 1992 Code tried to reflect the conscience of a broader computing profession; it was designed for every member of the ACM, not just voting members; and it established a model for the computing profession as a whole.

2018 Draft 1

In spite of the rapid changes in computing technology, it has taken many years for the ACM to revisit and update its 1992 Code of Ethics. The ACM Committee on Professional Ethics (COPE) is leading the “Code 2018” project to update the Code. The core ethical principles have largely remained constant, but the dramatic expansion of the role of technology in society requires a careful rethinking of how these principles are manifested. The updated Code will have essentially the same basic structure: a preamble, and sections on fundamental ethical principles, professional obligations, leadership obligations, and compliance.

Challenges

There are recognized challenges in writing and updating any code of ethics.² A rapidly changing society means that no code can ever be complete, so a code should contain suggestions to address new situations. It is a mistake to assume that the absence of a behavior from the code is an approval of that behavior. Codes must be carefully worded to provide enough detail for guidance and not so much detail that the code cannot speak to new situations. Another difficulty is that in some situations, a code’s imperatives may seem to conflict. Codes need to provide some help in making decisions in these situations.

Codes include statements about a member’s obligation to follow the code, and mention consequences for failing to follow the code. An organization needs clear procedures for enforcing its code, as well as articulating and distinguishing degrees of code violations and appropriate consequences. These due process documents are not part of the code, but are referenced by the code.

A code cannot list every sort of harm it rejects, so there need to be supporting documents, examples and case studies that help members understand the breadth of the word “harm” and help answer questions like “Does the code rule out ‘harm’ to criminals or terrorists?” “Can I work on military weapon software?” “How do I keep user data from falling into the wrong hands?” The 1992 Code talks about harm to data and files but does not address emotional and reputational harms that are now possible through social media. In an ideal situation, a professional organization would provide someone to help members to address ethical dilemmas and difficult situations where they may be required to prioritize stakeholders.

When an organization has a global membership (as the ACM does), there are additional difficulties. The code cannot appeal to a particular national document or position for its authority, and care must be taken to ensure that the code and its associated support documents are clear across different countries and cultures.

Some of the challenges in updating a code are prosaic, but still daunting. For example, people are sometimes embroiled in controversies about “ensure” versus “assure;” when to use “shall,” “should,” “must,” or “may,” and when to use an Oxford comma.

The members of Code 2018 Task Force considered comments from a broad spectrum of computing practitioners resulting in Draft 1 of the suggested updates, as described in the December 2016 Communications of the ACM³. The changes to the Code were sometimes isolated in a single section and sometimes a change to one section led to changes in other sections.

Discussing proposed changes to the Code in a public forum such as the CEPE/ETHICOMP Conference helps the Task Force to meet these challenges to updating the Code and helps to prepare for the next draft of the Code. This paper examines suggested changes to the code, to promote a thorough evaluation of proposed changes to Draft 1, and to educate ACM members about issues in building a computing code of ethics. We review the technical and ethical reasons for the changes to each section of the Code, considering suggested updates, discussing and evaluating the technical and ethical motivations behind the changes. The paper concludes by reviewing some of the suggestions that impact the entire Code. We hope this paper encourages further suggestions about the development of an updated ACM Code and promotes discussion about the ethical issues facing computing. You can use the comment form at <https://ethics.acm.org/code-of-ethics/code-2018/> to give us your comments on the drafts of the Code.

The Preamble

The 1992 Code was written as a collection of imperatives, each with a set of guidelines to provide assistance in applying the imperative. The drafters envisioned that the imperatives would be stable over time and that the guidelines would change “more frequently than the Code.” This turned out to be not the case. Draft 1 of the updated Preamble reflects this reality. Updating the Code is a significant undertaking and both the imperatives and the guidelines are intimately tied to one another. The guidelines “assist members in understanding and applying the imperatives” (Draft 1).

As the community reviewed Draft 1 of the update, they suggested four areas of concern surrounding the Preamble. One concern is the range of application of the Code. Should it apply only to ACM members or should it apply to all those involved in the computing profession? A second concern is the transparency of the ethical decision-making process. The updated Preamble in Draft 1 introduces “the public good [as] a primary consideration” in the ethical decision-making process. This led a commenter to question whether the Preamble ought to spell out what “the public good” entails, our third concern. Another consideration

is whether, and if so how, the public good is changing because of the pervasiveness of computing technology. The fourth concern is how the Preamble, and the Code in general, might deal with the professional conduct of members in areas that do not directly involve computing technology, such as those involved in establishing conference paper review and acceptance policies for ACM conferences. (This issue is also taken up in addressing changes to Section 3 of the Code.) The remainder of this section offers some ideas on each of these four topics from several different perspectives.

Universality

changes in Draft 1 were attempts to more clearly identify whether a statement applied to all computing professionals, only to ACM members, or only to ACM members who are acting in a computing professional or leadership capacity. Yet, based on the comments we received on Draft 1, this distinction is still not always clear. Adding language to the Preamble that indicates that where the terms are used, they were chosen intentionally to identify to which group the statement applies would mitigate some concerns. Addressing this confusion, however, does not address two broader questions: Even if the ACM were based in a single country, would it be appropriate for the ACM to establish a code of ethics for an entire profession? In the case that it is appropriate for a single country, can the relevant professional and ethical standards be used to articulate a code for an organization whose members are scattered throughout the world?

