

PART I

What is Computer Ethics?

It is not enough that you should understand about applied science in order that your work may increase man's blessings. Concern for man himself and his fate must always form the chief interest of all technical endeavors.

Albert Einstein

Editors' Introduction

In the 1940s and early 1950s, the field of study that is now called “computer ethics” was given a solid foundation by Professor Norbert Wiener of MIT. Unhappily, Professor Wiener’s works in computer ethics were essentially ignored for decades by other thinkers. In the 1970s and 1980s computer ethics was recreated and redefined by thinkers who did not realize that Wiener had already done so much work in the field. Today, more than 50 years after Wiener created computer ethics, some thinkers are still attempting to define the nature and boundaries of the subject. Let us briefly consider five different definitions that have been developed since the 1970s.

Maner’s Definition

The name “computer ethics” was not commonly used until the mid-1970s when Walter Maner began to use it. He defined this field of study as one that examines “ethical problems aggravated, transformed or created by computer technology.” Some old ethical problems, he said, were made worse by computers, while others came into existence because of computer technology. He suggested that we should use traditional ethical theories of philosophers, such as the *utilitarian* ethics of the English philosophers Jeremy Bentham and John Stuart Mill, or the *rationalist* ethics of the German philosopher Immanuel Kant.

Johnson’s Definition

In her book, *Computer Ethics* (1985), Deborah Johnson said that computer ethics studies the way in which computers “pose new versions of standard moral problems and moral dilemmas, exacerbating the old problems, and forcing us to apply ordinary moral norms in uncharted realms.” Like Maner

before her, Johnson adopted the “applied philosophy” approach of using procedures and concepts from utilitarianism and Kantianism. But, unlike Maner, she did not believe that computers create wholly new moral problems. Rather, she thought that computers gave a “new twist” to ethical questions that were already well known.

Moor’s Definition

In his influential article “What Is Computer Ethics?” (1985), James Moor provided a definition of computer ethics that is much broader and more wide-ranging than those of Maner or Johnson. It is independent of any specific philosopher’s theory; and it is compatible with a wide variety of approaches to ethical problem-solving. Since 1985, Moor’s definition has been the most influential one. He defined computer ethics as a field concerned with “policy vacuums” and “conceptual muddles” regarding the social and ethical use of information technology:

A typical problem in Computer Ethics arises because there is a policy vacuum about how computer technology should be used. Computers provide us with new capabilities and these in turn give us new choices for action. Often, either no policies for conduct in these situations exist or existing policies seem inadequate. A central task of Computer Ethics is to determine what we should do in such cases, that is, formulate policies to guide our actions. . . . One difficulty is that along with a policy vacuum there is often a conceptual vacuum. Although a problem in Computer Ethics may seem clear initially, a little reflection reveals a conceptual muddle. What is needed in such cases is an analysis that provides a coherent conceptual framework within which to formulate a policy for action. (Moor 1985, p. 266)

Moor said that computer technology is genuinely revolutionary because it is “logically malleable”:

Computers are logically malleable in that they can be shaped and molded to do any activity that can be characterized in terms of inputs, outputs and connecting logical operations. . . . Because logic applies everywhere, the potential applications of computer technology appear limitless. The computer is the nearest thing we have to a universal tool. Indeed, the limits of computers are largely the limits of our own creativity. (Ibid.)

According to Moor, the computer revolution will occur in two stages. The first stage is that of “technological introduction” in which computer technology is developed and refined. This already occurred during the first 40 years after the Second World War. The second stage – one that the

industrialized world has only recently entered – is that of “technological permeation” in which technology gets integrated into everyday human activities and into social institutions, changing the very meaning of fundamental concepts, such as “money,” “education,” “work,” and “fair elections.”

Moor’s way of defining computer ethics is very powerful and suggestive. It is broad enough to be compatible with a wide range of philosophical theories and methodologies, and it is rooted in a perceptive understanding of how technological revolutions proceed.

Bynum’s Definition

In 1989 Terrell Ward Bynum developed another broad definition of computer ethics following a suggestion in Moor’s 1985 paper. According to this view, computer ethics *identifies and analyzes the impacts of information technology on such social and human values as health, wealth, work, opportunity, freedom, democracy, knowledge, privacy, security, self-fulfillment, etc.* This very broad view of computer ethics employs applied ethics, sociology of computing, technology assessment, computer law, and related fields. It employs concepts, theories, and methodologies from these and other relevant disciplines. This conception of computer ethics is motivated by the belief that – eventually – information technology will profoundly affect everything that human beings hold dear.

Gotterbarn’s Definition

In the 1990s, Donald Gotterbarn became a strong advocate for a different approach to computer ethics. From his perspective, computer ethics should be viewed as a branch of *professional ethics*, concerned primarily with standards of good practice and codes of conduct for computing professionals:

There is little attention paid to the domain of professional ethics – the values that guide the day-to-day activities of computing professionals in their role as professionals. By computing professional I mean anyone involved in the design and development of computer artifacts. . . . The ethical decisions made during the development of these artifacts have a direct relationship to many of the issues discussed under the broader concept of computer ethics. (Gotterbarn 1991, p. 26)

With this “professional ethics” approach to computer ethics, Gotterbarn co-authored the 1992 version of the ACM Code of Ethics and Professional Conduct and led a team of scholars in the development of the 1999 ACM/IEEE Software Engineering Code of Ethics and Professional Practice. (Both of these codes of ethics are included in this book in Part III.)

