

HUMAN VALUES, ETHICS, AND DESIGN

Batya Friedman and Peter H. Kahn, Jr.
University of Washington

Introduction	1178	Value-Sensitive Design	1186
How Values Become Implicated in Technological Design	1178	Human Values With Ethical Import	1187
The Embodied Position	1178	Human Welfare	1187
The Exogenous Position	1179	Ownership and Property	1188
The Interactional Position	1179	Privacy	1188
Distinguishing Usability From Human Values With Ethical Import	1180	Freedom From Bias	1189
Moral, Philosophical, and Psychological Foundations	1181	Universal Usability	1189
Moral Theory	1181	Trust	1190
Moral Epistemology	1181	Autonomy	1190
Moral Variability and Universality	1182	Informed Consent	1191
Designing for Diversity	1183	Accountability	1191
Approaches to Human Values and Ethics in Design .	1183	Identity	1192
Computer Ethics	1183	Calmness	1192
Social Informatics	1184	Environmental Sustainability	1193
Computer-Supported Cooperative Work	1185	Professional Ethics	1193
Participatory Design	1186	Conclusions	1195
		Acknowledgment	1196
		References	1196

INTRODUCTION

Human values and ethical considerations no longer stand apart from the human-computer interaction (HCI) community—perhaps in some separate field called “computer ethics”—but are fundamentally part of our practice.

This shift reflects, at least in part, the increasing impact and visibility that computer technologies have had on human lives. Computer viruses have destroyed data on millions of machines. Large linked medical databases can, and often do, infringe on individuals' privacy. Accountability becomes a problem when large computer systems malfunction and result in human deaths. The outcome of the 2000 presidential elections in the United States may have hinged on the poor design of the computerized Florida election ballot (the “butterfly ballot”). On and on, the media portray such problems. Presidential reports (PITAC, 1999) ask that such problems be addressed. Corresponding research agendas on human values have been integrated into recent government-funded programs, such as the National Science Foundation's funding initiative on Information Technology Research. Thus, the writing decades ago by cyberneticists like Norbert Wiener (1953/1985) and computer scientists like Joseph Weizenbaum (1972) now seem prescient. They argued that humans fundamentally control technology, and that we must make wise and humane choices about its design and use.

But, if human values, especially those with ethical import—such as rights to privacy and property, physical welfare, informed consent, trust, and accountability, to name but a few—are important, they are no less controversial. What values count? Who decides? Are values relative? Or are some values universal, but expressed differently across culture and context? Does technology even have values? If not, how do values become implicated in design? Also, it is clear that values can conflict. For example, shared online calendars can enhance trust within an organization, but at the expense of individual privacy. Thus, on what basis do some values override others in the design of, say, hardware, algorithms, databases, and interfaces? Finally, how can HCI designers working within a corporate structure and with a mandate to generate revenue bring values and ethics into their designs?

In this chapter, we review how the field has addressed such questions. The second section reviews positions on how values become implicated in technological designs. The third section distinguishes usability from human values with ethical import. The fourth section reviews foundational issues that make ethics so controversial. We also show how scholars have bounded the moral controversies in ways that make it possible for practitioners to shape their work from an ethical stance. The fifth section reviews the major HCI approaches to the study of human values, ethics, and design: Computer Ethics, Social Informatics, Computer-Supported Cooperative Work, Participatory Design, and Value Sensitive Design. We discuss the strengths (and limitations) of each approach by drawing on exemplar projects. The sixth section calls attention to key values relevant for design: Human welfare, ownership and property, privacy, freedom from bias, universal usability, trust, autonomy, informed consent, accountability, identity, calmness, and environmental

sustainability. Each of these values could merit its own chapter; and, indeed, in this *Handbook*, some do. Yet by reviewing them here, however briefly, we highlight their ethical status and thereby suggest that they have a distinctive claim on resources in the design process. The seventh section reviews the special ethical responsibilities that HCI professionals incur by virtue of their professional standing. Finally, in the eighth section, we highlight a few suggestions for moving the field forward.

HOW VALUES BECOME IMPLICATED IN TECHNOLOGICAL DESIGN

Technological innovations implicate human values. Consider, for example, that in the early 1900s, missionaries introduced a technological innovation—steel ax heads—to the Yir Yoront of Australia, a native people. The missionaries did so without regard for traditional restrictions on ownership, and indiscriminately distributed the ax heads to men, women, old people, and young adults alike. In so doing, they altered relationships of dependence among family members and reshaped conceptions of property within the culture (Sharp, 1952/1980). Another example: About four decades ago, snowmobiles were introduced into the Inuit communities of the Arctic and have now largely replaced travel by dog sleds. This technological innovation thereby altered not only patterns of transportation, but symbols of social status, and moved the Inuit toward a dependence on a money economy (Houston, 1995; Pelto, 1973). Or a computer example. Electronic mail rarely displays the sender's status. Is the sender perhaps a curious layperson, system analyst, full professor, journalist, assistant professor, entry-level programmer, senior scientist, or high school student? Who knows until the e-mail is read, and maybe not even then. This design feature (and associated conventions) has thereby played a significant role in allowing electronic communication to cross traditional hierarchical boundaries and to contribute to the restructuring of organizations (Sproull & Kiesler, 1991).

But how exactly do values become implicated in technological designs? Three types of positions have been offered in the literature: Embodied, exogenous, and interactional.

The Embodied Position

The embodied position holds that designers inscribe their own intentions and values into the technology; and once developed and deployed, the resulting technology determines specific kinds of human behavior. Indeed, in the history of science literature (Smith & Marx, 1994), this position is sometimes referred to as technological determinism. To illustrate the idea, consider Latour's (1992) description of the “Berliner lock and key”: A device designed to deal with people who forget to lock the door behind them. The device initially works as a normal lock. You put your key into the lock, turn the key, and the door unlocks. But then to remove the key, it is necessary to push it through the key hole to the other side. Then, by moving to the other side of the door, you rotate the key one more turn,

thereby relocking the door. Thus, when used, this technology “determines” that people will lock their doors.

There is a hard and soft version of the embodied position. In the hard version, it is argued that the very meaning and intentions that designers and builders bring to their task literally become a part of the technology (cf. Appadurai, 1988; Cole, 1991). Many people have trouble with this version because it imputes mental states to things that do not seem to have the capacity to have them. As Smith and Marx (1994) write: “Critics of ‘hard’ determinism question the plausibility of imputing agency to ‘technology’ How can we reasonably think of this abstract, disembodied, quasi-metaphysical entity [technology], or of one of its artifactual stand-ins (e.g., the computer), as the initiator of actions capable of controlling human destiny?” (p. xii). Yet, it is possible that as the field of Artificial Intelligence moves forward—for example, as computational learning systems increasingly mimic human agency—the strong version of the embodied position will gain some currency.

In the more common soft version, it is recognized that objects themselves do not literally embody an intention or value. It is also recognized that designers themselves are shaped by organizational, political, and economic forces; and that because of such forces, a particular technology may never take hold in a society. But, it is argued, if the technology takes hold, then it becomes very difficult for an individual or society to override the values driven by the technology. Think about how the design and deployment of passenger-side front air bags have in a sense held children hostage in the back seat of millions of cars. Or consider that a virus program readily destroys data or grabs CPU cycles (or whatever it is programmed to do). Granted, you could perhaps use the program to teach about self-replicating code, but such activity would in effect be secondary to its primary function. In other words, to execute the computer virus code likely leads to the destruction of data. Similarly, if your goal in online interactions is to keep data secure, you can design technical security mechanisms, such as mandatory encryption, certificates, and anonymizers. Such designs ensure behavior. Thus, according to Akrich (1992), a large part of a designer’s work is that of “inscribing” or embodying his or her vision of “the world in the technical content of the new object” (p. 208).

The Exogenous Position

The exogenous position holds that societal forces—that involve, for example, economics, politics, race, class, gender, and religion—significantly shape how a deployed technology will be used. Consider, for example, the Internet. In many ways, its basic file-sharing functionality remains unchanged from its initial inception. Files are still broken into packets, and packets are routed and then reassembled. Why is it, then, that this underlying technology was primarily used for nearly a quarter of a century by the scientific and educational communities before becoming appropriated by the private sectors for commerce? Once appropriated, why did the development of e-commerce occur so rapidly? According to the exogenous theorists, answers to such questions cannot be found in the supposedly embodied values of the technology itself. Rather:

To understand the origin of a particular kind of technological power, we must first learn about the actors. Who were they? What were their circumstances? . . . Why was this innovation made by these people and not others? Why was it possible at this time and this place rather than another time and place? . . . Instead of treating “technology” per se as the locus of historical agency . . . [these advocates] locate it in a far more various and complex, social, economic, political and cultural matrix (Smith & Marx, 1994, p. xiii).

In turn, the exogenous position has been used and developed in various ways. Bulliet (1994), for example, examines three technologies that emerged in the Islamic world in the fifth century: Block printing, the harnessing of draft animals, and wheeled transport. Each technology failed to have an immediate transforming social and economic impact, not because of its lack of economic advantage, but because of “social filters” that involved race, class, and lifestyle. Perdue (1994) explores the comparative history of agrarian societies (medieval Western European, eighteenth-century Russian, and fourteenth-century Chinese). He argues that studies of technology need contextual accounts that integrate environmental, technological, social, and cultural elements. Scranton (1994) draws on postmodern theory to argue that local contingencies, diversities, disjunctions, multiple oppositions, and contrasting norms should be central to exogenous explanations. Hughes (1994) suggests that younger developing technologies tend to be more open to sociocultural influences, whereas older, more mature technologies tend to be more deterministic. Through many such studies, a recurring view is that technological systems are not value neutral, but invariably favor the interests of people with economic and political power (Noble, 1991; Smith, 1994; Winner, 1986).

The Interactional Position

The interactional position holds that whereas the features or properties that people design into technologies more readily support certain values and hinder others, the technology’s actual use depends on the goals of the people interacting with it. A screwdriver, after all, is well-suited for turning screws, and yet amenable as a poker, pry bar, nail set, cutting device, and tool to dig up weeds. Moreover, through human interaction, technology itself changes over time. On occasion, such changes (as emphasized in the exogenous position) can mean the societal rejection of a technology, or that its acceptance is delayed. But more often it entails an iterative process whereby technologies get invented and then redesigned based on user interactions, which then are reintroduced to users, further interactions occur, and further redesigns implemented. Typical software updates (e.g., of word processors, browsers, and operating systems) epitomize this iterative process.

Two sides of the interactional position have been emphasized in the literature. One side emphasizes the properties designed into the technology (Friedman, 1997a). For example, let us for the moment agree that disabled people in the work place should be able to access technology, just as they should be able to access a public building (Perry, Macken, Scott, & McKinley, 1997). As system designers, we can make the choice

to try to construct a technological infrastructure that disabled people can access. If we do not make this choice, then we single-handedly undermine the human value of universal access.

The other side of the interactional position emphasizes how users use the technology in the context of social-organizational structures. Orlikowski (2000), for example, studied the use of Lotus Development Corporation's *Notes* software by two groups within a large multinational consulting firm: Technologists and consultants. The technologists used *Notes* extensively. They used e-mail, maintained electronic discussions with *Notes* databases, and created their own database designs. Moreover, "supported by the cooperative norms of technical support, the technologists used many of the properties of *Notes* to promote their collective technical work, and to cooperate with each other" (p. 415). In contrast, the consultants used *Notes* minimally, sometimes even begrudgingly. Orlikowski found that what she calls such "technology-in-practice" was enacted for at least three different reasons. First, some consultants had doubts about the value of *Notes* for their performance. Second, in contrast to the technologists, the consultants were under a time-based billing structure. "Because many consultants did not see using *Notes* as an activity that could be billed to clients, they were unwilling to spend time learning or using it" (p. 416). Third, consultants feared that the collaborative properties of *Notes* would threaten their status within the company. Thus, Orlikowski proposes "a view of technology structures, not as embodied in given technological artifacts, but as enacted by the recurrent social practices of a community of users" (p. 421).