Inasmuch as codes of ethics are aspirational, the broader interpretations would seem appropriate on both counts. It seems reasonable for the ACM to have this aspiration for all computing professionals. While there may be parts of the Code that do not apply to non-ACM members, one still might argue that every computing professional ought to follow: “**Imperative 4.1** Uphold, promote, and respect the principles of this Code.” Any sort of enforcement actions might only come to bear on those computing professionals who are or are attempting to become ACM members. The argument might continue that the ACM bears some responsibility of spreading the word about this responsibility to all computing professionals. At the very least, this would entail making the Code available in many languages and establishing mechanisms by which computing professionals are educated about the Code and about applying the Code.

On the one hand, the ACM is a private organization and its members choose to become members at non-trivial cost. Not all computing professionals are in a position to make that investment. A further step the ACM might take in the name of the profession is to call to account computing professionals who are not members for violations of the code. Certainly, some would argue that this such an action would be overreach by the ACM. On the other hand, maintaining such a public list of offenders of the Code might also be viewed as a service to the computing profession and, more broadly, to the whole of society.

A case for applying the Code to non-members might be easier to make for those who organize, submit papers to, or even attend, ACM sponsored conferences. Here the conference organization team, including all members of the program committee, can be said to be representing the ACM, even if they are not ACM members. The ACM could take steps to make conference leaders aware of the Code, aware of the expectation that they adhere to the Code, and of their responsibility to ensure that all those with leadership positions for the conference are aware of those expectations.

Transparency

Often the evaluation of ethical decision-making centers on knowing the process that was used to make the decision. Attitudes such as “all’s well that ends well” do not fit in such an evaluation. Transparency of the ethical decision-making process includes identifying the stakeholders considered and the cultural context

used, and documenting the timing and nature of the process (was it a one-off evaluation or was there repeatedly more thorough evaluation as the project was taking shape?). This leads to the question, “Should the Code address transparency in the ethical decision-making process?”

This question can be addressed at the two levels mentioned in the Code: the individual level (computing professionals and managers) and the organizational level. At the individual level there is the question of transparency between management and staff. It seems reasonable for the Code to address this. On the other hand, when the transparency is between a provider of a system and its users, should there be different expectations for individual professionals and for organizations? Does the size of an organization matter? For example, is it reasonable for the Code to suggest that large organizations have in place a transparent ethical decision making process? Doing so invites questions of overreach again.

However, one might argue that such processes engender trust between the public and the computing profession, surely a worthy goal.

The Public Good

One of the most significant additions to the Draft 1 Preamble is language that the public good is a primary consideration when using the Code to deliberate a particular decision. A Code needs to provide guidance for situations when its principles appear to conflict. This problem is generally addressed in a paramountcy clause used to address any potential inconsistency or tension between code principles. The paramount consideration is the public good. An indication of the Code’s support for this principle is already in the structure of the Code, whose first principle is “Contribute to society and human well- being.” In many situations, asking which possible action ends in supporting this principle more resolves the conflict.

This also raised the question of whether the Code ought to clarify the meaning of “the public good.” A commenter suggested defining the public good as “inclusivity/affording equal opportunities to all, upholding civil rights and liberties of individuals, treating individuals with respect and upholding their right to autonomy.” Including such a definition raises a number of concerns. First, is defining the public good in the purview of this code of ethics? Doing so makes it clearer what is meant by the term within the bounds of the document. Those who are using the document to guide their analysis of ethical decisions will have a clearer understanding of what the term means to the Code’s authors. But including such a definition suggests that the ACM is claiming some special expertise in understanding and identifying the public good. Similarly concerning is that by defining the term, the meaning of the public good becomes fixed within the Code, perhaps limiting its future applicability. Other professions, medicine and engineering, address this “term flexibility” in external regularly reviewed documents which provide a model for evaluating “public good” in different socio-technical contexts related to multiple professional contexts.

“Imperative 3.7 Recognize when computer systems are becoming integrated into the infrastructure of society, and adopt an appropriate standard of care for those systems” was added to Draft 1 to highlight the significant way that technology has been integrated into the infrastructure of society. Such integration has led to debate about such issues as “digital rights” and the “digital divide.” These clearly need to be part of the conversation when one is discussing the meaning of the public good. Moving forward, issues such as “machine rights” are also likely to have an impact. These changes, while not rapid, suggest that the notion of the public good may be changing over time in response to technology and how people integrate it into their lives. Ideally, the Code should be sufficiently flexible to accommodate changes in the meaning of the public good, yet solid enough to protect any widely shared notion of the public good. One strategy to accomplish this is to include a definition of the public good that captures a minimal shared notion and have the Code leave open the possibility of expanding the definition. Another is to leave the term

undefined and rely on other aspects of the Code that remind the user that the Code is a tool for *guidance* in ethical decision-making, and it should not be treated as a document that is complete in its capturing of the complexities of ethical decision making.