Each of these definitions of computer ethics has influenced this textbook to some extent. Part I makes special use of the ideas of Moor and Maner; later parts of the book bring in other ideas as well.

References

- Gotterbarn, D. (1991). "Computer Ethics: Responsibility Regained." *National Forum: The Phi Beta Kappa Journal*, 71: 26–31.
- Johnson, D. G. (1985). *Computer Ethics*. Prentice-Hall (2nd edn 1994; 3rd edn 2001).
- Maner, W. (1980). *Starter Kit in Computer Ethics*. Helvetia Press (published in cooperation with the National Information and Resource Center for Teaching Philosophy). (Originally self-published by Maner in 1978.)
- Moor, J. H. (1985). "What Is Computer Ethics?" In T. W. Bynum (ed.), *Computers and Ethics*. Blackwell, pp. 266–75. (Published as the October 1985 issue of *Metaphilosophy*.)

CHAPTER 1

Reason, Relativity, and Responsibility in Computer Ethics

James H. Moor

Searching for Ethics in the Global Village

As computing becomes more prevalent, computer ethics becomes more difficult and more important. As Terry Bynum and Simon Rogerson put it,

We are entering a generation marked by globalization and ubiquitous computing. The second generation of computer ethics, therefore, must be an era of “global information ethics.” The stakes are much higher, and consequently considerations and applications of Information Ethics must be broader, more profound and above all effective in helping to realize a democratic and empowering technology rather than an enslaving or debilitating one. (1996, p. 135)

I heartily concur with the concern that Bynum and Rogerson express about the global impact of computing. The number and kinds of applications of computing increase dramatically each year and the impact of computing is felt around the planet. The ubiquitous use of electronic mail, electronic funds transfer, reservations systems, the World Wide Web, etc. places millions of the inhabitants of the planet in a global electronic village. Communication and actions at a distance have never been easier. We are definitely in a computer revolution. We are beyond the introduction stage of the revolution in which computers are curiosities of limited power used only by a few. Now, entire populations of developed countries are in the permeation stage of the revolution in which computers are rapidly moving into every aspect of daily life.

James H. Moor, “Reason, Relativity, and Responsibility in Computer Ethics.” This chapter was originally presented as the keynote address at ETHICOMP96 in Madrid, Spain and later published in *Computers and Society*, 28:1 (March 1998), pp. 14–21. © 1998 by James H. Moor and reprinted by permission of the author.

The computer revolution has a life of its own. Recently [i.e., in 1996], in northern California about one-sixth of the phone calls didn't connect because of excessive use of the Internet. People are surging to gain access to computer technology. They see it as not only a part of their daily lives but a necessary venue for routine communication and commercial transactions. In fact, the surge has become so great that America On Line, a prominent Internet service provider, offered its customers refunds because the demand for connection overwhelmed the company's own computer technology. The widespread desire to be wired should make us reflect on what awaits us as the computer revolution explodes around the world. The digital genie is out of the bottle on a worldwide scale.

The prospect of a global village in which everyone on the planet is connected to everyone else with regard to computing power and communication is breathtaking. What is difficult to comprehend is what impact this will have on human life. Surely, some of the effects will be quite positive and others quite negative. The question is to what extent we can bring ethics to bear on the computer revolution in order to guide us to a better world or at least prevent us from falling into a worse world. With the newly acquired advantages of computer technology, few would want to put the genie completely back into the bottle. And yet, given the nature of the revolutionary beast, I am not sure it is possible to completely control it, though we certainly can modify its evolution. Aspects of the computer revolution will continue to spring up in unpredictable ways – in some cases causing us considerable grief. Therefore, it is extremely important to be alert to what is happening. Because the computer revolution has the potential to have major effects on how we lead our lives, the paramount issue of how we should control computing and the flow of information needs to be addressed on an ongoing basis in order to shape the technology to serve us to our mutual benefit. We must remain vigilant and proactive so that we don't pillage the global village.

Although almost everyone would agree that computing is having a significant, if not a revolutionary, impact on the world, and that ethical issues about applications of this surging technology should be raised, there is disagreement about the nature of computer ethics. Let me describe two positions with which I disagree. These two positions are both popular, but represent opposite extremes. I believe they mislead us about the real nature of computer ethics and undercut potential for progress in the field. The first view I will call the "Routine Ethics" position. According to the Routine Ethics position, ethical problems in computing are regarded as no different from ethical problems in any field. There is nothing special about them. We apply established customs, laws, and norms, and assess the situations straightforwardly. Sometimes people steal cars and sometimes people steal computers. What's the difference? The second view is usually called "Cultural Relativism." On this view, local customs and laws determine what is

right and wrong, but, because computing technology such as the World Wide Web crosses cultural boundaries, the problems of computer ethics are intractable. Free speech is permitted in the United States but not in China. How can we justify a standard for or against free speech on the World Wide Web? Routine Ethics makes computer ethics trivial and Cultural Relativism makes it impossible.

I believe that the views of both Routine Ethics and Cultural Relativism are incorrect, particularly when used to characterize computer ethics. The former underestimates the changes that occur in our conceptual framework and the latter underestimates the stability of our core human values. The problems of computer ethics, at least in some cases, are special and exert pressure on our understanding. And yet our fundamental values, based on our common human nature, give us an opportunity for rational discussion even among cultures with different customs. The purpose of this chapter is to explain how it is possible to have both reason and relativity in computer ethics. Only with such an understanding is responsibility in computer ethics possible.