Regardless of emphasis, from the interactional position it should be clear that design and social context matter, dialectically. Moreover, users are not always powerless when faced with unwelcomed value-oriented features of a technology.

DISTINGUISHING USABILITY FROM HUMAN VALUES WITH ETHICAL IMPORT

The language and conceptualizations within HCI currently provide solid means by which to pursue issues of usability (Adler & Winograd, 1992; Nielsen, 1993; Norman, 1988). Usability refers to characteristics of a system that make it work in a functional sense, including that it is easy to use, easy to learn, consistent, and recovers easily from errors.

Some HCI designers, however, have a tendency to conflate usability with human values with ethical import. (In the next section, we discuss what exactly is meant by the term moral. Here, it can be understood to involve issues of fairness, justice, human welfare, and virtue.) This conflation arises insofar as usability is itself a human value, although not always a moral one. But when it is, both can be addressed by the same design. For example, systems that can be modified by users to meet the needs of specific individuals or organizations can both (a) enhance usability and (b) help users to realize their goals and intentions: The moral value of autonomy. Other times, however, usability can conflict with human values with ethical import. Nielsen (1993), for example, asks us to imagine a computer system that checks for fraudulent applications of people who are applying for unemployment benefits. Specifically, the system

asks applicants numerous personal questions, and then checks for inconsistencies in their responses. Nielsen's point is that even if the system receives high usability scores, some people may not find the system socially acceptable, based on the moral value of privacy.

In terms of a general framework, four pairwise relationships exist between usability and human values with ethical import.

1. *A design is good for usability and independently good for human values with ethical import.* This relationship is exemplified previously, where adaptable systems can also promote user autonomy. Another example involves browser designs that offer users more efficient cookie management than currently offered in Netscape or Internet Explorer and also promote values of informed consent and privacy (Friedman, Howe, & Felten, 2001).

2. *A design is good for usability, but at the expense of human values with ethical import.* This relationship is exemplified in the previous example offered by Nielsen, where a highly usable system undermines the value of privacy. Tang (1997) provides another example by means of a case study of a team designing a workstation. At one point, the team was trying to decide how to power a microphone, and finally decided to power it directly from the workstation. Consequently they eliminated a separate hardware on/off switch on the microphone. From a usability perspective, Tang points to benefits of this design. Users, for example, no longer had to remember to turn the microphone off when it was not in use (to conserve battery power), or be inconvenienced by replacing a dead battery. However, despite these usability advantages, Tang found that some users thought the design—by not allowing them direct control over the microphone—undermined their privacy and security.

3. *A design is good for human values with ethical import, but at the expense of usability.* This relationship is exemplified by choosing a browser preference with the setting "accept or decline each cookie individually." The values of privacy and informed consent are well supported; but, for most people, the nuisance factor is too high. Another example (the flip side of an earlier one) arises when a system is purposefully made adaptable to promote autonomy, but in the process the system becomes unwieldy and difficult to use.

4. *A design good for usability is necessary to support human values with ethical import.* This relationship is exemplified by security systems that are not so cumbersome to use that either critical features of the security system are disabled or individuals invent workarounds that compromise the security. For example, if a security system requires 3 passwords and 10 steps to login, a user's workaround might involve taping a yellow sticky to the side of the computer screen with the passwords. Other instances of this relationship arise when morality requires that all individuals in a specified group be able to successfully use the system. For example, to have a fair national election using a computerized voting system, all citizens of voting age must be able to use the system.

HCI professionals are often responsible for usability. Accordingly, it is important that we be aware of the possible relationships between usability and human values with ethical import.

At times, the two support one another. Other times we need to give ground on usability to promote human values with ethical import, or, conversely, give ground on human values with ethical import to promote usability. Such optimizations require judicious decisions, carefully weighing and coordinating the advantages of each.

MORAL, PHILOSOPHICAL, AND PSYCHOLOGICAL FOUNDATIONS

HCI professionals—like people in many fields—may sometimes wonder if morality is too controversial to be integrated in a principled way into their work. After all, who is to say for another person what is right or wrong, or good or bad? How do we make sense of the seemingly different moralities among people? Is not morality relative to person or at least culture?

Such questions form the backdrop to a wide range of issues that moral philosophers and psychologists have pursued. In this section, we review relevant literature to help bound moral controversies such that HCI researchers and practitioners can, with legitimate grounding, move forward proactively in shaping their work from a moral stance.

Moral Theory. One common starting point in moral discussions is to raise the question, “What do you mean by morality?” Or, “How do you define it?” In response, moral philosophers have offered what can be viewed as three overarching approaches toward developing moral theory: Consequentialist, deontological, and virtue-based. Briefly stated, consequentialist theories maintain that a moral agent must always act to produce the best available outcomes overall (see Scheffler, 1982, for an analysis). Utilitarianism is the most common form of consequentialism, wherein a moral agent should act to bring about the greatest amount of utility (e.g., happiness) for the greatest number of people. In contrast, deontological theories maintain that there are some actions that a moral agent is forbidden to do, or, in turn, must do, regardless of consequences (Dworkin, 1978; Gewirth, 1978; Kant, 1785/1964; Rawls, 1971). For example, a deontologist might maintain that it is immoral to torture an innocent child even if such an act would bring about great good (e.g., to prevent a bombing). Along similar lines, it can be argued that vendors who collect and sell information about individuals’ web-browsing activity without individuals’ knowledge and consent violate individuals’ right to privacy. Indeed, most any argument for a basic right makes a deontological move.

Both consequentialist and deontological theories are centrally concerned with answering the fundamental question, “What ought I to do?” In turn, answers are often viewed to be morally obligatory, meaning that the action is required of every moral agent “regardless of whether he wants to accept them or their results, and regardless also of the requirements of any other institutions such as law or etiquette” (Gewirth, 1978, p. 1). Moreover, such obligatory prescriptions are often framed in the negative (as in “Thou shall not steal”), such that it is possible to fulfill the prescription (e.g., just don’t steal). In turn, virtue-based theories are centrally concerned with answering the fundamental question, “What sort of person ought

I to be?,” in which the focus is on long-term character traits and personality (see Louden, 1984). This tradition dates back to Aristotle’s delineation in *Nicomachean Ethics* of the ethical virtues (e.g., courage, temperance, friendship, wisdom, and justice), and developed, for instance, by Foot (1978), MacIntyre (1984), and Campbell and Christopher (1996). In virtue ethics, the prescription is often viewed to be morally discretionary, meaning that whereas an action is not morally required, it is conceived of as morally worthy and admirable based on considerations of welfare and virtue (Kahn, 1992). Moreover, such discretionary prescriptions are often framed in the positive (as in “be charitable”), such that it is not possible to completely satisfy the declaration. After all, one cannot continuously practice charity without soon becoming destitute oneself (Fishkin, 1982).

With these core distinctions in mind—between a theory of the right (both consequentialism and deontology) that is largely viewed as obligatory, and a theory of the good (virtue theory) that is largely viewed to be discretionary—it is possible to characterize a wide range of HCI designs. For example, consider Mattel Corporation’s “Barbie doll” Web site that was up and running during the latest census collection in the United States. On the Web site, Barbie posed as a census taker and asked children to provide information about their family. Presumably, Mattel sought to use the resulting information to market their products more effectively. Is this action moral? Most people would presumably say no. But why? From a deontological position, such a deceptive web design violates a moral obligation (e.g., the obligation not to intentionally deceive others, let alone to deceive children for a corporation’s material gain), and should not have been designed and deployed.

Other human values articulated in HCI fit within a virtue orientation. For example, SeniorNet is an organization that brings seniors together via computer networks. Mynatt, Adler, Ito, Linde, and O’Day (1999) found that “SeniorNetters repeatedly commented on the warmth and friendliness of the community as something that differentiated SeniorNet from other net communities, and as a reason for their participation and comfort with the community” (p. 232). Are values such as warmth and friendliness moral values? For a consequentialist and even more so for a deontologist the answer is probably no; a person, after all, is presumably not under a moral obligation to engage in “warm-hearted” actions. But for a virtue ethicist the answer is fundamentally yes insofar as such values are conceived to be central to moral personhood. Or consider Web sites (such as Alfie.com) that offer children a menu of computer games, like “pinball” or “miniature golf,” that children can play. Most of these games do not promote such virtues as friendship, caring, or compassion. Must they? Presumably, even virtue theorists would say that not every game has to. Rather, promoting such virtues is usually conceived as discretionary—perhaps morally praiseworthy if done, and perhaps contributing to an account of human flourishing—but not morally obligatory.

Moral Epistemology. This discussion raises the question, “Who is to say what’s moral?” or “How do we know?” Such questions move us into the field of moral epistemology: The study of the limits and validity of moral knowledge. Often at

stake is whether it is even possible for a moral statement to be objectively true or false, and for a moral value to be objectively right or wrong, or good or bad. A wide variety of positions have been taken. For instance, some believe that moral knowledge corresponds to or approaches a correspondence with a moral reality that exists independent of human means of knowing (Boyd, 1988; Sturgeon, 1988). Others believe moral knowledge can be objectively grounded by constructing rational principles that strive for coherence and consistency while building on the common ground and specific circumstances of a society (Dworkin, 1978). Others believe that the only thing that can be said of moral knowledge is that it can be true subjectively for an individual or culture depending on that individual's or group of individuals' desires, preferences, and goals (Mackie, 1977; Rorty, 1982). Finally, others believe that any moral knowledge is unattainable, even in a weak sense (the full skeptic's position; see Nagel, 1986, for a characterization).

The skeptic's position—that moral knowledge is unattainable—reflects the morally relativistic position described previously: That no one can say what is really right or wrong and so ethics and values become, at best, a personal choice in terms of HCI design. But consider: No one can prove that at this moment you are not really just a brain in a vat being stimulated by electrical impulses to think that you are reading this chapter. Still, you have compelling reasons to think that proposition false. Similarly, with morality. Although the skeptic's position cannot be proved conclusively false, moral philosophers have provided compelling reasons not to believe it (Nagel, 1986; Williams, 1985).

Moral Variability and Universality. Anthropologists often document moral differences between cultures. For example, from some anthropological accounts, we learn that devout Hindus believe that it is immoral for a widow to eat fish, or for a menstruating woman to sleep in the same bed with her husband (Shweder, Mahapatra, & Miller, 1987). Other accounts document that members of the Yanomamo tribe of Brazil at times practice infanticide and that the women are "occasionally beaten, shot with barbed arrows, chopped with machetes or axes, and burned with firebrands" (Hatch, 1983, p. 91). Per our discussion, there are three ideas to understand about such examples.

1. *Variability in human practice does not prove or disprove the moral skeptic's position.* Imagine going to a culture where the people there did not believe in logical transitivity. (If $A = B$ and $B = C$, then $A = C$.) That finding would be interesting psychologically and culturally, but presumably has no bearing on whether logical transitivity is true or false. It is or it is not. If it is true, then people who think otherwise are simply wrong. So, too, with the moral life. People can believe that a certain act (such as shooting women with barbed arrows) is moral; but documenting such a belief does not make it so.