The Code and the Meaning of Professional Work

The Preamble states that “The Code is intended to serve as a basis for ethical decision making in the conduct of professional work.” It then goes on to focus on “one’s conduct as a computing professional.” There is a lack of guidance on the application of the Code to members (and non-members) whose professional activities are outside of the computing profession and those who are in charge of members of the computing profession. In addition to those involved in ACM’s conferences, two large subsets of members come to mind, those whose professional responsibilities include education and those with professional responsibilities in computing research. The general moral imperatives clearly apply to these members and, arguably, to any professional. The imperatives in Section 2 of the Code are almost all written in such a way that they apply to any professional. Yet, much of the guidance offered in the section (in addition to the imperatives themselves) focuses more specifically on the *computing* professional. Should this guidance be expanded to include members whose responsibilities may not be computing per se?

Section 3 of the Code turns to organizational leaders. “**Imperative 3.4** Ensure that the public good is a central concern during all professional computing work” by its very wording does not seem to directly apply to educators, scholars and conference organizers who are members of the ACM. Including them might lead to positive changes, such as conference paper acceptance policies that encourage analysis of the public good as part of every submission.

Wording such as that found in **Imperative 3.4** may suggest that the Code does not apply uniformly to all members. Even so, members who are not computing professionals may benefit from this imperative and any guidance it offers. Alternatively, one might argue that placing the public good as a central concern comes to members who are educators and scholars via a different route. Educators at public institutions, for example, are funded by taxpayers and are necessarily beholden to the public good. Later in the paper, **Imperative 3.4** is taken up for different reasons. Here, it serves as means to ask the question of how the Code might apply to those professional roles that surround computing but are not actually computing.

Section 1 Fundamental Ethical Principles (Also Known as General Moral Imperatives)

The first section of Draft 1 contains 7 principles, down from 8 in the 1992 Code. The 1992 Code contained separate imperatives on intellectual property and proper attribution that had significant overlap, and these were combined.

An overarching theme in the comments on Section 1 was discomfort with the idea that the Code is not an algorithm for decision-making, but a set of principles, and that ACM members are required to do thoughtful ethical analysis. The guidance often contains lists of examples of good or bad behaviors, and many commenters took these to be definitions, not examples. As with the “public good” in the preamble, commenters often wanted definitions of key terms: harm, unfair, system, conflict of interest, and proper authorities. They also often wanted clear lines drawn. What is a good enough attempt? What is sufficient consideration? These issues are intentionally left open in the Code, because the facts of each case, the

context of the situation, and the weighing of tradeoffs all matter. We cannot (and should not try to) write a code that is a collection of rules that can be applied without critical thought. Rather, Section 1 is a statement of overarching fundamental moral principles of the computing profession. They remind us of factors that must be given weight and considered in our deliberations, but they cannot (and should not be mistaken for) “get out of jail free cards” for avoiding ethical dilemmas. Sadly, it is just not that easy.

One overarching suggestion for Section 1 was to retitle it from “General Moral Imperatives” to “Fundamental Ethical Principles.” The argument is that the items in this section function better as principles than as rules. In “Profession, Code and Ethics” (p. 40) Michael Davis (Davis, 2002) defines *principles* as “requir[ing] specific factors to be given weight in deliberations,” *rules* as “requir[ing] or forbid[ding] a class of acts,” and *ideals* as “set[ting] a target good to achieve but not necessarily wrong not to achieve.” Much of the confusion about Section 1 might be cleared up if it were understood to be a collection of principles rather than a collection of rules.

In Draft 1 the **Principle 1.1** language was revised from “Contribute to society and to human well-being” to “Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing and its artifacts.” This reflects the view of the Task Force and several commentators that computing professionals are often tempted to omit many stakeholders from their analysis. The phrase “computing and its artifacts” is meant to remind practitioners that it is not just the programming statements that they write that matter, but also those things that emerge from those statements. In particular, the Task Force was trying to address growing concerns about algorithms that emerge from machine learning rather than directly from algorithm designers.

The guidance for **Principle 1.1** has also been updated to emphasize that one cannot satisfy it without paying appropriate attention to accessibility, broadly construed. Public comment identified a possible problem with the wording of the guidance, which currently says “When designing or implementing systems, computing professionals must attempt to ensure that the products of their efforts will be used in socially responsible ways, will meet social needs...” This may be too high a bar: Would it be wrong to release a photo editing program that could be used to forge documents? Most likely not, but the wording of the guidance might appear to support such a claim. Understanding this as a principle rather than a rule reduces some of the tension.

“**Principle 1.2** Avoid harm to others” is unchanged from the 1992 Code, but the guidance has several updates. The 1992 Code encourages us to think about users, the general public, employees, and employers. Draft 1 adds the phrase “and any other stakeholders.” While it is true that “the general public” might technically contain all other stakeholders, the additional emphasis is useful. Public comment on Draft 1 identified a potential weakness of the guidance for **Principle 1.2**. It is quite long, and it currently tries to define “harm” and gives many examples of things that can cause harm, ways to mitigate harm, ways to avoid causing harm, and so on. Some commenters wanted this section expanded even further, with careful and legalistic definitions of all terms. Others wanted cuts because the wordiness of the guidance distracts from the fact that this is meant to be a moral principle, not a rule.