Logical Malleability and Informational Enrichment

Computers are *logically malleable*. This is the feature that makes computers so revolutionary. They are logically malleable in that they can be manipulated to do any activity that can be characterized in terms of inputs, outputs, and connecting logical operations. Computers can be manipulated syntactically and semantically. Syntactically, one can alter what the computer does by changing its program. And semantically one can use the states of a computer to represent anything one chooses, from the sales of a stock market to the trajectory of a spacecraft. Computers are general purpose machines like no others. That is why they are now found in every aspect of our lives and that is why a computer revolution is taking place.

Computers are also *informationally enriching*. Because of their logical malleability, computers are put to many uses in diverse activities. Once in place, computers can be modified to enhance capabilities and improve overall performance even further. Often, computerized activities become informationalized; i.e., the processing of information becomes a crucial ingredient in performing and understanding the activities themselves. When this happens, both the activities and the conceptions of the activities become informationally enriched.

The process of informational enrichment is gradual and is more manifest in some activities than in others. What is striking is how often and the extent to which it does occur. In a typical scenario a computer is introduced merely as a tool to perform a job or to assist in an activity. Gradually, the computer

becomes an essential part of the methodology of doing the job or performing the activity. To do it properly is to use a computer. Over time, the job or activity is viewed increasingly as an informational phenomenon, so that information processing is taken as a salient or even defining feature.

Consider some examples of informational enrichment. At one time in the United States money was backed by gold. There was an exchange of paper bills, but the bills were merely coupons that could, at least in principle, be redeemed for gold or perhaps silver. For some time the US remained on the gold standard so that paper bills were markers for money. Monetary transactions were grounded in gold. Then the gold standard was dropped and the paper bills became the money. To have money was to have the paper, presumably backed by the good faith and trust in the government. Now paper has been augmented with credit cards and debit cards that can be read by computers. Of course, these cards are not the real money because one can always exchange the credits for paper money. But, it is likely that the use of paper money will decrease and the electronic tokens on the cards or in a bank's computer will become the money. Some cards now have chips embedded in them so that they can be loaded with electronic money which is then transferred as information to a merchant at the point of sale. We are headed for a cashless society. Monetary transactions are increasingly grounded in information. Money may come to be conceived as an elaborate computable function among people. In the computer age the concept of money is becoming informationally enriched.

As another example of informational enrichment, consider the evolving nature of warfare. Traditionally, in warfare different sides send people into battle who fight with each other at close quarters until one side has killed or captured so many that the other side surrenders. People are still sent to the battlefield, but warfare is rapidly becoming computerized. The stealth bomber used by the United States during the Gulf War [in 1991] was the result of computerized engineering. Computers designed the shape of the aircraft so that it would be nearly invisible to radar. The aircraft's design deprived Iraq of information. The Gulf War was about information and the lack of it. Bombs were dropped and guided by lasers and computers. Missiles were launched from ships and sought their targets by reading the terrain using computer guidance systems. The first objective of the armed forces under General H. Norman Schwarzkopf's command was to eliminate the ability of Iraq to communicate among its own forces or to use its aircraft detection systems. Schwarzkopf remarked after the war that it was the first time an enemy was brought to his knees by denial of information. As war becomes increasingly computerized, it may be less necessary or desirable to send men and women into the battlefield. Wars ultimately will be about the destruction of information or the introduction of misleading information. One side surrenders when it is not able to obtain and control certain kinds of

information. This may not be a bad result. Better that data die, than people. As warfare becomes increasingly computerized, our concept of war becomes informationally enriched. The information processing model is seizing the high ground.

Informational enrichment can also affect ethical and legal practices and concepts. Consider the concept of privacy as it has evolved in the United States as an example (Moor 1990). Privacy is not explicitly mentioned in the Declaration of Independence or in the Constitution of the United States, though there are portions of these documents which implicitly support a notion of privacy as protection from governmental intrusion, particularly the physical invasion of people's houses. The notion of privacy has been an evolving concept in the US. For instance, in the 1960s and '70s the legal concept of privacy was expanded to include protection against government interference in personal decisions about contraception and abortion. Today, the concept of privacy includes these earlier elements but increasingly focuses on informational privacy. This shift in emphasis has been brought about because of the development of the computer and its use in collecting large databases of personal information.

The computer, originally viewed by many as little more than an electronic filing cabinet, rapidly revealed its potential. Once data is entered into a computer it can be sorted, searched, and accessed in extraordinarily easy ways that paper files cannot be in practical amounts of time. The activity of storing and retrieving information has been enhanced to the extent that all of us now have a legitimate basis for concern about the improper use and release of personal information through computers. The computerization of credit histories and medical records for use in normal business provides an ongoing possibility for misuse and abuse. Because of the widespread application of computer technology, our concern about privacy today goes far beyond the original concern about the physical intrusion of governmental forces into our houses. Now concerns about privacy are usually about improper access and manipulation of personal information by the government and many others who have access to computerized records. The original concept of privacy in the United States has become informationally enriched in the computer age.