2. *When moral differences occur between peoples, it is not necessarily the case that the practices are believed legitimate by the victims.* For example, in Hatch's (1983) report on the Yanomamo, he also reports that the women did not appear to enjoy such physically abusive treatment, and were seen running

in apparent fear from such assaults. Psychological data of a similar kind can be found in a study by Wainryb (1995) on the Druze population in Israel. The Druze largely live in segregated villages, are of Islamic religious orientation, and organized socially around patriarchal relationships. The father, as well as brothers, uncles, and other male relatives, and eventually a woman's husband, exercise considerable authority over women and girls in the family, and restrict their activities to a large degree. However, when these women were interviewed, a majority of them (78%) unequivocally stated that the husband's or father's demands and restrictions were unfair. Thus, Yanomamo and Druze women—like many women in Western societies—are often unwilling victims within what they themselves perceive to be an uncaring or unjust society. In such situations, it is less the case that societies differ on moral ground, and more that some societies (Western societies included) are involved explicitly in immoral practices.

3. *Moral variability may be much less pervasive than many people suppose.* Reconsider, for example, the Shweder et al. (1987) report that devout Hindus believe it is immoral for a widow to eat fish, or for a menstruating woman to sleep in the same bed as her husband (but two of many dozens of their examples). At first glance, for Western eyes at least, such moral beliefs do seem different. However, when Shweder et al.'s data were reanalyzed by Turiel, Killen, and Helwig (1987), they showed bases for not only difference, but also moral commonality. In their reanalysis, for example, Turiel et al. found that devout Hindus believed that harmful consequences would follow from a widow who ate fish (the act would offend her husband's spirit and cause the widow to suffer greatly), and from a menstruating woman who sleeps in the same bed with her husband (the menstrual blood is believed poisonous and can hurt the husband). Although such beliefs, themselves, differ from those in Western culture, the underlying concern for the welfare of others is congruent with it. More generally, Turiel et al.'s claim is that when researchers differentiate informational and metaphysical assumptions about the world from moral judgments based on those assumptions, then the moral life often takes on a greater universal cast (cf. Turiel, 1998, in press).

When analyzing moral variability, conceptualizations of morality that entail abstract characterizations of justice and welfare tend to highlight moral universals, whereas definitions that entail specific behaviors or rigid moral rules tend to highlight moral cross-cultural variation. Typically, theorists who strive to uncover moral universals believe they are wrestling with the essence of morality, with its deepest and most meaningful attributes. In contrast, theorists who strive for characterizing moral variation argue that, by the time you have a common moral feature that cuts across cultures, you have so disembodied the idea into an abstract form that it loses virtually all meaning and utility. For instance, in the example of devout Hindus who believe that by eating fish a widow hurts her dead husband's spirit, is the interesting moral phenomenon that Hindus, like ourselves, are concerned with not causing others harm? Or is the interesting moral phenomenon that Hindus believe in spirits that can be harmed by earthly activity?

In our view, both questions have merit, and a middle ground provides a more sensible and powerful approach for the HCI

community: One that allows for an analysis of universal moral values, as well as allowing for these values to play out differently in a particular culture at a particular point in time (Friedman, 1997b; Kahn, 1991, 1999; Kahn & Lourenço, 1999).

Designing for Diversity. In this chapter, we will continue to draw on the preceding analysis to help us review the HCI literature, and to offer morally principled design methods that respect culture and context. As a case in point, imagine you are designing a computerized voting booth. At what height would you place the electronic ballot? A reasonable answer might go something like: "It depends—tell me, how tall are the people who vote? Moreover, at a later time will the voting booth be used by other people of different height? After all, a voting booth designed only for players of the National Basketball Association will disenfranchise most voters in Japan." In other words, the universal value is to enfranchise all voters, but the specific mechanism by which to do so may need to be adaptable to specific contexts and cultures. The general principle then is that designs need to be robust enough to substantiate the value under consideration and yet adaptable enough so that different cultures (or subcultures) can use the designs in their own way.

Many problems occur when this principle is ignored or when unanticipated users of a system emerge. For example, as described by Friedman and Nissenbaum (1996), since the early 1970s, the computerized National Resident Medical Match Program has placed most medical students in their first jobs. In the system's original design, it was assumed that only one individual in a family would be looking for a residency. At the time, such an assumption was perhaps not out of line, because there were few women residents. But as women have increasingly made their way into the medical profession, marriages between residents have become more frequent, and bias against placing couples emerged. Another example involved a dog icon that was used to indicate the printing orientation of a printer, landscape or vertical. When the printer was shipped to Islamic countries, the vendor discovered that people in such countries often considered dogs as offensive animals. Or consider the case of competitive educational software that was exported to Micro Asia, with dismal results because the value of competition promoted by the software clashed too strongly with the culture's emphasis on cooperation. Or consider that data protection laws and policies differ across national boundaries (Agre & Rotenberg, 1998; Rotenberg, 2000). For example, Bennett (1998) describes how most of the European countries have applied the same data-protection policies to both the public and private sector. However, the United States, Canada, Australia, and Japan have preferred "to regulate only the public sector's practices and to leave the private sector governed by a few sectoral laws and voluntary codes of practice" (p. 100). Thus, lenient designs from the private sector of these latter countries need to be adaptable to transfer readily to European countries.

Granted, building value adaptability into systems requires additional time and financial resources. Also, from a user's perspective, additional options add further complexity and challenges. Thus value adaptability is not always the best way to go. But it often is, especially when we can anticipate that similar values will play out in different ways for different users. In such cases,

not only are a larger number of human lives enhanced from an ethical standpoint, but also from an economic standpoint such systems increase market share and generate profits.

APPROACHES TO HUMAN VALUES AND ETHICS IN DESIGN

The computing field has addressed issues of human values and ethics by means of a handful of approaches. To some extent, these approaches overlap with one another. For example, sociotechnical analyses, which are central to Social Informatics, often form the front end of efforts in Participatory Design, and are incorporated into the empirical investigations of Value-Sensitive Design. Both Computer-Supported Cooperative Work and Participatory Design share particular interests in collaboration in the work place. But, that said, the approaches differ significantly when considering how each integrates their respective positions on moral epistemology, methods, and contexts studied.

In this section, our goal is not to review each approach comprehensively (which is beyond the scope of this chapter). Rather, we seek to show how each approach contributes to our understanding of how to integrate human values and ethics in design.

Computer Ethics

When applied moral philosophers brought their talents and energies to bear on understanding the impact of computer technologies on social life, the field of Computer Ethics was born (Bynum, 1985; Johnson, 1985; Moor, 1985).

In their resulting examinations, computer ethicists have embraced two complimentary goals. One goal has been to utilize existing moral theory to bring clarity to issues at hand, and—at appropriate times—to proscribe norms of behavior. For example, in the computer science literature, the term trust has often been used synonymously with security (Schneider, 1999). Yet, drawing on ethical theory, Nissenbaum (1999) has shown that these two terms need to be distinguished. For example, if, as HCI professionals, our goal is to create a place where people feel safe in their online interactions, we can move in two design directions. We can move forward with technical solutions that involve security features like locks, keys, passwords, and encryption. Or we can understand how trusting relationships are created and fostered, and design them into our online systems. Each direction leads to a very different online environment. The strength of this type of philosophical contribution is that it helps translate moral abstractions into crisp working conceptualizations that HCI professionals can use.

A second goal of computer ethicists builds on the innovations of the technology itself. For the technologies have not only generated new entities (computer programs, the Internet, Web pages), but have also enlarged the scale of activities (data mining), increased the power and pervasiveness of its effects (ease of communication), and often become invisible to human purview (using computers for surveillance). Thus, computer ethicists have been interested in understanding how such

innovations extend the boundaries of traditional ethical concepts. Moor (1985), for example, examines the ways in which the invisibility of computers affect human lives. *Invisible abuse* is the intentional use of the invisible operations of a computer to engage in unethical conduct, such as the invasion of the property and privacy of others. *Invisible programming values* are those values that are embedded in a computer program, such as bias. *Invisible complex calculations* refer to calculations that are too complex for human inspection and understanding. In all three situations, Moor argues that computer technologies raise special ethical issues, different from other forms of technology, and that the task of computer ethics is to fill in what he refers to as the resulting “conceptual vacuum” and “policy vacuum” (p. 266).

One theoretical debate in this literature is whether or not the technological innovations fundamentally challenge ethical theory (Johnson & Miller, 1997; Spinello & Tavani, 2001). In other words, is it the case that the innovations are so qualitatively different and far-reaching from past ones that traditional ethical theories have to be, if not abandoned, at least significantly revised? Or is it the case that the technology simply offers a new domain within which traditional ethical theory works? Toward understanding this debate, a helpful analogy can be made to literature, wherein there exist the categories (genres) of fiction and nonfiction. These categories seem clear enough until we encounter a new form of writing that blurs the boundaries, such as historical novels. Does that new form undermine the traditional categories? Philosophers like Searle (1983) argue no: That the categories of fiction and nonfiction fundamentally remain, and what the new literary form shows is that writing can embody aspects of both categories. Similarly, it can be argued that whereas the technology can blur traditional ethical boundaries, and demand further refinements and clarity of moral theory, that ethics itself has not changed, nor its fundamental precepts.

Regardless of how one views this debate, it is clear that the field of Computer Ethics advances our understanding of key values that lie at the intersection of computer technology and human lives. However, there are some limitations to the field's contributions to HCI. For one, Computer Ethics often remains too divorced from technical implementations. How exactly, for example, can HCI professionals build interfaces that enhance trust within a community of users? Or how exactly do we address the value problems that arise through invisible computing? In addition, Computer Ethics has focused too often on a single value at a time. Yet, as HCI professionals, we commonly wrestle with design trade-offs between competing moral values. Collaborations between computer ethicists and HCI professionals would be a fruitful way to address these limitations.

Social Informatics

In the second section, we reviewed embodied, exogenous, and interactional positions on how values become implicated in technology. Recall that the embodied position holds that designers inscribe their own intentions and values into the technology; and once developed and deployed, the resulting technology is said to determine specific kinds of human behavior. As also

noted, many researchers find such a position highly untenable, and instead emphasize the social context in which information systems are used. Recall, for example, Orlikowski's (2000) study, reviewed earlier, where she investigated the effects when Lotus Development Corporation's *Notes* was introduced into a large corporation. She found that, at least in part, the incentive systems in the corporation did more to influence how, and whether, *Notes* was used than the mere capability of the software itself.

Over the last 25 years, this emphasis on the social context of information technologies has been the subject of systematic research (Attewell, 1987; Borgman, 2000; Iacono & Kling, 1987; Kiesler, 1997; King, 1983; Kling, 1980; Orlikowski, 1993; Poltrock & Grudin, 1994). The research has appeared under many labels, including social analysis of computing, social impacts of computing, information systems, sociotechnical systems, and behavioral information systems. In more recent years, this overarching enterprise has begun to coalesce within a new field called Social Informatics (Kling, Rosenbaum, & Hert, 1998; Kling & Star, 1998; Sawyer & Rosenbaum, 2000). As defined by participants at the 1997 National Science Foundation-sponsored workshop on this topic, Social Informatics is the interdisciplinary study of the design, uses, and consequences of information technologies that takes into account their interaction with institutional and cultural contexts.

To illustrate the value of Social Informatics, Kling (1999) contrasts the design and functioning of two electronic journals: *Electronic Transactions of Artificial Intelligence* (ETAI) and *Electronic Journal of Cognitive and Brain Sciences* (EJCBS). Both journals envision attracting high-quality papers. Both journals also have implemented technology that works effectively. However, Kling argues that differences in their sociotechnical designs lead ETAI to prosper and EJCBS to wane. For example, articles submitted to ETAI are reviewed in a two-phase process. In phase 1, the article is open to public online discussion for a period of 3 months. Based on the resulting discussions, authors have an opportunity to revise their papers. In phase 2, the article undergoes a quick confidential peer review. In contrast, articles submitted to EJCBS are evaluated by their general online readership. Articles that receive a minimum score are then transferred to an archive of accepted papers. EJCBS has been designed as much as possible as an autonomous system that would run on its own after it was programmed. It removes editorial attention from the publishing process, and instead relies on a readers' plebiscite. But such a design, according to Kling, misconstrues the social context of successful academic journal publishing, one that requires a lively group of authors and readers, and attention from senior scientists in the field.