Multiple commenters interpreted **Principle 1.2** as a blanket ban on various kinds of potentially harmful activity: research on brain-machine interfaces, weapons development, or even law enforcement (because, of course, it tends to harm the criminal being punished). This interpretation is not the proper application of the principle because it fails to take into account all stakeholders and to weigh the potential harm of taking a particular action against the benefits of taking that action, or against the harms and benefits of *not* taking the action. Putting murderers in jail does not violate **Principle 1.2** because while it does harm murderers, there would be a worse harm to society in letting murderers go free.

“**Principle 1.3** Be honest and trustworthy” is also unchanged from the 1992 Code, but the guidance has been updated to caution against fabrication and falsification of data. There was very little comment about this principle.

Principle 1.4 has been modified, and with some controversy. In the 1992 Code it reads “Be fair and take action not to discriminate.” In Draft 1 it reads “Be fair and take action not to discriminate *unfairly*.” At issue, here is how to resolve the problem of the word *discriminate* having two different connotations. The first is simply to differentiate between two things, with no negative connotation. The second (and more popular) is an unjust differentiation between two things. It would generally not be wrong to hire a skilled programmer instead of someone that is not a programmer at all for a programming job. This is a discrimination under the first definition, but it does not seem unjust or unfair. Some commenters fear that adding the word “unfairly” to **Principle 1.4** will lead to people thinking that some forms of unfair prejudice are okay.

The guidance for **Principle 1.4** has a greatly expanded list of types of impermissible prejudice. Some commenters reacted very favorably, while others were very negative about this expansion. Some commenters suggested that the list should consist of only legally protected groups. Others argued that the list ought to go beyond what the law requires because it would not be necessary to have the guidance otherwise. Some worried that the list was not expansive enough and that it should be made as complete as possible. Several commenters suggested that sexual harassment needed to be definitively addressed in the guidance for **Principle 1.4**, while others were passionately against doing so, and passionately against any formal recognition of the existence of diverse gender identities or sexual preferences that might be the basis for discrimination or harassment.

“**Principle 1.5** Honor property rights including copyrights and patent” and “**Principle 1.6** Give proper credit for intellectual property” from the Code have been merged into a single new **Principle 1.5**, which says “Honor intellectual property rights and give proper credit.” There was too much overlap of concern between the two principles and that separating them made them more about laws than about ethics. Combining them allows for the rewrite of the guidance to focus more on the ethical concerns. The guidance was also enhanced with support for fair uses of copyrighted material and for the open source community.

Several commenters felt that the guidance language in Draft 1 mixed up plagiarism with copyright. Others felt that the language is still too much about US laws and not enough about the core ethical concerns.

Many in our community are vehemently against copyright and patent, and several commenters were deeply critical of the inclusion of intellectual property rights (IPR) in the Code, feeling that it is a legal concern (not an ethical one), or that copyright and patent are inherently unethical in themselves, or that the current IPR regime in the US and Europe is so broken that we should not support it in any way. This is one of the few legal issues that the Code has taken a stand on, and in part it is likely because there is so much controversy about it within the computing profession.

The last two principles in section 1 are “**Principle 1.6** Respect privacy” and “**Principle 1.7** Honor confidentiality.” The confidentiality principle is unchanged from the 1992 Code, and the privacy principle shortened from “Respect the privacy of others.” A paragraph that was about confidentiality was moved from the guidance of **Principle 1.6** to **Principle 1.7**. The privacy guidance has also been augmented with a statement about aggregation. Several commenters pointed out that “organizations” do not really have “privacy,” and suggested additional clarifications for the guidance of the two principles. Others requested specific mention of the Fair Information Practice Principles, a United States Federal Trade Commission

Standard. Several commenters pointed out the ambiguity in the guidance for **Principle 1.7** that says:

“User data observed during the normal duties of system operation and maintenance must be treated with strict confidentiality, except in cases where it is evidence for the violation of law, organizational regulations, or this Code. In these cases, the nature or contents of that information must be disclosed only to appropriate authorities.”

Some wanted the Code to define who the appropriate authorities are. Some interpreted this guidance to mean that there was a duty to report any and all suspicions of lawbreaking to the police. The intention was to say that in cases where it is necessary to break confidentiality, the confidential data should only be shared very sparingly with those that the computing professional judges to be appropriate authority

Section 2 (More Specific) Professional Responsibilities

Section 2 of Draft 1, consists of 8 specific principles related to the practice of computing. These principles can be grouped into general positions about ethics and computing, and more specific issues about practice. When considered together, they indicate ways in which the general moral principles of Section 1 apply to computing.