Even concepts that begin as informational concepts can be informationally enriched. As an example, consider the legal concept of copyright. Legislation protecting the products of authors and inventors is authorized by the Constitution of the United States. Early copyright laws were passed to protect literary works, and patent laws were passed to protect inventions. Copyright laws in the US have been amended over the years to extend the length of protection to authors and to protect a wider and wider range of materials including music and photographs. But until the computer age the underlying conception of copyright was that it was intended to protect those

items which could be read and understood by humans. For example, in the early part of the twentieth century an attempt to protect piano rolls by copyright was denied on the grounds that piano rolls were not in human readable form.

In the 1960s programmers began to submit copies of printouts of their programs for copyright protection. The printouts were in human readable form. But what programmers wanted to protect was not the printouts of programs but the programs as they existed on computers. However, the programs, as they existed on computers, were not in human readable form. If the human readable printouts were to count as surrogates to protect the machine versions of programs, copyright law had to be stretched. Moreover, if machine-readable programs were protectable by copyright, then it would seem that programs as instantiated on computer chips might be protectable by copyright as well. Copyright protection was so extended. Through the development of computing, the concept of copyright has become informationally enriched. Copyright extends not only to computer languages, but to computer languages in forms readable only by machines. Indeed, what is copyrightable today sometimes looks more like an invention than a literary work.

I have used the concepts of money, war, privacy, and copyright as examples of informational enrichment. There are many more. It is difficult to think of an activity now being done extensively by computers that has not been informationally enriched. In some cases this enrichment is so salient that our concepts shift somewhat. They too become informationally enriched. In the computer age, we live in a different world.

The Special Nature of Computer Ethics

I maintain that computer ethics is a special field of ethical research and application. Let me begin by describing computer ethics and then making a case for its special nature.

Computer ethics has two parts: (i) the analysis of the nature and social impact of computer technology and (ii) the corresponding formulation and justification of policies for the ethical use of such technology. I use the phrase “computer technology” because I take the subject-matter of the field broadly to include computers and associated technology, including software, hardware, and networks (Moor 1985).

We need thoughtful analyses of situations in which computers have an impact, and we need to formulate and justify policies for using them ethically. Although we need to analyze before we can formulate and justify a policy, the process of discovery often comes in the reverse order. We know that computing technology is being employed in a given situation, but we are

puzzled how it should be used. There is a *policy vacuum*. For example, should a supervisor be allowed to read a subordinate's email? Or should the government be allowed to censor information on the Internet? Initially, there may be no clear policies on such matters. They never arose before. There are policy vacuums in such situations. Sometimes it may be simply a matter of establishing some policy, but often one must analyze the situation further. Is email in the workplace more like correspondence on company stationary in company files or more like private and personal phone conversations? Is the Internet more like a passive magazine or more like an active television? One often finds oneself in a *conceptual muddle*. The issues are not trivial matters of semantics. If someone's health status is discovered through email or an impressionable child is exposed to distressing material on the Internet, the consequences may be very damaging. Obtaining a clear conception of the situation from which to formulate ethical policies is the logical first step in analysis, although chronologically one's uncertainty about the appropriate policy may precede and motivate the search for conceptual clarification. Given a tentative understanding of the situation, one can propose and evaluate possible policies for proper conduct. The evaluation of a policy will usually require a close examination and perhaps refinement of one's values. Such policy evaluation may lead one back for further conceptual clarification and then further policy formulation and evaluation. Eventually, some clear understanding and justifiable policy should emerge. Of course, with the discovery of new consequences and the application of new technology to the situation, the cycle of conceptual clarification and policy formulation and evaluation may have to be repeated on an ongoing basis.

Because computers are logically malleable, they will continue to be applied in unpredictable and novel ways, generating numerous policy vacuums for the foreseeable future. Moreover, because computerized situations often become informationally enriched, we will continue to find ourselves in conceptual muddles about how precisely to understand these situations. This is not to say that we can't achieve conceptual clarity and that we can't formulate and justify reasonable policies. Rather, it is to point out that the task of computer ethics is, if not Sisyphean, at least ongoing and formidable. No other field of ethics has these features to the degree that computer ethics does. Computer ethics is not simply ethics rotely applied to computing. Typically, problems in computer ethics require more than straightforward application of ethical principles to situations. Considerable interpretation is required before appropriate policies can be formulated and justified. Of course, to say that computer ethics is a special field of ethics does not mean that every ethical problem involving computers is unique or difficult to understand. Stealing a computer may be a simple case of theft. A straightforward application of an ethical principle is appropriate. In such a situation there are no policy vacuums and no conceptual muddles. And to say that

computer ethics is a special field of ethics does not mean that other fields of applied ethics do not have some instances of policy vacuums and conceptual confusion. Medical technology raises questions about what policy to follow for brain-dead patients and conceptual questions about what counts as life. What is special about computer ethics is that it has a continually large number of evolving situations which are difficult to conceptualize clearly and for which it is hard to find justified ethical policies. Doing computer ethics is not impossible, but doing it typically involves much more than rote application of existing norms.