Work in Social Informatics has been successful in providing sociotechnical analyses of deployed technologies. Yet, in terms of its contributions to HCI, Johnson (2001), for example, writes: “One aspect that still confounds me is how to reconcile the basic premise of social informatics—that it is critical to gain knowledge of the social practices and values of the intended users—with the basic work of system developers. How, if at all, can programmers practice and apply social informatics?” (p. 18). Granted, Kling (1999) recognizes the importance, for example, of “workplace ethnography, focus groups, user participation in design teams, and participatory design strategies.” But then he

says "to discuss it [these methods] in detail here would lead us away from our focus on the structural elements of a socio-technical analysis." Thus, it appears that Social Informatics can move in at least two directions. One direction leads to developing the sociotechnical analyses, and viewing this work as complementing work in design (and other areas, such as Computer Ethics). Another direction leads to an expansion of Social Informatics such that it fundamentally embraces design (and other areas) into its theoretical framework.

Computer-Supported Cooperative Work

Although social informatics has emphasized the sociotechnical analyses of deployed technologies, the field of Computer-Supported Cooperative Work (CSCW) has, for some time, focused on the design of new technologies to help people collaborate effectively in the workplace (Galegher, Kraut, & Egido, 1990; Grief, 1988; Grudin, 1988). Groupware is the name often used for software that seeks to facilitate CSCW goals. The history of CSCW goes back to early work in various research laboratories, like Xerox PARC, EuroPARC, Bell Lab, and IBM. There, computer professionals, working within relatively small groups themselves, and sometimes remotely, sought to improve their collaborations.

Typically, then, the values considered in CSCW designs have been closely tied to group activities and workplace issues. Cooperation has been, of course, an overarching value. But, in addition, the field has paid attention to such values as privacy, autonomy, ownership, commitment, security, and trust. Isaacs, Tang, and Morris (1996), for example, designed a system to support informal interactions in the workplace, using the piazza (the plaza) as the metaphor. The application provided means for workers to opt out of piazza interactions, thus protecting workers' privacy and autonomy. Olson and Teasley (1996) report on the planning, implementation, and use of groupware tools over the course of a year in a real group with remote members. One of their key findings was that "social responsibility and commitment appeared diminished or missing when people did not meet face-to-face" (p. 425). Dewan and Shen (1998) discuss an access-control model that accounts for joint ownership of shared objects, different ownership rights for different types of users, and the delegation of access rights (security). Hudson and Smith (1996) sought after methods that allow workers to share video information about themselves to their colleagues while providing protections for privacy. One solution was to shadow images on the video screen so that colleagues can tell that you are in your office, but not what you are doing. Fuchs (1999) developed a notification service for awareness information that also addresses potential conflicts between awareness and privacy. Van House, Butler, and Schiff (1998) examined how in a workplace with environmental planning data sets that trust is created and assessed, and "how changes in technology interact with those practices of trust" (p. 335).

CSCW has traditionally focused on the workplace, as its name implies. Yet if the recent Conference Proceedings of CSCW are any indication, then the field is quickly expanding to include non-workplace settings. For example, in the CSCW 1996

Proceedings, about 6% of the papers involved a non-workplace setting. In 1998, it was 13%. In 2000, it was 31%. The CSCW 2000 conference co-organizers note: "the fact that these [work] place technologies are available not only in working settings but also in homes means that the focus of our attention [in the CSCW community] has broadened to encompass a much wider range of activities that we could have imagined when the CSCW conference series began in 1986" (Dourish & Kiesler, 2000, p. v).

Correspondingly, the range of values that the CSCW field has begun to analyze has also started to expand. For example, Mynatt et al. (1999) delineate the values of safety, civility, warmth, and friendship that are fostered by SeniorNet, an organization that brings seniors together through computer networking technologies. O'Neill and Gomez (1998) describe a project that links middle and high school science students with working scientists, as mentors for the students' science projects. The researchers "illustrate the unique dynamics of these relationships, consider their technical and social demands, and discuss the potential for CSCW systems to help sustain long-term help relationships by better accommodating their needs" (p. 325). McCarthy and Anagnost (1998) explore the social ramification of a group preference agent for music in a shared environment, a fitness center. By technical means, they thus seek to democratize the music selection process. Palen, Salzman, and Youngs (2000) tracked 19 new cell mobile phone users for 6 weeks. In their discussion, they call attention to issues of unfairness that arose in the context of use.

It remains unclear how much more the field of CSCW can expand to include non-workplace settings before its very name (that has "work" in its title) becomes an historical artifact rather than a description of its current activity. But that is a quibble with nomenclature. The direction of the field seems intellectually vibrant and exciting. Thus, as CSCW continues to expand into a broader range of human activity, it will increasingly encounter and thereby take hold of a broader range of human values with ethical import.

A substantive question thus emerges: How should we understand the epistemic standing of moral values within a CSCW framework? One way of currently reading the field is that moral values are simply relative to culture (and to the "culture" of any particular work group). Harrison and Dourish (1996), for example, argue that "privacy is relative, not a set of psychological primitives" (p. 71). As such, these values should be respected only if the workers themselves think the values are important. Yet, as the field continues to broaden, the tensions embedded in this perspective will continue to become cause of concern. Greenbaum (1996), for example, argues that CSCW needs to consider the political dimensions of labor and not assume that work is cooperative. Suchman (1994), too, emphasizes the power hierarchies within social organizational groups. In such situations—which presumably arise within every organization, at least to some extent—then values often will conflict. Perhaps management (which holds the power) seeks efficiency over privacy, whereas workers seek the converse. What does a designer do? Thus, once CSCW analyses move beyond largely homogeneous groups and into organizational structures, then potentially a principled position on the moral standing of human values will be required.

Participatory Design

In Norway, in the early 1970s, there was a general consensus that computer systems should not deskill workers, but enhance skill, protect crafts, and foster meaningful work. At that time, strong labor unions also helped enact into law a national *codetermination agreement*. This agreement entitled workers along with management to determine which technologies are introduced into the workplace (Kuhn & Winograd, 1996). Thus, emerging from this social structure was a new approach to system design—Participatory Design—that fundamentally sought to integrate workers' knowledge and a sense of work practice into the system design process (Bjerknes & Bratteteig, 1995; Bodker, 1990; Ehn, 1989; Greenbaum & Kyng, 1991; Kyng & Mathiassen, 1997).

In light of value considerations, then, Participatory Design has embedded within it a commitment to democratization of the workplace and human welfare. It also has what can be viewed as virtue-based moral commitments insofar as it seeks to account for meaningful activity in everyday lives.

At least five important methods have emerged from, or have been elaborated by, the field of Participatory Design. (1) *Identifying Stakeholders*—Toward achieving designs that work, it is often necessary to identify the people they directly and indirectly effect (cf. Korpela et al., 1998). (2) *Workplace Ethnography*—Ethnographies document the practices and beliefs of a group from the group's perspective. Methods include analyzing documents and artifacts in the group's environment, participant observation, field observations, surveys, and formal and informal interviews (cf. Kensing, Simonsen, & Bodker, 1998; Mackay & Fayard, 1999). (3) *Future Workshops*—A future workshop is a method to uncover common problems in the workplace and to solve them. This method is divided into three phases. As described by Kensing and Masden (1991), "the Critique phase is designed to draw out specific issues about current work practice; the Fantasy phase allows participants the freedom to imagine 'what if' the workplace could be different; and the Implementation phase focuses on what resources would be needed to make realistic changes" (p. 157). (4) *User Participation in Design Teams*—Four ways users participate in design teams can be characterized (Kuhn & Winograd, 1996): Directness of interaction with designers; length of involvement in the design process; scope of the participation; and the degree of control over the design decisions. User participation has been central to Participatory Design projects conducted in Europe and North America. In addition, Korpela et al. (1998) write that their 7-year experience in Nigeria suggests that user participation is also a must in developing countries. (5) *Mock-ups and Prototypes*—Both mock-ups and prototypes create physical representations of technological designs. A mock-up looks roughly like the artifact it represents, but completely lacks the artifact's functionality. Mock-ups usually occur very early in the design process. As noted by Ehn and Kyng (1991), mock-ups encourage "hands-on" experience, and are understandable to the end-user, cheap to build, and fun to work with. In turn, prototypes incrementally embed functionality into the artifact through successive iterations. Both methods help end-users envision the potential for the proposed technology

and the resulting changes in work practice (Greenbaum & Kyng, 1991).

Some HCI practice in the United States has followed closely in the Scandinavian style of Participatory Design. But what has been embraced even more commonly is *Pragmatic Participatory Design*, a term we use to refer to the above rich constellation of methods and design techniques, but largely stripped of their moral commitments.

Two reasons may help account for why the HCI community in the United States has been resistant to embrace the moral commitments of Participatory Design, while embracing many of its methods. One reason is that although the United States is politically committed to the value of participatory democracy, its capitalistic business culture is not. Thus, the moral values that lie at the core of Participatory Design (participation, democracy, and moral personhood) run counter to values in the U.S. workplace. A second reason can be viewed in light of the cultural homogeneity in the Scandinavian countries. Compared with the United States, for example, these countries historically have been more homogeneous in terms of race, ethnicity, and religion, and thus have encountered fewer opportunities for corresponding prejudices and hostilities. Thus, when applied in more diverse contexts, Participatory Design has too little to say when divisive constituencies argue on the basis of narrowly conceived self interests and hostile prejudices. After all, at least in theory, Participatory Design values each participant's voice, even those that appear uncaring and unjust. This problem has been of concern within the field (Gross, Parker, & Elliott, 1997; Muller, 1997).

Value-Sensitive Design

Given the limitations of the other approaches in integrating ethics and sociotechnical analyses with actual design, another approach has recently emerged called Value-Sensitive Design. This approach seeks to design technology that accounts for human values in a principled and comprehensive manner throughout the design process (Friedman, 1997a). Value-Sensitive Design is primarily concerned with values that center on human well-being, human dignity, justice, welfare, and human rights. This approach is principled in that it maintains that such values have moral epistemic standing independent of whether a particular person or group upholds such values. At the same time, Value-Sensitive Design maintains that how such values play out in a particular culture at a particular point in time can vary, sometimes considerably.

Value-Sensitive Design articulates an interactional position for how values become implicated in technological designs. From this position, Friedman and Nissenbaum (1996) have analyzed bias in computer systems. Cranor and Resnick (2000) have analyzed anonymity in ecommerce. Agre and Mailloux (1997) have analyzed privacy and computerized vehicle-highway systems. Thomas (1997) has analyzed universal access within a communications company. Ackerman and Cranor (1999) have analyzed interface components to safeguard users' privacy on the Internet. Abowd and Jacobs (2001) have called attention to how the design of advanced sensing technologies within and

outside the home can impinge on the individual's right (as established by the Fourth Amendment) to be protected against illegal search and seizure by the government. Shneiderman and Rose (1997) have proposed social impact statements for information systems.