The changes made to Draft 1, Section 2 clarified the ethical components of some of the principles. **Principle 2.1**, which had asserted the relation of quality to ethics, now describes quality as an ethical concern because of the potential negative impacts of products on anyone directly or indirectly affected by the work. **Principle 2.2**, which had described the importance of technical competence in a domain before undertaking a task, now makes it clear that technical competence is not enough, but that quality requires an understanding and application of potential ethical issues in the development and deployment of an artifact.

A focus in the 1992 Code was on skilled production and moral responsibility to deliver a quality product. **Principle 2.2** is modified to dispel the mistaken belief that mere competent creation of an artifact is an adequate model of professional behavior. **Principle 2.2** was consistent with this narrow competent creation view. This is a view of software developers merely as skilled craftspeople who focuses narrowly on the object of creation. This view misses the additional need to address the socio- technical artifacts and their positive and negative societal impacts. **Principle 2.2** only addressed technical knowledge, ignoring the responsibility to use their knowledge to advise the client of the values and virtues of building the system in the way requested. The Draft 1 addition to “**Principle 2.2** Professional competence includes technical knowledge, awareness of the social context in which the work will be deployed, and competence in recognizing and navigating ethical challenges” is intended to dispel the myth that competent creation alone is adequate for computing professionals.

“**Principle 2.4** Accept and provide appropriate professional review” is an example of what is required to move away from an overly narrow reliance on competent creation. Reviewing work helps to improve its final quality and in catching potential risks. **Principle 2.4** originally advocated this form of helpful critical review. Unfortunately, some understood the word “critical” to mean “severe, negative, or destructive,” and such reviews have contributed to destructive bullying, harassment, and the delivery of lower quality work. To reduce this misinterpretation, “constructive” was added as an adjective to “critical review.”

“**Principle 2.5** Give comprehensive and thorough evaluations of computer systems and their impacts, including analysis of possible risks” focuses on a particular type of review. There was minimal change to this clause but it provoked discussion in the Task Force and web commenters surrounding the scope of applicability of the Code. This was taken up in the discussion of the Preamble and again in the discussion

of Section 3.

Commenters raised an additional concern that several elements in this section are impractical, if not impossible, and too idealistic. Reading this principle as a rule, some are concerned that the Code requires the impossible in that practitioners are expected to foresee every possible problem with a system when they may only see part of the system. Those concerned requested a change to limit the scope of this anticipatory requirement to “reasonable knowledge” or “direct and foreseeable” consequences.

There are two problems here, besides the rule/principle confusion. First adding the constraint that we only have to anticipate “direct” consequences moves back toward mere competent creation, removing the requirement to think of our work in socio-technical contexts. **Principle 1.2** also emphasizes addressing indirect harm to others. Second, there is no requirement for computing professionals to foresee every consequence, but rather to address the anticipatable negative ethical impacts of their work.

Others are concerned with the practical issue of how one calls their employer unethical, without getting fired. There is a request for training how to do this. There are clear mechanisms for doing this as illustrated by the recent refusal of a group of some computing technologists to develop an ethnic registry. The historical reasons for their refusal are given and they have taken a pledge (Neveragain, 2016).

We, the undersigned, are employees of tech organizations and companies based in the United States. We are engineers, designers, business executives, and others whose jobs include managing or processing data about people. We are choosing to stand in solidarity with Muslim Americans, immigrants, and all people whose lives and livelihoods are threatened by the incoming administration’s proposed data collection policies. We refuse to build a database of people based on their Constitutionally-protected religious beliefs. We refuse to facilitate mass deportations of people the government believes to be undesirable.

“**Principle 2.6** Honor contracts, agreements, and assigned responsibilities” categorizes reasons for not doing a project unless the problems identified are resolved. Reasons for not doing a project include a lack of technical skill or available domain expertise, addressable problems, theoretically impossibility, or as in the above Neveragain case, there are overriding ethical issues.

“**Principle 2.7** Improve public understanding of computing and its consequences” indicates that computing professionals, consistent with the principle of helping society, have a responsibility to create an awareness of computing’s impacts, limitations, vulnerabilities, and opportunities. This promotion of educated users reduces the potential for unanticipated uses of a product.

Section 2 of the Code also includes some more practical matters. When dealing with computing practice, legal and ethical principles sometimes overlap and sometimes they are in contention. **Principle 2.3** addresses that tension and recognizes that in some cases you must challenge a law that may be ethically problematic or may conflict with a more important law. This clause presumes a hierarchy of laws, some of which have solid ethical foundation while others may be purely pragmatic. One of the critical issues with this clause is how to clarify this distinction for the reader.

Section 3 Leadership Obligations

Section 3 of the Code applies to organizational leaders. Draft 1 of the update left the first two imperatives almost untouched. The guidelines for **Principle 3.1** now mention transparency as a component of reducing harm to the public. The guidelines for **Principle 3.2** explicitly mention psychological well-being as a

consideration for workers in the implementation of a system. **Principle**

3.3 has been clarified to include leadership's responsibility for the establishment of appropriate rules regarding the use of the organization's computing resources. Additionally, that principle establishes that the rules must also cover the information that the company has on those resources. These additions did not engender many comments other than declarations that they were welcome.