I have argued that computer ethics is special but is the subject-matter truly *unique*? The answer depends upon what one means by “the subject-matter.” If by “the subject-matter” one means “computing technology,” then computer ethics is unique, for computing technology possesses unique properties (Maner 1996). I believe its most important property is logical malleability, which explains the ongoing wave of revolution and generation of ethical problems. If by “the subject-matter” one has in mind the occurrence of some novel ethical issues, then computer ethics is not unique because other fields of ethics sometimes consider novel situations which require revisions of conceptual frameworks and new policy formulation. If by “the subject-matter” one means “the overall range, depth and novelty of ethical issues generated by a technology,” then computer ethics is unique. No other technology, as revolutionary as it may be for a given area, has and will have the scope, depth, and novelty of impact that computing technology has and will have. There is no mystery why computer ethics has a prominence that toaster ethics, locomotive ethics, and sewing machine ethics do not have.

In summary, what is unique about computer ethics is computing technology itself, and what makes computer ethics different as a field of ethics is the scope, depth, and novelty of ethical situations for which conceptual revisions and policy adjustments are required. Deborah Johnson, in her excellent introduction to computer ethics, avoids taking sides on the issue of the uniqueness of computer ethics and suggests that ethical issues surrounding computers are “*new species of old moral issues.*” Johnson goes on to say:

The metaphor of species and genus encompasses the element of truth on each side of the debate in that a new species has some unique characteristics making it different from any other species, but at the same time, the species has generic or fundamental characteristics that are common to all members of the genus. (1994, p. 10)

Perhaps, the ambiguity in the question about the uniqueness of computer ethics suggests this middle ground approach. But I believe that Johnson’s characterization of a problem of computer ethics as just another species of a

fixed ethical genus is somewhat misleading because the conceptual uncertainty generated by some problems in computer ethics affects not only our understanding of the particular situation but also the ethical and legal categories that apply to it. As I have suggested, ethical and legal categories, such as privacy and copyright, can shift in meaning as they become informationally enriched. The novelty of the species sometimes infects the genus! Whether or not one regards computer ethics as unique, computer ethics is definitely a demanding field of ethics which requires more than routine application of principles.

Reasons within Relative Frameworks

I have been arguing against understanding computer ethics in terms of Routine Ethics because the application of computing technology regularly produces policy vacuums and informational enrichment which promotes conceptual shifts, if not outright conceptual muddles. Computer ethics is not rote. But, the rejection of Routine Ethics leaves many people uncomfortable. If ethics is not routine, how can it be done at all? Retreating to a position of Cultural Relativism will not solve the problem. According to Cultural Relativism, ethical issues must be decided situationally on the basis of local customs and laws. Two problems immediately confront us with such a position with regard to computer ethics. First, because computing activity is globally interactive, appealing to local customs and laws will not in general provide us with an answer to what we should do when customs and laws conflict. On the World Wide Web information flows without regard to particular customs. Which customs should we apply in regulating it? To pick the customs of any one culture seems arbitrary. Do we pick the customs of the culture in which the information appears on the computer screen or the customs of the culture from which the information originates? Second, all of the difficulties with Routine Ethics continue to apply. A policy vacuum may occur for every culture. A computing situation may be so novel that there are no customs or laws established anywhere to cope with it. Initially, an appeal to Cultural Relativism may seem like a sophisticated and plausible attempt to escape the parochial limits of Routine Ethics, but on closer inspection it has the limitations of Routine Ethics and more.

The shortcomings and difficulties with Routine Ethics and Cultural Relativism may make one cautious about doing applied ethics at all. If people differ in their ethical judgments, how can disagreements be avoided or resolved? It is for this reason, I think, that computer scientists and others are sometimes reluctant to teach computer ethics. Ethical issues seem to be too elusive and vague. It is more comfortable to talk about algorithms, data structures, memory locations, and networks because there are facts of the

matter on these topics. The realm of values seems hopelessly virtual, never to be as substantial as the real realm of facts. But a safe retreat to a realm of pure facts where everything is black or white, true or false, is not possible. Every science, including computer science, rests on value judgments. If, for example, truth is not taken as a critical value by scientists, the enterprise of science cannot begin.

My position is that all interesting human enterprises, including computing, are conducted within frameworks of values. Moreover, these frameworks can be rationally criticized and adjusted. Sometimes they are criticized externally from the vantage point of other frameworks and sometimes they are critiqued internally. Some value frameworks, such as those in an emerging science like computer science, undergo rapid evolution. Other value frameworks are more stable. Value frameworks provide us with the sorts of reason we consider relevant when justifying particular value judgments. Human values are relative, but not simply in the shallow sense of Cultural Relativism. Our most basic values are relative to our humanity, which provides us with a shared framework in which to conduct reasoned arguments about what we ought to do.

My intent is not to search for a way to eliminate value disputes altogether, which I do not think is possible, but to show how some reasoned discussion about value issues is possible even when customs may be absent or in conflict. To say that values are relative means that they are not absolute; it does not mean they are random or uncommon or uncriticizable. Perhaps, reflecting about reasoning with relative values is like thinking about swimming for the first time. It seems impossible. Why doesn't one sink to the bottom? How can one move if the water moves when pushed? Why doesn't one drown? But, swimming without drowning is possible and so is reasoning with relative values. In fact, not only is it possible; we do it all the time. Given the relativity of values, is there any hope for rational discussion in computer ethics. Absolutely!

My presentation will be in two steps. First, I will discuss the ubiquity of non-ethical values and emphasize their use in every aspect of human activity – we cannot escape value decision-making even if we want to do so. I will use computer science itself as an example, though any interesting human enterprise could serve as an illustration. And, second, I will discuss the use of values in making ethical decisions. My position is that an accommodation between reasoned argument and relativity of values is possible. We can acknowledge the difference in values among people and among cultures and still engage in rational discussions about the best policies for using computer technology.