Methodologically, at the core of Value-Sensitive Design lies an iterative process that integrates conceptual, empirical, and technical investigations (Friedman & Kahn, in press). *Conceptual investigations* involve philosophically informed analyses of the central constructs and issues under investigation. Questions include: How are values supported or diminished by particular technological designs? Who is affected? How should we engage in trade-offs among competing values in the design, implementation, and use of information systems? *Empirical investigations* involve social-scientific research on the understandings, contexts, and experiences of the people affected by the technological designs. *Technical investigations* involve analyzing current technical mechanisms and designs to assess how well they support particular values, and, conversely, identifying values, and then identifying and/or developing technical mechanisms and designs that can support those values. As mentioned, these investigations are iterative and integrative. For example, results from the empirical investigations may reveal values initially overlooked in the conceptual investigations, or help to prioritize competing values in the design trade-offs between technical mechanisms and values considerations.

To illustrate this methodology, consider a current project by Friedman, Felten, and their colleagues as they have sought to understand how to design web-based interactions to respect informed consent (Friedman, Howe, & Felten, 2001; Friedman, Millett, & Felten, 2000; Millett, Friedman, & Felten, 2001). They began their project with a conceptual investigation of informed consent itself. What is it? How can it be garnered in diverse online interactions in general, and web-based interactions in particular? To validate and refine their resulting conceptual analysis, and initiate their technical investigation, they conducted a retrospective analysis of existing technology. Namely, they examined how the cookie and web-browser technology embedded in Netscape Navigator and Internet Explorer changed—with respect to informed consent—over a 5-year period, beginning in 1995. (These results are summarized in *Informed Consent*.) Then, the design work began. Their design improvements are being implemented in the Mozilla browser (the open-source code for Netscape Navigator) and undergoing empirical investigations (usability studies and formative evaluation), which will then reshape the initial technical and conceptual work.

The National Science Foundation recently sponsored two workshops to help shape a research agenda for Value-Sensitive Design. The recommendations from the workshops' final reports (Friedman, 1999; Friedman & Borning, 2001) have focused on the need for (1) *theoretical and conceptual analyses* that study not only particular values in the online context, but also complexities that arise when trade-offs among competing values are required in a design; (2) *translations* of well-analyzed values into technical implementations; (3) *proof-of-concept projects* in which multidisciplinary teams apply Value-Sensitive Design to a particular technology, domain, or design

problem; (4) *contextual analyses* that investigate the impact of different stakeholders who influence the design and use of a technology, and who may have different goals and priorities that, in turn, lead to different value trade-offs; (5) *integrating the methodology* into organizational structures and work practices; and (6) *criteria and metrics*—both qualitative and quantitative—that can guide the design process and assess the success of particular designs.

HUMAN VALUES WITH ETHICAL IMPORT

We review and discuss 12 specific human values with ethical import. Some of these values have garnered individual chapters in this handbook. But, by including these values here, we highlight their ethical status and thereby suggest that they have a distinctive claim on resources in the design process.

Two caveats. First, not all the values we review are fundamentally distinct from one another. Nonetheless, each value has its own language and conceptualizations within their respective fields, and thus warrants separate treatment here. Second, this list is not comprehensive. Perhaps no list could be, at least within the confines of a chapter. Peacefulness, compassion, love, warmth, creativity, humor, originality, vision, friendship, cooperation, collaboration, purposefulness, devotion, diplomacy, kindness, musicality, harmony—the list of other possible values could get very long very fast. Our particular list comprises many of the traditional values that hinge on the deontological and consequentialist moral orientations reviewed previously: Human welfare, ownership and property, privacy, freedom from bias, universal usability, trust, autonomy, informed consent, and accountability. In addition, we have chosen several nontraditional values within the HCI community: Identity, calmness, and environmental sustainability. Our goal here is not only to point to important areas of future inquiry, but also to illustrate how an overarching framework for human values and ethics in design can move one quickly and substantively into new territory.

Human Welfare

Perhaps no value is more directly salient to individuals at large as their own welfare, and the welfare of other human beings. Moreover, societal interest in the moral dimensions of computing often arises when such harms occur. The faulty computer technology and interface design of Therac-25 led to the physical suffering and death of cancer patients (Leveson & Turner, 1993). Faulty computer technologies are implicated in nuclear accidents like Chernobyl. Indeed, the Risks literature (Neumann, 1995) is full of almost daily examples.

Some people in the HCI community, like Leveson (1991) refer broadly to harms that impact people and objects as harms to physical welfare. But there is merit in demarcating three categories of welfare claims (Friedman, 1997b; Kahn, 1992, 1999; Turiel, 1983). *Physical welfare* appeals to the well-being of individuals' biological selves, which is harmed by injury, sickness, and death. Many of the previous examples—that involve

Therac-25, “smart” missiles gone astray, nuclear accidents, and so forth—involve harms of this kind. In many such instances, physical harms occur because of faulty hardware or software. But it is important to recognize that technological designs working correctly can themselves lead to corresponding harms or benefits. For example, by means of the connectivity of the web, stalkers can more easily find their “prey,” and pedophiles can more easily find unwitting children to contact and potentially abuse. On the positive side, the mediating quality of the web can enhance anonymity (e.g., in a Chat Room) and buffer individuals from physical harm. *Material Welfare* appeals to physical objects that humans value and human economic interests. It goes hardly without saying that a major driving force behind computer technologies in general has been to enhance the material welfare of humans. In turn, a good deal of effort is spent protecting against material harms that computers can engender, such as damaged data from computer viruses and stolen financial information. *Psychological Welfare* refers to the higher order emotional states of human beings, including comfort, peace, and mental health. The connectivity offered by the web again provides a clear example of psychological benefits and harms. A benefit accrues when the web enhances friendships; and a harm occurs when the web allows for new forms of betrayal (e.g., the “friend” you thought you found in a chat room turns out to be a Bot).

Ownership and Property

From the anthropological literature, it would appear that all cultures embrace the idea that people can own property, although the form of such conceptions can vary considerably (Herskovits, 1952). In perhaps its most stringent form, ownership can be understood as a general right to property, which, in turn, entails a group of specific rights, including the right to possess an object, use it, manage it, derive income from it, and bequeath it (Becker, 1977).

This basic concept of a property right seems simple enough when applied to tangible objects. If you own a table, for example, you can keep it in your house, eat on it, let others eat on it, rent it, and give it away. But current computational technologies blur the boundaries between the tangible and intangible. Herein lie difficult questions. Can users, for example, legitimately copy software for their personal use? After all, to do so involves no loss of a physical object from the original seller. Or can a programmer modify part of the code from an operating system and then sell the modified system? Provisional legislative answers in the United States have hinged on nuanced distinctions between copyright, patent, licensing, and trade secrets. For example, unlike a patent, a copyright does not protect ideas, but only the expression of an idea once it is fixed in a tangible medium. The courts have largely granted developers copyrights (but not patents) to their software. But even here further controversies arise. Let us say that a developer does not copy any code of a competitor's software, but copies its “look and feel”? Is that an infringement? The courts have said yes if other technical means readily exist for implementing a different “look and feel” for the interface (Lipinski & Britz, 2000).

While case law is solidifying for issues that focus on computer software, it is hardly keeping pace with other computational developments that impact ownership and property. Consider, for example, a workplace discussion group. Who gets to decide whether the comments from the online discussion forum are made accessible to the public at large? In other words, do the workers have equal ownership of the compiled contributions? Or is ownership divided by level of participation? Or is the owner the organizer of the discussion group? Or the president of the company? Can one member delete the contribution from another member (if, for example, he or she finds the contribution offensive)? Or consider Goldberg's (2000) Telegarden, an installation that combines robotics and a real-time web camcorder, such that users interact from a remote distance with a real physical environment (a garden), and plant seeds and water them. Who owns each actual live plant? The user? The owner of the installation? Goldberg himself? Who decides how to respond to an infestation of pests, and whether pesticides will be used? If the garden becomes overcrowded, do the people who physically maintain the site have the right to uproot all the plants that were planted and cared for remotely by others? Or consider whether a cookie on a user's machine belongs to the user or to the Web site that set the cookie. If the cookie belongs to the user, and the browser does not allow the average user to delete the cookie, then by virtue of a technical mechanism HCI designers have in effect deprived an average user of the capability to exercise one of his or her specific property rights.

Our point is that through conventions and technical mechanisms, HCI designers shape answers to questions of ownership and property that lie at the forefront of social discourse and legal codification.

Privacy

Privacy refers to a claim, an entitlement, or a right of an individual to determine what information about himself or herself can be communicated to others (Schoeman, 1984). Historically, a good deal of our privacy protections arose because it was simply too much effort to collect and sort through relevant information about other individuals. But as computer technologies increasingly garner large amounts of information about specific individuals, and increasingly link information databases, our historical protections for privacy are being eroded.

In the HCI community, three general approaches have emerged for privacy protections. One approach informs “people when and what information about them is being captured and to whom the information is being made available” (Bellotti, 1998, p. 70). For example, a video monitor next to a surveillance camera in a convenience store can inform customers of the information being recorded. A second approach allows “people to stipulate what information they project and who can get hold of it” (Bellotti, 1998, p. 70). Tang's (1997) example of an on/off switch for a video conferencing workstation (noted earlier) maps onto this approach. A third approach applies privacy-enhancing technologies (PETs) that prevent sensitive data from being tagged to a specific individual in the first place. PETs have become “extraordinarily successful” in taxing

“the capacities of even the most powerful surveillance institutions” (Phillips, 1998, p. 245), and thus have provoked sharp conflicts in attempts to disseminate them. Indeed, Agre (1998) argues that PETs “disrupt the conventional pessimistic association between technology and social control. No longer are privacy advocates in the position of resisting technology as such” (p. 4).

As we move toward an era of ubiquitous computing, where embedded computation in our everyday physical objects are linked, the amount of information known about us will increase by many orders of magnitude. Thus, here, as in other domains of applications, we in the HCI community face difficult questions. How do we inform users about the risks to privacy they incur through the use of various information systems? How do we inform users about the technical protections for privacy available to them? How do we help users to understand the risks from aggregated vs. individual data? How much protection is afforded by anonymity? What should the default on systems be?—toward greater privacy protection? Or toward greater access to information? Many of these questions form part of current inquiries in HCI and, particularly CSCW, that seek to balance individual privacy with group awareness (Boyle, Edwards, & Greenberg, 2000; Cohen, Cash, & Muller, 2000; Godefroid, Herbsleb, Jagadeesan, & Li, 2000; Jancke, Venolia, Grudin, Cadiz, & Gupta, 2001; Milewski & Smith, 2000; Shoemaker & Inkpen, 2001; Svensson, Hook, Laaksolahti, & Waern, 2001).

Freedom From Bias

Bias refers to systematic unfairness perpetrated on individuals or groups. Three forms of bias in computer systems have been identified (Friedman & Nissenbaum, 1996): pre-existing social bias, technical bias, and emergent social bias.

Pre-existing social bias has its roots in social institutions, practices, and attitudes. It occurs when computer systems embody biases that exist independently of, and usually prior to, the creation of the software. For instance, the work by Nass and his colleagues (Nass & Gong, 2000; Reeves & Nass, 1996) has shown that people respond to a computer’s “gender” along stereotypical lines. Male voice interfaces are rated more competent and persuasive, and more knowledgeable about technical subjects. Female voice interfaces are viewed as more knowledgeable about topics such as love and relationships. As Nass points out, as a designer it is all too easy to inadvertently build on these pre-existing social biases when building interfaces. Thus, toward minimizing pre-existing bias designers must not only scrutinize the design specifications, but must couple this scrutiny with a good understanding of relevant biases out in the world. In addition, it can prove useful to identify potential user populations that might otherwise be overlooked and include representative individuals in the field test groups. Rapid prototyping, formative evaluation, and field testing with such well-conceived populations of users can be an effective means to detect unintentional biases throughout the design process.