The remaining principles underwent more substantial change. **Principle 3.4** and its guidelines, which focused on system users and those affected by a system, were struck in their entirety and replaced with a principle and guidelines that call on leadership to think much more broadly. The principle now puts the public good at the center of all professional computing work. Furthermore, this responsibility is given leadership as it is explicitly placed in the section of the Code directed at leaders. Given the adoption of a variety software development techniques that have supplanted the traditional waterfall methodology, an interesting comment is that this responsibility needs to address more than organizational leadership. There are now other leadership roles, including team leaders, open source leaders, collaborative leaders and crowd-sourced leaders. So, not only have changes in the capabilities and uses of computing systems impacted the Code, but so have changes in the approaches to software development. Since most programming and design is now done in teams, team leadership responsibilities need to be specifically addressed.

Changes to **Principle 3.5**, while not directly focusing on the public good, instead focus on the individuals who use or are affected by the use of a computing system. **Principle 3.5** in conjunction with **Principle 3.4**, call upon leadership to consider both the individual, the individual in connection with society, and the good of society. One commenter on **Principle 3.5** questioned the inclusion of a definition of "dignity." The argument regarding whether the definition ought to be included is very much like the argument surrounding the definition of the "public good" mentioned in the section on the Preamble, as well as those considered in Section 1.

Changes to **Principle 3.6** include the addition of language that is explicit about accountability shared by all members of the organization for the limitations and impacts of software systems. The guidelines make explicit the seriousness of this responsibility by articulating that software errors are inevitable, and that software systems have impacts and are impacted by the contexts in which they are deployed.

The inclusion of **Principle 3.7** reflects what is perhaps the most fundamental change in the use of software since the 1992 Code. Software has increasingly become part of the infrastructure of society. As such, Draft 1 calls on leadership to recognize when this is happening. It is not always an easy thing to judge or to be aware of, but simple thought experiment can give some guidance. If this piece of software were to disappear, would there be serious immediate economic consequences beyond the software company that produced it? If so, that system is a candidate for infrastructure status.

Harder to measure, but just as important, is the social value a particular system delivers. What sort of social and personal hardship would the elimination of the software bring? If there are genuine concerns, then it is worth looking for a way to move those impacted onto some other system that can provide similar social or personal value.

Another commenter brought up a closely related notion. Certain companies have become gatekeepers for what computer programs can be distributed to and even used on certain devices. Those devices are part of the social infrastructure of their users. Companies that are gatekeepers can block particular apps, or completely bans certain developers, with no due process or recourse. Does the Code address this? Should the Code have another principle added that indicates that companies that control access to key infrastructure services must provide adequate due process? Can this concern be accommodated by enhancing the

guidelines for **Principle 3.7**?

An important part of professionalism for many ACM members surrounds the work in computing that is presented at conferences, especially ACM-sponsored conferences, or published in ACM journals. This raises the question of whether the Code should speak more directly to these sorts of activities in addition to the development of software systems. The case for the Code to address ACM journal publications might be more difficult to make. ACM journals engage in either a blind or double-blind review process and the journal's editorial team can monitor the review process. In addition, the ACM's Publications Board has adopted numerous policies that are consistent with the Code and may be sufficient.

Nonetheless, the question remains whether this section of the Code dedicated to Organizational Leadership Principles ought to more directly address this important function of the ACM.

A second area of concern is ACM-sponsored conferences. By their very nature, conferences are more dynamic than journals, and often, program review committees consist of both members and non-members. In addition, the acceptance or non-acceptance of one's work at a particular conference can have a significant impact on one's career. It is easy to suggest that the Organizational Leadership Imperatives ought to apply to those organizing conferences, yet very little of the language used in the principles and guidelines makes it clear how the principles might apply. Adding language to identify a more direct application of the principles to conference organizers and reviewers would be helpful.

Such language might also address a concern raised by a commenter who suggested that Section 3 of the Code seemed to apply only to those working in industry, rather than all ACM members. Essentially, the commenter was concerned that the Code suggested different levels of applicability depending on whether one worked in academia or in industry. A step further in the direction of showing that the Code applies to all organizational leaders with professional responsibilities related to computing would be for the guidelines to be explicit about applying to the leadership within the ACM.

Section 4. Compliance

Section 4 is the shortest section, both in the 1992 Code and in Draft 1. This does not mean that compliance with the Code is unimportant; compliance with the Code is a vital aspect of the interaction of the Code, the ACM, and its members. However, compliance, as described in the Code, is only part of the whole picture; there are also ACM policies that are separate from the Code that specify how the Code is to be enforced. Those policies are also being revised, although that revision is in an earlier stage of development than the update to the Code.

Section 4 of the Code is a guideline that establishes the goals of Code compliance. There are ongoing revisions of the relevant bylaws which establish procedures to assure that ACM members and officers observe due process when a member is accused of violating the Code. The current bylaws establish a single penalty for a code violation, namely expulsion from the ACM, regardless of the degree of violation. In the ACM regulations on plagiarism, several degrees of plagiarism are identified and assigned different penalties. The Task Force is considering mirroring that model.