Let me begin by emphasizing the ubiquity of values in our lives. In every reasonably complex human activity decisions are made which require value choices at least implicitly. Cooks make value decisions about what

constitutes a good meal. Businesspeople make value decisions about good investments. Lawyers make decisions about good jurors. All of these endeavors utilize facts, but the facts are always in the escort of values. Each discipline has its own cluster of values which members of the discipline use in making decisions. Even scientists, who pride themselves in establishing facts, must utilize values at least implicitly. In order to gather the facts, scientists must know what counts as good evidence, what counts as good methodology, and what counts as good explanation. Values permeate our lives. I am not speaking here primarily of ethical values. Rather, these are the values of daily activities that make our activities purposeful. Values are so much a part of what we do that we often don't reflect on the fact that values are at work when we make ordinary decisions. Value judgments cannot be escaped by any of us in work or play. Values saturate our decision-making and are necessary for the flourishing of the activities of life.

Even if one agrees that non-ethical values cannot be escaped in doing ordinary activities, there is still the concern that the relativity of values makes it impossible to have reasoned disputes. After all, cooks, businesspeople, lawyers, and scientists disagree among themselves. To examine the problem of relativity of values, let us use the activity of computer science as an example. In doing computer science, like other sophisticated human activities, one must make decisions and these decisions utilize, often implicitly, sets of non-ethical values. These are the values of the discipline. For instance, a computer scientist knows what makes a computer program a good program. Here I am using "good" primarily in a non-ethical sense. A good computer program is one that works, that has been thoroughly tested, that doesn't have bugs, that is well structured, that is well documented, that runs efficiently, that is easy to maintain, and that has a friendly interface. All of the properties of a good program reflect values. They are the features that make one computer program *better* than another. Moreover, this set of related values, that constitutes a set of standards within computer science, is widely shared among computer scientists. Given these standards, rational discussions can be conducted about how to improve a particular computer program. Moreover, policies regarding good programming techniques can be reasonably justified relative to the set of standards. For instance, one might argue for a policy of using object-oriented programming on the grounds that it leads to fewer bugs and computer code that is easier to maintain.

Computer scientists, like everyone else, can have disagreements, including disagreements about the standards. But disagreements which might appear to be about values are sometimes merely disagreements about facts. If there is a disagreement about the justification of the policy to use object-oriented programming, the real disagreement may be about whether or not object-oriented programming really leads to fewer bugs and code that is

easier to maintain. Such a dispute might be put to an empirical test. In this situation it is not a dispute about the importance of bug-free, easily maintainable code, but about how well object-oriented programming achieves these valued goals. Thus, disputes that initially may strike us as irreconcilable disputes about values may really be disputes about the facts of the matter subject to empirical adjudication.

Naturally, computer scientists can also disagree about the values that make up a good computer program as well. Some may rank documentation as essential and others may take it to be a less important optional feature. Depending upon the ranking of the different values, different judgments can be made regarding which programs are better than others and which policies about constructing computer programs are the most important. What I want to emphasize, however, is the degree of consensus that exists among computer scientists about what constitutes a good computer program. The specific rankings may differ somewhat from person to person, but a pattern of agreement emerges about the types of program that are the best. No computer scientist regards an ineffective, untested, buggy, unstructured, undocumented, inefficient, unmaintainable code with an unfriendly interface as a good program. It just doesn't happen. In a sense, the shared standards define the field and determine who is qualified and, indeed, who is in the field at all. If one prefers to produce buggy, "spaghetti code" programs, one is not doing serious computer science at all.

Discussions of the relativity of values sometimes engage in the "Many/Any Fallacy". This fallacy occurs when one reasons from the fact that many alternatives are acceptable to the claim that any alternative is acceptable. There are many acceptable ways for a travel agent to route someone between Boston and Madrid. It doesn't follow that any way of sending someone between these cities is acceptable. Traveling through the center of the Earth and going via the North Star are not included. Many different computer programs may be good, but not just any computer program is good.

To summarize, non-ethical values play a role in our decision-making in all interesting human activities, including computer science. No escape to a safe realm of pure facts, even in science, is ever possible. The standards of value of a discipline may be widely shared, implicit, and go unnoticed, but they are always there. Moreover, every discipline has sufficient agreement upon what the standards are to conduct its business. Without some consensus on what is valuable, progress in a discipline is impossible.

Core Values

Given that some consensus about values within communities with shared preferences exists, is there any basis for consensus about values among

communities? Ethical judgments are made beyond the narrow bounds of special interest communities. Given differences among communities, let alone differences among cultures, how is it possible to ground ethical judgments? Ethical judgments about computing technology may seem even more dubious. Because computing technology generates policy vacuums, i.e., creates situations in which there are no established policies based on custom, law, or religion, we are confronted with the difficult task of justifying ethical policies about novel applications of computing technology even within a community.