Technical bias occurs in the resolution of technical design problems. For example, imagine a database for matching organ

donors with potential transplant recipients. If certain individuals retrieved and displayed on initial screens are favored systematically for a match over individuals displayed on later screens, technical bias occurs. Technical bias also originates from attempts to make human constructs such as discourse, judgments, or intuitions amenable to computers. For example, consider a legal expert system that advises defendants on whether or not to plea bargain by assuming that law can be spelled out in an unambiguous manner and is not subject to human interpretations in context. Toward minimizing technical bias, designers often need to look beyond the features internal to a system and envision the design, algorithms, and interfaces in use.

Emergent social bias emerges in the context of the computer system’s use, often when societal knowledge or cultural values change, or the system is used with a different population. Toward minimizing emergent bias, designers need to plan for not only a system’s intended contexts of use, but also its potentially emergent contexts. Yet, given limited resources, such a proposal cannot be pursued in an unbounded manner. Thus, three practical suggestions are as follows: First, designers should reasonably anticipate probable contexts of use and design for these. Second, where it is not possible to design for extended contexts of use, designers should communicate to users the contextual constraints. As with other media, we may need to develop conventions for communicating the perspectives and audience(s) assumed in the design. Third, system designers and administrators can take responsible action if bias emerges with changes in context. The National Resident Medical Match Program offers a good example. Although the original design of the Admissions Algorithm did not deal well with the changing social conditions (when significant numbers of dual-career couples participated in the match), those responsible for maintaining the system responded conscientiously to this societal change and modified the system’s algorithm to place couples more fairly (Roth, 1990).

Universal Usability

Universal usability refers to making all people successful users of information technology. Or, if this requirement of “all people” is too stringent, it can be reframed as “all people who so desire,” or in some other way. Shneiderman (1999), for example, says that universal usability means “having more than 90 percent of all households as successful users of information and communication services at least once a week.” In many respects, universal usability comprises a special case of freedom from bias, one that focuses on usability as a means to systematically prevent unfair access to information systems.

Three challenges are often addressed by proponents of universal usability: (1) *technological variety*—supporting a broad range of hardware, software, and network access; (2) *user diversity*—accommodating users with differences in, say, skills, knowledge, age, gender, disabilities, literacy, culture, and income; and (3) *gaps in user knowledge*—bridging the gap between what users know and what they need to know (Shneiderman, 1999, 2000b). Toward addressing these challenges, current work includes Aberg and Shahmehri (2001);

Beard and Korn (2001); Cooper and Rejmer (2001); Jacko, Dixon, Rosa, Scott, and Pappas (1999); and Oviatt (1999). Stephanidis (2001) provides an anthology of current concepts, methods, and tools.

Universal usability, of course, is not necessarily always a moral good insofar as it depends on what is being used or accessed. Virtually no one would suggest, for example, that there is a moral imperative to provide universal access to reruns of "I Love Lucy." Nonetheless, universal usability often is a moral good. Hert and her colleagues (Hert, Marchionini, Liddy, & Shneiderman, 2000), for example, shows that if all U.S. citizens have the right to access federal statistics, and if those statistics are only available in an online format, then to be able to exercise their right, the online federal statistics system must be usable by all U.S. citizens. Moreover, often universal access with ethical import provides increased value to a company. For example, based on his case study within a large communications corporation (NYNEX), Thomas (1997) shows how making communications systems more accessible leads to three direct corporate benefits. One, increasing access makes a communications device more pervasive in social life, and thereby more valuable for everyone. Second, increasing access increases market share. Third, increasing access "forces technologists, developers, marketers, and executives to think 'out of the box'" (p. 271).

Trust

People sometimes use the term trust broadly to include expectations of natural phenomena or machine performance. It is in this sense that people trust that "the sun will rise tomorrow" or that "brakes will stop a car." Indeed, the Computer Science and Telecommunications Board in their thoughtful publication *Trust in Cyberspace* (Schneider, 1999) adopted the terms "trust" and "trustworthy" to describe systems that perform as expected along the dimensions of correctness, security, reliability, safety, and survivability. However, equating the term "trust" with expectations for machine performance (or physical phenomena) misconstrues fundamental characteristics of this value. Specifically, trust is said to exist between people who can experience good will, extend good will toward others, feel vulnerable, and experience betrayal (Baier, 1986; Friedman, Kahn, & Howe, 2000; Kahn & Turiel, 1988). Moreover, on the societal level trust enhances our social capital (Putnam, Leonardi, & Nanetti, 1993).

In their analysis of trust online, Friedman et al. (2000) suggest that it is important for designers to distinguish two overarching contexts: e-commerce and interpersonal relationships. In e-commerce, for example, certain characteristics of the technology—such as those that concern security, anonymity, accountability, and performance history—can make it difficult for consumers to assess their possible financial harms and the potential good will of the company. One solution that has emerged has been a form of insurance (usually through the credit card companies) that limits a person's financial liability. Trust in e-commerce has received a good deal of attention in the last several years in the HCI and larger community (Dieberger, Hook, Svensson, & Lonnqvist, 2001; Egger, 2000; Fogg & Tseng, 1999;

Greenspan, Goldberg, Weimer, & Basso, 2000; Jones, Wilkens, Morris, & Masera, 2000; Olson & Olson, 2000; Shneiderman, 2000a).

In online interpersonal interactions, violations of trust make us most vulnerable psychologically: for example, hurt feelings or embarrassment. Rocco (1998) suggests that interpersonal trust online succeeds best when preceded by face-to-face interaction. In response, Zheng, Box, Olson, and Olson (2001) suggest that when users do not meet physically but engage in online chat to get to know one another that they can establish the same kinds of interpersonal trust that are established in face-to-face interactions (cf. Bos, Gergle, Olson, & Olson, 2000). Interestingly, whereas the characteristic of anonymity works against establishing financial trust in e-commerce, it is double-edged in interpersonal relationships. On the negative side, online anonymity can limit the depth of interpersonal interactions in so far as we engage in a singular means of expression (written). On the positive side, online anonymity opens up new interpersonal opportunities. For example, a gay teenager in an intolerant family and community might rely on the anonymous characteristics of the web to find and interact with like-minded peers. Thus, toward enhancing interpersonal trust online, we need to build tools that allow users to have easy and refined control about what personal information is made known to others.

Autonomy

People decide, plan, and act in ways that they believe will help them to achieve their goals. In this sense people value their autonomy. It might appear relatively easy for HCI designers to support user autonomy. The idea would be that whenever possible provide users with the greatest possible control over the technology. However, the task is harder than that. After all, most users of a word processor have little interest, say, in controlling how the editor executes a search and replace operation or embeds formatting commands. In other words, autonomy is protected when users are given control over the right things at the right time. The hard work arises in deciding what those features are and when those conditions occur.

Toward this end, Friedman and Nissenbaum (1997) identify four aspects of systems that can promote or undermine user autonomy. The first involves system capability. Recall Isaacs, Tang, and Morris's (1996) design of a system to support informal interactions in the workplace, using the piazza as the metaphor. Their application provided a means for workers to opt out of piazza interactions, and thus supported worker autonomy. The second involves system complexity. For example, as previously noted, software programs can proliferate features at the expense of usability and user autonomy suffers. The third involves misrepresentation of the system. For example, hyperbolic advertising claims can lead users to develop inaccurate expectations of the system and thereby frustrate users' goals. The fourth involves system fluidity. Over time, user's goals often change. Thus, systems need to provide ready mechanisms for users to review and adapt their systems.

Many systems depend on categorization, and it is here that Suchman (1994) worries that organizations can run roughshod

of user autonomy. For, according to Suchman, whomever determines the categories—and how those categories can be used—imputes their own personal values into the system and has power over the user. In response, Winograd (1994) points to the frequent need for socially coordinated activity through which groups of people seek to share information and technology. In such activity, Winograd argues, we need some degree of standardization wherein designers impose categories. That is their job. The key, according to Winograd, is to cultivate regularized activity without becoming oppressive. In turn, Malone (1994) responds to Suchman by arguing that not only are categories often useful, but to some extent they are necessary given the structure of human cognition. At the same time, Malone suggests that no category system is complete and that designs need to be adaptable, often by their users.

Informed Consent

Informed consent provides a critical protection for privacy, and supports other human values such as trust and autonomy. Yet, currently, there is a mismatch between industry practice and the public's interest. According to a recent report from the Federal Trade Commission (2000), for example, 59% of Web sites that collect personal identifying information neither inform Internet users that they are collecting such information nor seek the user's consent. Yet, according to a Harris Poll (2000), 88% of users want Web sites to garner their consent in such situations. The Federal Trade Commission (2000, p. iv) hopes that industry will continue to make progress on this problem, in conjunction with its proposed legislation. In turn, the HCI community has been recognizing the need to understand better what constitutes informed consent, and to realize it in online interactions.

Friedman, Millett, & Felten (2000) offer an analysis of what constitutes informed consent, and show how both words—"informed" and "consent"—have import (Faden & Beauchamp, 1986; *The Belmont Report*, 1978). The idea of "informed" encompasses disclosure and comprehension. Disclosure refers to providing accurate information about the benefits and harms that might reasonably be expected from the action under consideration. Comprehension refers to the individual's accurate interpretation of what is being disclosed. In turn, the idea of "consent" encompasses voluntariness, competence, and agreement. Voluntariness refers to ensuring that the action is not controlled or coerced, and that an individual could reasonably resist participation should he or she wish to. Competence refers to possessing the mental, emotional, and physical capabilities needed to give informed consent. Agreement refers to a reasonably clear opportunity to accept or decline to participate. Based on this account, Friedman, Millett, and Felten (2000) offer general design principles for informed consent online. Namely: (1) Decide whether the capability is exempt from informed consent. (2) Be particularly careful when invoking the sanction of implicit consent. (3) Defaults settings should err on the side of preserving informed consent. (4) Put the user in control of the "nuisance factor." (5) Avoid technical jargon. (6) Provide the user with choices in terms of potential effects rather than in terms of technical mechanisms. (7) Field test to help ensure

adequate comprehension and opportunities for agreement. (8) Design proactively.

In conjunction with this analysis, Millett et al. (2001) have examined how cookie technology and web browser designs have responded to concerns about informed consent. Specifically, they document relevant design changes in Netscape Navigator and Internet Explorer over a 5-year period, starting in 1995. Their retrospective examination led them to conclude that while cookie technology has improved over time regarding informed consent, some startling problems remain. They specify six problems and offer design remedies. For example, in both Netscape Navigator and Internet Explorer, the information disclosed about a cookie still does not adequately specify what the information will be used for or how the user might benefit or be harmed by its use. One remedy is to redesign the browser's cookie dialog box to include three additional fields, one for stating the purpose for setting the cookie, one for a brief statement of benefits, and one for a brief statement of risks.

Accountability

Medical expert systems. Automated pilots. Loan approval software. Computer-guided missiles. Increasingly, computers participate in decisions that affect human lives. In cases of computer failure, there is a common response to blame the computer—"it's the computer's fault." Indeed, Friedman and Millett (1995) found that 83% of undergraduate computer science majors she interviewed attributed aspects of agency—either decision making and/or intentions—to computers. In addition, 21% of the students consistently held computers morally responsible for error. Yet, if we accept that humans, but not computational systems, are capable of being moral agents, then such blame is fundamentally misplaced (Searle, 1980).