Both the 1992 Code and Draft 1 have two principles in Section 4. The update adds some detail to the guidelines. For example, in the 1992 Code **Principle 4.1** is "Uphold and promote the principles of this Code." Draft 1 retains "uphold" and "promote" and adds "respect." This addition emphasizes that a member should not merely avoid direct violations of the Code, but should aspire to fulfill the spirit of the Code. Subsequent changes to **Principle 4.1** reflected the view that ACM members, as computing professionals

should “encourage and support adherence” to these principles by all of those working in computing. This change reflects the obligation of professionals to encourage quality work in all computing domains.

The update to **Principle 4.2** makes the text shorter and more direct. The “voluntary” nature of compliance is no longer mentioned, and instead, **Principle 4.2** states that “If an ACM member does not follow this code, membership in the ACM may be terminated.” The revision is intended to be as clear as possible about the ACM’s commitment to Code compliance. A commenter suggested that the ACM consider cooperating with other professional organizations pursuing sanctions against violators of the Code.

General Comments on the Code

A number of general comments were submitted about the proposed revision. These comments spanned the breadth of the Code, or were outside the realm of any single principle. Some of these comments were philosophical in nature, but some were also quite practical, focusing on the context of the Code and methods for making it more accessible.

Some respondents were keen on keeping the current state of the Code. Although we do agree that the Code still reads “remarkably well” today and has stood the test of time, there are ambiguities and inconsistencies in language that do not hold up in the modern computing environment. Others suggested that they appreciated the update and already found Draft 1 a noticeable improvement.

Some comments were in the form of suggestions for new principles, such as suggestions for protecting workers, ensuring fair and equal pay, a prohibition against assisting authoritarian regimes in restricting freedom, prohibiting sexual harassment, and requiring educational outreach to enable the general public to make more informed decisions about technology use. Before pursuing the study needed to ensure that the right essence is captured in the wording of a principle, we will explore whether it is already captured in existing principles of the Code.

General comments on the Code continued to express concern about the complexity of words and phrases. In addition to those mentioned in previous sections, “informed decision,” “privacy,” and “quality” are terms that commenters identified as being ambiguous. However, as suggested above for other terms, there is a risk in defining such terms too narrowly. And, as a practical matter, extensive definitions would lengthen the Code significantly. Additionally, one of the things that have made the Code so applicable over the years is that the breadth of the terms allows them to be applied in many more situations as the sociotechnical systems that they are related to change. Of course, this could also open the Code to potential “loopholes”, where the ambiguity of a concept could potentially leave out specific instances of similar behavior, but this change when the Code is viewed as guidance, rather than a set of rules to be followed. Unfortunately, even if we were to attempt to “patch” these loopholes, we can never identify all current and future applications or all contexts for particular concepts. So instead we endeavor to be as comprehensive as we can, allowing for the clear language but future-proofing through general concepts whose specific meanings might change over time.

Specific technologies were also raised as special cases that could need highlighting in the Code. Artificial intelligence is one such technology; commenters are concerned about the significant impact it could have on all aspects of computing and society. Who would be responsible for mistakes that artificially intelligent agents make? Would such agents have ethical or moral agency? Should there be a specific code (as the rules Asimov proposed) for robots? The idea of a code of ethics for specific technologies is outside the scope of the Code, but it may be that some of the key underlying aspects raised by this technology should

be incorporated into the Code. For example, perhaps there should be a principle for intelligent agents. (As before, this raises the problem of defining that term.) Perhaps, instead, clear guidelines for the professional and organizational responsibility that articulate that professionals involved in implementing artificial intelligence must take responsibility for the actions of their creations. These are not new arguments or discussions, but the implications for society are potentially significant, and thus should be considered for this round of the Code.

Another question that was raised in this process was: what responsibility do professionals have to actively raise awareness about ethics? (Or, as the commenter put it, “be vocal?”) The Code requires professionals to act in accordance with the Code and report violations; but should the Code also require proactively promoting ethical thinking in the profession? Another commenter suggested that professionals have an ethical responsibility for ongoing “reflective self-examination” (as per Vallor’s work on virtue ethics (Vallor, 2016)), which could complement this proactive approach. Is this, however, something that should go into a code of ethics, or could it be more an approach that education could foster? The danger of scoping in wider philosophical reflection on behavior is that it becomes difficult to codify in a way that can apply to all situations. On the one hand, teaching up-and-coming professionals self-reflection and taking an ethical perspective could be more effective in motivating an industry to be critical about its own behavior and actively promote ethical thought and action. On the other hand, we see breaches of ethical practice frequently enough to challenge whether educators are doing enough to promote professional practice with the currently existing strategies. Could the encouragement, or even requirement for “vocal” action be a “next generation” solution? Clearly, efforts external to the Code that provide support materials for those educators would facilitate such efforts.