To address these challenges, we must begin by asking whether we share any values as human beings. What do we have in common? I believe that there is a set of core values which are shared by most, if not all, humans. They are familiar to all of us. Life and happiness are two of the most obvious such values. At the very least, people want to avoid death and pain for themselves. Of course, in some situations people give up their lives and suffer pain to accomplish certain objectives. But, generally speaking, people do not intentionally hurt and kill themselves for no reason. There is a *prima facie* value on life and happiness for humans. Other core values (or core goods) for humans include ability, freedom, knowledge, resources, and protection. These values are articulated in different ways in different cultures, but all cultures place importance on these values to some extent. Obviously, some cultures may distribute these goods unequally among their members, but no culture disregards these values completely. No culture or individual human could continue to exist and disregard the core values completely. Humans need nourishment and cultures need to raise their young to survive. These kinds of activity require at least some ability, freedom, knowledge, resources, and protection. The fact that humans share some basic values is not surprising. These values provide some evolutionary advantages. Individuals and cultures that completely neglect the core goods will not exist for very long.

The core values provide standards with which to evaluate the rationality of our actions and policies. They give us reasons to favor some courses of action over others. They provide a framework of values for judging the activities of others as well. As we become acquainted with other cultures, differences often strike us. The members of other cultures eat different meals, wear different clothing, and live in different shelters. But at a more abstract level people are remarkably alike. Initially, we may find the habits of others to be strange, silly, or bizarre, but after investigation we don't find them to be unintelligible. Activities that may appear at first to be random or purposeless are in fact ordered and purposeful. This doesn't make the practices of others uncriticizable, any more than our own are uncriticizable, but it does make them understandable.

Discussions of relativism in ethics often include examples of the Many/Any Fallacy. Many different customs exist, and, so it is argued, any custom

may exist. Not so! Some possible practices are ruled out and other practices (in some form or other) are required if a culture is to exist. Core human values are articulated in a multitude of delightful ways, but they also constrain the realm of possibilities. “Relative” doesn’t mean “random.”

To say that we share the core values is only a first step in the argument toward grounding ethical judgments. The most evil villain and the most corrupt society will exhibit core human values on an individual basis. Possessing core human values is a sign of being rational, but is not a sufficient condition for being ethical. To adopt *the ethical point of view*, one must respect others and their core values. All things being equal, people do not want to suffer death, pain, disability, interference, deception, loss of resources, or intrusion.

If we take as an ethical guideline to avoid harming others without justification, then the core values give us a set of standards by which to evaluate actions and policies. The core values provide a framework for analysis in computer ethics. By using the core-value framework, some policies for applying computer technology can be judged to be better than others. Let us consider a set of possible policies for the activities of a web browser as an example.

Possible policies for a web site

- 1 Destroy information on the user’s hard disk by leaving a time bomb on the user’s hard disk.
- 2 Remove information from the user’s hard disk without the user’s knowledge.
- 3 Leave a cookie (information about the user’s preferences) on the user’s hard disk without informing the user.
- 4 Leave a cookie on the user’s hard disk and inform the user.
- 5 Do not leave or take any permanent information from the user’s hard disk.
- 6 Give the user the information and ability to accept or decline cookies.

If we respect others and their core values, i.e., take the ethical point of view, then these policies can be ranked at least roughly. Policies 1 and 2 are clearly unacceptable. Nobody contacts a web site wishing or expecting to have his or her hard disk erased or information stolen. The information found on a hard disk is a resource of the user that requires respect and protection. Policy 3 is better than 1 or 2. People may benefit from having their preferences recorded so that the web site can tailor its responses more effectively the next time it is visited. Yet, information is being left on the users’ hard disks without their knowledge. Some deception is involved. Policy 4 is better than 3 in that the user is informed about the activity. Policy 6 is better still in that the user has both the knowledge and the ability to allow or refuse the cookies.

Given these advantages, policy 6 is better than 5, though 5 would be a perfectly acceptable policy in that no harm is being caused to the user.

This analysis of the comparative strengths and weaknesses of these policies could be elaborated, but enough has been said to make several points. People may not agree on exactly how to rank these policies. Some may believe that the theft of information is worse than its destruction and so policy 2 is worse than policy 1. Some may believe that policy 6 creates some risks because of possible misunderstandings about what is being placed on a hard disk and so policy 5 is better than policy 6. But nobody would argue from an ethical point of view that policy 1 or 2 is acceptable. Most would agree that some of the other policies are acceptable and that some are better than others. Moreover, even when there is disagreement about the rankings, the disagreements may have as much to do with factual matters as with value differences. As a matter of fact, does the loss of information cause more damage than its destruction, and, as a matter of fact, do misunderstandings occur about what is or is not left on a hard disk? Apparent value differences may be open to empirical resolution.

The situation is parallel to the evaluation of computer programs. Computer scientists have substantial agreement that some computer programs are terrible and some are very good. There are disagreements about the rankings of some in the middle. Often reasons can be given about why some are better than others. Similarly, some policies for using computers are ethically not acceptable whereas others clearly are. People may have different rankings, but these rankings, assuming an ethical point of view, will have significant positive correlation. Moreover, people can give reasons why some policies are better than others. The core values provide a set of standards by which we can evaluate different policies. They tell us what to look for when making our assessments about the benefits and harms of different policies. They give us the reasons for preferring one policy over another. They suggest ways to modify policies to make them better.

Responsibility, Resolution, and Residue

There are many levels of relativity in value judgments. Some of our values are relative to our being human. If we were angels or creatures from another dimension, our core values might be different. And then, of course, different cultures articulate the core human values differently. And different individuals within a culture may differ in their assessments of values. Indeed, some values of one individual may change over time. I have been arguing that such relativity is compatible with rational discussion of ethical issues and resolution of at least some ethical disputes. We are, after all, human beings, not angels or creatures from another dimension. We share core values. This

provides us with a set of standards with which to assess policies even in situations in which no previous policies exist and with which to assess other value frameworks when disagreements occur.