How can HCI designers minimize this tendency of users to attribute blame to computational systems? Suggestions have been offered based on minimizing two types of distortions (Friedman & Kahn, 1992). In the first type of distortion, the computational system diminishes or undermines the human user's sense of his or her own moral agency. In such systems, human users are placed into largely mechanical roles, either mentally or physically, and frequently have little understanding of the larger purpose or meaning of their individual actions. To the extent that humans experience a diminished sense of agency, human dignity is eroded and individuals may consider themselves to be largely unaccountable for the consequences of their computer use. Conversely, in the second type of distortion the computational system masquerades as an agent by projecting intentions, desires, and volition. Strikingly, even when computer interfaces only minimally mimic human agency, people appear predisposed—at least in certain regards—to treat computers as social agents. For example, as Nass and his colleagues have shown, people respond to multiple voices from a single computer as though they were separate entities, respond to a computer's "gender" along stereotypical lines, are less likely to criticize a computer directly (i.e., if the computer itself asks for an evaluation) than to criticize the computer to a third party (a different computer or human), and respond to computer flattery

(Nass & Gong, 2000; Reeves & Nass, 1996). Thus, to the extent that humans inappropriately attribute agency to computational systems, humans may well consider the computational systems, at least in part, to be morally responsible for the effects of computer-mediated or computer-controlled actions.

Identity

The idea of personal identity embraces two seemingly contradictory ideas. On the one side is the obvious fact that each one of us has many roles. A single person can be, for example, a father, lover, poker player, gourmet cook, computer geek, and animal lover. Indeed, William James says that a person "has many social selves as there are individuals who recognize him" (quoted in Rosenberg, 1997, p. 23). On the other side, virtually all of us feel like we live reasonably coherent lives, and that the person we are today is pretty much whom we were yesterday and last week, if not last year. Thus, identity appears to be multiple and unified, and both aspects are essential to human flourishing. Too far toward multiplicity and we end up schizophrenic; too far toward being unified and there are too few mechanisms for psychological growth, and too little basis for healthy social functioning (e.g., it makes little sense to present the same "persona" to one's boss as one does to one's child).

In terms of HCI, it is important that the field as a whole substantively support both manifestations of identity, and that in our designs we not swing too far one direction or the other. To date, the networked personal computer has easily supported multiplicity. Thus, the same person can easily communicate with many unrelated groups (chats, list serves, etc.), easily establish a different identity within many of these groups, or even multiple identities within a single group (Turkle, 1996). This trend is being checked in interesting ways. For example, individuals in some online communities link their comments to their personal homepage ("to find out more about me, click here"). Some chats (e.g., SeniorNet) require single identities. Indeed, with the advent of increasingly linked online databases coupled with ubiquitous computing, we believe that the pendulum will soon shift powerfully and pervasively toward the unification of identity. If so, following Bers, Gonzalo-Heydrich, and DeMaso (2001) and Schiano and White (1998), the challenge for HCI designers will be to find ways such that individuals have flexibility to establish and reveal not only an integrated self, but a multiplicity of identities.

Calmness

Weiser and Brown (1997) suggest that in the last 50 years, there have been two great trends in the human relationship with computing: The first was with the mainframe computer and the second is currently with the personal computer. They suggest that the next great trend is toward ubiquitous computing, characterized by deeply embedded computation in the world. In turn, they argue that the "most potentially interesting, challenging, and profound change implied by the ubiquitous computing era is a focus on *calm*. If computers are everywhere, they had better

stay out of the way, and that means designing them so that the people being shared by the computers remain serene and in control" (p. 79).

The central design mechanism Weiser and Brown put forward is that which conveys information in our peripheral awareness, but wherein the information is moved to the center of our awareness at appropriate times. One architectural example they provide is of a glass window between offices and hallways. Such a window, they suggest, "extends our periphery by creating a two-way channel for clues about the environment. Whether it is motion of other people down the hall (it's time for lunch; the big meeting is starting) or noticing the same person peeking in for the third time while you are on the phone (they really want to see me; I forgot an appointment), the window connects the person inside to the nearby world" (pp. 81-82). A computational example involves an Internet multicast, a continuous video from another location that provides not so much videoconferencing, but "more like a window of awareness" of another location (p. 82). Another example involves a "dangling string," "an eight-foot piece of plastic spaghetti that hangs from a small electric motor mounted in a ceiling. The motor is electronically connected to a nearby Ethernet cable, so that each bit of information that goes past causes a tiny twitch of the motor. A very busy network causes a madly whirling string with a characteristic noise; a quiet network causes only a small twitch every few seconds" (p. 83).

The idea that ubiquitous computing will need to take up the challenge of how to preserve calmness in human lives seems on the mark. There are, however, some limitations in the solution Weiser and Brown put forward in terms of designing for the periphery. Perhaps the most notable is that very quickly the periphery itself can become overloaded with information, especially if work place conventions expect workers to be continuously aware of such information. Imagine, for example, working in an office with a window out onto the hallway, with four Internet multicasts playing in the background, a dangling string in the corner, five other peripheral information streams on each of several pieces of software you are using, and then add in peripheral information streams on one's regular phone, cell phone, personal digital assistant, and any and all other pieces of embedded computation in the office of today (or the future). Then try to get some work done. It will not be easy.

An alternative design approach (Kahn & Friedman, in preparation; cf. Farrell, 2001) toward preserving calmness has as its starting point the psychological literature that shows that direct experiences with nature have beneficial effects on people's physical, cognitive, and emotional well-being. For example, Ulrich (1984) found that postoperative recovery improved when patients were assigned to a room with a view of a natural setting (a small stand of deciduous trees) versus a view of a brown brick wall. More generally, studies have shown that even minimal connection with nature—such as looking at a natural landscape—can reduce immediate and long-term stress, reduce sickness of prisoners, and calm patients before and during surgery (see, for reviews, Beck & Katcher, 1996; Kahn, 1999; Ulrich, 1993). Building on this literature, the question becomes in what ways computer technologies can augment the human relationship with the natural world with beneficial effects. One

implementation might involve video plasma display "windows" in inside offices that stream in real-time views of a psychologically restorative local nature scene. Another implementation might involve robotic "pets" as possible companions for the elderly. The point is that while such augmented natural interactions may not be as psychologically beneficial as real nature, it may be more so than no nature, in which case it becomes a plausible area for design applications.

Environmental Sustainability

We pollute air and water, deplete soil, deforest, create toxic wastes, and through human activity are extinguishing over 27,000 species each year (a conservative estimate). Such problems have generated a tremendous amount of attention among the populace, and have slowly been coming under the purview of the computing community (cf. IEEE and ACM: Software Engineering Code of Ethics and Professional Practice, 1998). On the production end, there is concern about the resources used in producing computer technologies, as well as the resulting toxic wastes (e.g., of making computer chips). On the consumption end, there is concern about the resources computer technologies use. The electrical demands, for example, are particularly high, especially when energy sources are scarce.

Another question arises of how we can design systems that foster a healthier and more life-affirming connection with the natural world (cf. Kahn & Kellert, 2002). The more straightforward approach has been simply to harness computer technologies in the service of environmental science, in computer models, for example, of global warming or earth tectonics. Environmental educators have also been using networked computers such that students can share environmental data or narratives with other students (or scientists) in diverse geographical locations. Or consider an approach being taken by Borning and his colleagues (Noth, Borning, & Waddell, 2000) where he is using a computer simulation ("UrbanSim") of an urban environment to help residents, politicians, and planners visualize the effects of their proposed land-use plans. Another approach might involve designing calming technologies, as previously discussed, that integrate restorative aspects of nature into human lives.

PROFESSIONAL ETHICS

Computer professionals not uncommonly experience ethical conflicts in the workplace. Consider, for example, the situation where, by means of his or her specialized knowledge, a computer professional knows of a harm that can result from the implementation of a computer technology. Should he or she inform others of such impending harm even if it will jeopardize his or her job security? Or consider the situation where a client asks a software developer to develop an expert system that recommends against loans in certain neighborhoods. Although such a policy (called "redlining") may serve the economic interests of the loan-granting organization, it unfairly discriminates against individuals on the basis of where they live. Should a computer

professional deny such a request? Or consider the situation offered by Quesenbery (Molich, 2001):

You have set up early usability tests of a paper prototype with nurses at a medical facility. The test was difficult to schedule because nurses' time is guarded carefully and marketing carefully guards the relationship with customers. The nursing managers insist on 'taking the test' themselves first, and then on being present in the room during the test with the other participants to 'be sure they do it right.' You believe that the manager's presence will be intimidating to the nurses, altering the results of the test. Do you continue with the usability tests? (p. 218)

Quesenbery offers two possible answers. One answer is yes. She argues that if you turn down this opportunity, another one might not arise; and although you would have to carefully evaluate the test results for bias caused by the managers, valuable data about the prototype would still be generated. The other answer is no. She argues that it is unethical to put nurses into such a stressful situation where they are studied in front of their obviously critical managers.

In response to ethical issues that arise within the computing professions, numerous organizations have developed ethical codes of conduct. Codes have been developed, for example, by the ACM, IEEE Computer Society, DPMA, and ICCP. Berleur and Brunnstein (1997) provide a comparison of 30 different codes of ethics. At the end of 1998, the IEEE and the ACM together adopted a revised Software Engineering Code of Ethics and Professional Practice (<http://www-cs.etsu.edu/seeri/secode.htm>). At the highest level of abstraction (see Gotterbarn, 1999, for a review), the code states eight principles wherein software engineers shall (1) act consistently with the public interest; (2) act in a manner that is in the best interests of their client and employer consistent with the public interest; (3) ensure that their products and related modifications meet the highest professional standards possible; (4) maintain integrity and independence in their professional judgment; (5) subscribe to and promote an ethical approach to the management of software development and maintenance; (6) advance the integrity and reputation of the profession consistent with the public interest; (7) be fair to and supportive of their colleagues; and (8) participate in life-long learning regarding the practice of their profession and shall promote an ethical approach to the practice of their profession. In turn, each principle is elaborated upon with specific guidelines. For example, under the first principle (to act consistently with the public interest), the code states that software engineers should "disclose to appropriate persons or authorities an actual or potential danger to the user, the public, or the environment, that they reasonably believe to be associated with the software or related documents."

Before moving forward with this topic, it is important to take hold of a question often discussed in the literature, of whether the computer profession is even a profession (Weckert & Adeney, 1997). If it is not, then concerns about professional ethics—special moral requirements above and beyond what are applied to ordinary people—disappear.

What, then, is a profession? According to Bayles (1981), a profession is comprised of members who (1) have extensive training, (2) have an intellectual component such that the professional's primary role is to advise the client about things the

client does not know, and (3) provide an important service to society. Bennion (1969) also points out that professionals are typically self-employed, and establish a client-agent relationship based on trust. In addition, as Johnson and Miller (1997) write: "Being a professional means more than just having a job. The difference is commitment to doing the right thing because you are a member of a group that has taken on responsibility for a domain of social activity—a social function. The group is accountable to society for this domain, and for this reason, professionals must behave in ways that are worthy of public trust" (p. 22).

Based on such criteria, in some respects the computer field constitutes a profession. Computer personnel have intellectual knowledge and extensive training, and they provide an important service to society. Computer personnel are also especially well positioned to understand how new technologies may impact human lives; accordingly, many people argue that they incur the responsibility to communicate such impacts to the general public. Too, the media, industry, and academy widely use the term "computer professionals." That said, the computer field certainly does not constitute a canonical profession, like medicine or law. After all, many computer personnel work for businesses rather than as a consultant, and create artifacts (e.g., software, hardware, algorithms, interface designs) rather than offer advice. Too, computer personnel vary greatly in the amount and type of their training. Some computer personnel are self-taught, without even a high-school diploma. Other computer personnel complete 6-month to 2-year programs, or a bachelor's, master's, or doctoral program.