Others questioned whether the Code could capture the idea of a requirement for critical analysis of technology prior to its release. Such critical analysis is reminiscent of methods used in technology assessment, a common practice in European circles (to the point that there is a European Parliamentary Technology Assessment network (2017)) in evaluating the potential impact of a new technology on society. Whether frameworks such as technology assessment, responsible innovation, or other current structured methods for critically reflecting on technologies and building ethically and socially aware innovation and research models are worth articulating in the Code is debatable. On the one hand, some of these methods are quite useful as general good practice and approaches to determining the ethical and social impact of technologies; on the other hand, it may be that the principles are already in the Code and the more practical implementation of these principles should be kept separately from the Code itself. Whichever way this goes, the Code should reflect the key principles identified by the commenter: engendering trust in the professional sector through ensuring that the artifacts developed are socially and ethically desirable, perform reliably and to specification, and that any potential risks, dual uses, or other issues are identified and mitigated before release.

A discussion of conflicts between ethics and law arose in the comments. The professional organization for psychologists in the US (the American Psychological Association) requires that psychology professionals hold to the higher ethical standard in ethical-legal conflicts, but if no reconciliation between the law and the ethical standard can be made they are permitted to fall back to the law, while making it known that they are committed to the APA Code of Ethics. Should the Code have a similar statement? The ACM has members in many different countries with different legal structures and standards, and obviously, the laws in these countries are likely to be different. However, the ethical standard for the Code should be as universal as possible – concentrating on the practice of being a computing professional rather than attempting to codify social or political norms of specific countries.

Even though it is an international organization, the ACM sometimes appears to some as if it is an American

organization. Care must be taken to ensure that the values and norms of the US do not have undue influence on the Code. For example, free speech is understood differently across different polities. Indeed, the same Internet publication may be identified as protected free speech in one country and prohibited hate speech in a different country. It is important that those developing the Code are sensitive to these nuances. Perhaps, however, it would be sufficient, as in the APA Code, to let the professional determine for herself whether it is appropriate to uphold the Code in the face of a legal dispute or to fall back on the law if adhering to the Code is too difficult in the circumstances. (We note that this strategy could be abused as an ethical loophole described previously.) Another commenter suggested that technology is limited by stereotypes of masculinity and that the Code and thus the profession would be well-served by including new values not found in the Code; the ethics of (1) prudence, (2) active responsibility, (3) protection and care, (4) saving, and (5) commitment. This addition would increase the breadth of values expressed in the Code.

The general comments we received often raise more questions than they solve – but these are important aspects of the feedback received that should be the topic of discussion for the area as a whole so that we can successfully develop a theoretically sound, yet accessible and practical Code of Ethics for computing professionals.

Possible New Imperatives

Any code of ethics in the areas of computing and technology must consider where to adapt to the ever-increasing scope, depth, and integration of computing into society. In the twenty-six years between the 1992 Code and Draft 1, we've gone from floppy disks to cloud computing and from the Nokia 1011 to the iPhone. Commenters noted this transformation of technology, and its transformative role in our daily lives and professions, with suggestions of some specific areas for inclusion in the 2018 code:

- Should organizations that control infrastructure be required to provide a measure of transparency regarding application availability and review processes, as well as a commitment to due process should someone want to appeal?
- Should this code include guidelines for how to apply the principles in realistic situations, particularly in places where these guidelines are not being followed?
- Should organizations or companies be required to ensure that all people involved in their production processes are paid appropriately for their work, and that their work environments are safe and clean?
- Should the Code make recommendations regarding the active inclusion of minorities and under-represented groups in the development and discussions around technology and computing?
- Should this code provide boundaries regarding the use of user data/metadata, particularly regarding potential usages by authoritarian regimes that do not honor accepted human rights conventions?

Conclusion

One of the most difficult issues for the Task Force in suggesting updates to the Code is dealing with the limiting notion that a code of ethics is a set of fixed rules rather than a set of extensible, scalable guidelines that rely on human intelligence and thought in application. The Code needs to provide guideline and reminders for anyone working in computing that when they are making complex decisions they are affecting themselves, the people they work for, their users, and most importantly society.

References

- ACM. (1992, October 16). ACM Code of Ethics and Professional Conduct. Retrieved July 22, 2017, from <http://www.acm.org/about-acm/acm-code-of-ethics-and-professional-conduct>
- Davis, M. (2002). *Profession, code, and ethics*. Aldershot, Hants, England ; Burlington, VT: Ashgate.
- European Parliamentary Technology Assessment EPTA. (2017). Welcome to the network of parliamentary technology assessment! Retrieved January 27, 2017, from <http://www.eptanetwork.org/>
- Ladd, J. (1980). The quest for a code of professional ethics: an intellectual and moral confusion. In R. Chalk, M. S. Frankel, & S. B. Chafer (Eds.), *AAAS Professional Ethics Project*. Washington, DC: American Association for the Advancement of Science. Retrieved from <http://dl.acm.org/citation.cfm?id=2569.2570>
- Neveragain. (2016). Our Pledge. Retrieved July 22, 2017, from <http://neveragain.tech/>
- Vallor, S. (2016). *Technology and the virtues: a philosophical guide to a future worth wanting*. New York, NY: Oxford University Press.