Ethical responsibility begins by taking the ethical point of view. We must respect others and their core values. If we can avoid policies that result in significant harm to others, that would be a good beginning toward responsible ethical conduct. Some policies are so obviously harmful that they are readily rejected by our core-value standards. Selling computer software which is known to malfunction in a way which is likely to result in death is an obvious example. Other policies easily meet our standards. Building computer interfaces which facilitate use by the disabled is a clear example. And of course, some policies for managing computer technology will be disputed. However, as I have been emphasizing, some of the ethical policies under dispute may be subject to further rational discussion and resolution. The major resolution technique, which I have been emphasizing, is the empirical investigation of the actual consequences of proposed policies. For instance, some people might propose a limitation on free speech on the Internet on the grounds that such freedom would lead to an unstable society or to severe psychological damage of some citizens. Advocates of free speech might appeal to its usefulness in transmitting knowledge and its effectiveness in calling attention to the flaws of government. To some extent these are empirical claims that can be confirmed or disconfirmed, which in turn may suggest compromises and modifications of policies.

Another resolution technique is to assume an impartial position when evaluating policies. Imagine yourself as an outsider not being benefited or harmed by a policy. Is it a fair policy? Is it a policy which you would advocate if you were suddenly placed in a position in which you were affected by the policy? It may be tempting to be the seller of defective software, but nobody wants to be a buyer of defective software. And finally, analogies are sometimes useful in resolving disagreements. If a computing professional would not approve of her stockbroker's withholding information from her about the volatility of stock she is considering buying, it would seem by analogy she should share information with a client about the instability of a computer program which the client is considering purchasing.

All of these techniques for resolution can help form a consensus about acceptable policies. But when the resolution techniques have gone as far as they can, some residue of disagreement may remain. Even in these situations alternative policies may be available which all parties can accept. But, a residue of ethical difference is not to be feared. Disputes occur in every human endeavor and yet progress is made. Computer ethics is no different in this regard. The chief threat to computer ethics is not the possibility that a residue of disagreements about which policies are best will remain after debates on the issues are completed, but a failure to debate the ethical issues

of computing technology at all. If we naively regard the issues of computer ethics as routine or, even worse, as unsolvable, then we are in the greatest danger of being harmed by computer technology. Responsibility requires us to adopt the ethical point of view and to engage in ongoing conceptual analysis and policy formulation and justification with regard to this ever evolving technology. Because the computer revolution now engulfs the entire world, it is crucial that the issues of computer ethics be addressed on a global level. The global village needs to conduct a global conversation about the social and ethical impact of computing and what should be done about it. Fortunately, computing may help us to conduct exactly that conversation.

References

- Bynum T. W. and Rogerson, S. (1996). "Introduction and Overview: Global Information Ethics." *Science and Engineering Ethics*, 2/2: 131–6.
- Johnson, D. G. (1994). *Computer Ethics*, 2nd edn. Prentice Hall, Inc.
- Maner, W. (1996). "Unique Ethical Problems in Information Technology." *Science and Engineering Ethics*, 2/2: 137–54.
- Moor, J. H. (1985). "What is Computer Ethics?" *Metaphilosophy*, 16/4: 266–75.
- Moor, J. H. (1990). "Ethics of Privacy Protection." *Library Trends*, 39/1 & 2: 69–82.

Basic study questions

1. What is the "Routine Ethics" position regarding the nature of computer ethics? Why does Moor believe that this view "undercuts potential for progress" in computer ethics?
2. What is the "Cultural Relativism" position regarding the nature of computer ethics? Why does Moor believe that this view "undercuts potential for progress" in computer ethics? Why does the global nature of the World Wide Web make Cultural Relativism an ineffective approach to computer ethics?
3. What is the "Many/Any Fallacy"? How does the Cultural Relativism position commit this fallacy?
4. Explain the meaning of "logical malleability." Why does this feature of computer technology, according to Moor, make it revolutionary?
5. What does Moor mean by the term "informational enrichment"?
6. How has the concept of money become informationally enriched?
7. How has the concept of warfare become informationally enriched?
8. How has the concept of privacy in the USA become informationally enriched?
9. How has the concept of copyright become informationally enriched?
10. According to Moor, computer ethics has two parts. What are these two components of computer ethics?
11. What, according to Moor, is a policy vacuum? How does computer technology generate policy vacuums?

12. What is a conceptual muddle? How is informational enrichment related to conceptual muddles?
13. What, according to Moor, is a “core value”? List the core values that Moor mentions.
14. According to Moor, to make an ethical judgment one must do more than use core values; one must also “take the ethical point of view.” What is “the ethical point of view”? (See also p. 66 below.)

Questions for further thought

1. What is the difference between a disagreement about facts and a disagreement about values? Give three examples that illustrate the difference.
2. Based upon Moor’s description of the nature of computer ethics, describe a step-by-step procedure for making computer ethics decisions regarding the right thing to do in a given case of computer use. Be sure to take account of the role of core values.
3. Given Moor’s account of the nature of computer ethics, why is computer ethics an especially important branch of applied ethics?