If we accept—as most people do—that in at least some important regards the computer field constitutes a profession, then how has the profession understood and conveyed its correlative ethical responsibilities? One approach has been through codes of ethics, noted previously. Such codes serve multiple purposes (Gotterbarn, 1999; Anderson, Johnson, Gotterbarn, & Perrolle, 1993). They serve to educate computer employees and managers. They help garner the trust of the general public. They provide computer employees with a formal document to turn to (and appeal to) when they face conflicts that pit ethical decisions against the economic benefit of their company. They can function as a means of deterrence and discipline. And they can enhance a profession's public standing.

Another approach for advancing the profession's ethical responsibilities involves using hypothetical scenarios to flesh out the meaning of the codes (Parker, 1979; Parker, Swope, & Baker, 1990). For example, Anderson et al. (1993) provide an analysis of nine scenarios that illustrate how the 1992 ACM Code of Ethics bear out in practice. They offer, for example, the hypothetical situation of a consultant, named Diane, who is designing the database management system for the personnel office of a medium-sized company. The database will store sensitive information, including performance evaluations, medical records, and salaries. In an effort to cut costs, the CEO of the company rejects Diane's suggestions and opts for a less secure system. Diane remains convinced that a more secure system is needed. What should she do? Anderson et al. (1993) then point to specific passages in the code that pertain to the importance of maintaining privacy and confidentiality. Based on the code, Diane's first obligation is to try to educate the company officials on

the ethics of the situation. "If that fails, then Diane needs to consider her contractual obligations . . . on honoring assigned responsibilities" (p. 100).

Another possible approach toward advancing the profession's ethical responsibilities involves that of licensing its members. The model draws from the American Medical Association or the American Bar Association. Both associations can revoke their members' licenses given serious violations of their respective professional organization's code of ethics. But there are good reasons why most professional computing organizations have resisted this move. Computing is an enormously diverse field; and it is difficult to formally demarcate all its areas, let alone the relevancy of one area to others. What exactly must an independent Web site designer know, for example, beyond whatever the designer and employer believe necessary? Does the field really want to exclude people all along its fringes? There has been concern that through licensing the profession might become a closed shop and thereby enhance the status and incomes of those admitted at the expense of those excluded. There has also been concern that licensing could stifle creativity, originality, and excellence (see the discussion by Neumann, 1991). Licensing also runs counter to the history of computing that has emphasized ease of entry, entrepreneurship, and a culture not often found in professional licensing organizations.

Setting aside licensing, and recognizing the value of codes of ethics and analyses of hypothetical dilemmas, what else can the computing profession do to advance its ethical responsibilities? Johnson and Miller (1997) offer two suggestions. One is that corporations could have ombudspersons to whom computing professionals could report their concerns, anonymously if desired. Another is that professional societies (like the ACM) could maintain "hotlines that professionals could call for advice on how to get their concerns addressed" (p. 23). Both suggestions, if implemented, could prevent many problems from escalating into "whistle-blowing" affairs and help integrate ethical considerations into the workplace.

The computing profession is also recognizing the importance of analyzing not only hypothetical scenarios, but real life events. Neumann (1995) is a pioneer in this regard, publishing ongoing chronicles of the risks that arise through computing. For the more interesting cases, Neumann's approach could be extended. What we have in mind here builds on a leading mountaineering journal that each year publishes brief accounts of the major mountaineering accidents of the previous year. To read these accounts, year in and year out, is to build up a rich repertoire of how mistakes happen in the field, and how they can be avoided. Similarly, it would be possible to provide comprehensive accounts of the major ethical problems that have arisen each year in the computing profession. Moreover, these accounts could be classified based on type of ethical value (e.g., privacy, trust, human welfare, security, or universal usability), user population (e.g., workers, elderly, children, or people in developing nations), and applications (e.g., online communities, information visualization, augmented reality, wearable computing, or groupware). With such a searchable database, computing professionals could gain ready access to real-life ethical case studies directly relevant to their own endeavors.

Most of the ethical issues of HCI professionals are subsumed under those encountered by computer professionals in general.

But HCI professionals incur unique responsibilities whenever they involve human subjects. At universities, such studies come under the purview of a Human Subjects Institutional Review Board (an IRB) that seeks to protect the rights and welfare of human subjects. But, in industry, such oversight is almost always lacking. Consider the following questions:

- You conduct usability studies with people within your company. The vice-President later asks you which people did well. Do you tell (Molich, 2001)?
- You collect video footage of people engaged in your usability study. Can you show the footage at a conference (Mackay, 1995)? Have you obtained written consent from the subjects? What if the subject initially gives consent but then comes off looking foolish in the video. Have you given the subject a chance to review the video footage and to opt out?
- You obtain informed consent, and then collect keystroke data on employees over a period of months. How do you keep your subjects aware that your data collection is on-going?
- You collect data from subjects living in an "aware home" (that has ubiquitous computation embedded throughout the living environment). How do you convey to subjects the extent of personal information (e.g., time spent in the bathroom) that can be mined from the resulting database?
- You collect data from employees in your company, and promise to keep identifying information confidential. But if there is only one disabled person in a given department, how can you maintain confidentiality?
- You are involved in a safety critical situation. You recognize that often usability studies in your company involve only a handful of subjects, which allows for fast prototyping and more quickly bringing products to market. But, in this situation, is a more rigorous psychological study of usability needed, even given that it will require additional resources and time (Molich, 2001)?
- You collect information from an online chat focused on sexual abuse, and then publish the information widely, quoting key emotional passages that now become visible on a societal level. Even though you report the passages without identifying names, others can now go to this chat, search on the text, and easily identify the people of interest. Have you adequately protected the privacy of these online "subjects"? How does using subjects garnered from the Internet differ from those garnered from more traditional subject pools (Frankel & Siang, 1999)?

These questions and many other related issues can be addressed by drawing on the rich literature on protecting the rights and welfare of human subjects (*The Belmont Report*, 1978) and its application to the internet (Frankel & Siang, 1999). Indeed, often HCI professionals may be the only source for giving voice to these ethical considerations.

CONCLUSIONS

During the early periods of computerization, around the 1950's, cyberneticist Norbert Wiener (1953/1985) argued that technology could help make us better human beings and create a more

just society. But for it to do so, he argued, we have to take control of the technology. We have to reject the "worshiping [of] the new gadgets which are our own creation as if they were our masters" (p. 678). Similarly, a few decades later, computer scientist Joseph Weizenbaum (1972) wrote:

What is wrong, I think, is that we have permitted technological metaphors... and technique itself to so thoroughly pervade our thought processes that we have finally abdicated to technology the very duty to formulate questions... Where a simple man might ask: "Do we need these things?" technology asks "what electronic wizardry will make them safe? Where a simple man will ask "is it good?," technology asks "will it work?" (pp. 611-612).

As HCI professionals, we have profound opportunities to shape the designs and implementations of computer technologies from an ethical stance.

In this chapter, we have reviewed varying approaches, projects, and ideas that offer us important ways of bringing human values and ethics into our design practice. Several ideas are worth emphasizing. First, as is well known, it is much easier to design systems right initially than to attempt to retrofit poor systems after they have become entrenched within organizations and other social systems. Thus, it is imperative that we take a proactive stance on human values, ethics, and design. Second, many of the difficult problems in this area require multidisciplinary collaborations. Third, we need to hold out human values with ethical import as a design criterion—along with the traditional criteria of reliability, efficiency, and correctness—by which systems may be judged poor and designers negligent. As with the traditional criteria, we need not require perfection, but commitment.

This *Handbook* has an entire section devoted to issues of diversity. Individual chapters include, for example, designing for gender differences, children, the elderly, internationalization, and motor, perceptual, and cognitive impairments. Yet, in our review of the literature, current HCI approaches to human values and ethics do not always fare well when used in diverse contexts. For example, Participatory Design can be particularly effective when a community shares many deeply held sensibilities, such as a commitment to participatory democracy in the workplace and to the idea of meaningful work itself. But this approach is more difficult to apply when divisive constituencies argue on the basis of narrowly conceived self-interests or hostile prejudices. What happens, for example, when a manager values accountability over a worker's privacy? Or efficiency over a worker's autonomy? If each value has equal weight, on what basis does a designer move forward with a particular design?

As the field of HCI seeks to bring moral commitments into diverse contexts, it needs a principled moral means to adjudicate competing value claims. It is for this reason that we provided early on a moral philosophical and psychological framework for approaching this problem. We suggested, for example, that when moral conflicts occur between diverse groups, that such variability by itself does not prove or disprove the moral relativist's position. People can believe that a certain act (such as racial discrimination) is moral; but documenting such a belief does not make it so. We also suggested that when moral differences appear in another culture, that a question to ask

is whether there are victims within that culture. If there are, then it is probably less the case that societies differ on moral ground, and more that some societies (Western societies included) may be involved explicitly in immoral practices. We suggested that close attention must be paid to the level of moral analysis. Namely, definitions of morality that entail abstract characterizations of justice and welfare tend to highlight moral universals, whereas definitions that entail specific behaviors or rigid moral rules tend to highlight moral cross-cultural variation. In our view, both levels of analysis have merit, and a middle ground provides an epistemically sensible and powerful approach in HCI research: One that allows for an analysis of universal moral constructs (such as justice, rights, welfare, and virtue), as well as allowing for the ways in which these constructs play out in a particular culture at a particular point in time. By embracing this moral theoretical framework, we are not saying that all moral problems can be solved. But such problems will be discussed (and argued about) in ways that respect diversity and prevent oppression.

In the future, certain trends—technologically and societally—will pose particular challenges in terms of human values, ethics, and design. From our view, three trends stand out. First, computational technologies will increasingly allow for the erosion of personal privacy. Even today, for example, surveillance cameras capture our images in banks and airports, and in many stores, malls, and even streets. In cars, GPS navigation systems not only receive positioning data, but can broadcast one's position. Businesses can (and sometimes do) monitor workers' electronic communications. Indeed, think of perhaps the last bastion of the private space, the home; and recognize that "aware homes" of the future will have the potential to record virtually every movement an individual makes within his or her home, and to link that data to large networked databases. Thus, protections of individual privacy will need to become an even more central concern to the HCI community. The second trend is that computational technologies will increasingly provide means for government to erode civil

liberties. Moreover, the public may unwittingly accept such consequences based on the assumption that these technologies will substantially enhance our nation's physical security. But such an assumption is not always warranted. Imagine, for example, if the government required each of us to obtain biometric national identification cards, and to use these cards for all of our financial transactions, air travel, entry into government buildings, and so forth. Secure? Hardly. For the systems that store this data can be cracked. Indeed, once biometric data is stolen, it is not so easy to get a new face, fingerprint, or DNA. Third, in our increasingly linked communicative technological infrastructure, the pervasiveness and speed of information can undermine people's psychological well-being and quality of life. How can we respond to this problem of information overload? One solution is that we need, at times, to turn away from our technological devices and thus to check their encroachment into human lives. Another solution is that, in our designs, we need to find ways to increase the ratio of quality over quantity of information. We can also profit by designing "calming" technologies (that go beyond placing information in our peripheral awareness, which itself can become overloaded).

In the last decade, the HCI community has made tremendous progress in integrating human values and ethics into the practice of design. As the field continues to move forward, the challenge remains how to design technology wisely, ethically, so as to create essential conditions by which humans live and flourish.

ACKNOWLEDGMENT

We thank Alan Borning and Gaetano Borriello for sustained discussions on human values and computational technologies. Valerie Wonder assisted with library research. The ideas presented in this paper are, in part, based on work supported by the National Science Foundation under Award Nos. IIS-0000567, IIS-9911185, and SES-0096131.

References

- Aberg, J., & Shahmehri, N. (2001). An empirical study of human web assistants: Implications for user support in web information systems. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2000)* (pp. 404-411). New York: Association for Computing Machinery Press.
- Abowd, G. D., & Jacobs, A. (2001, September/October). The impact of awareness technologies on privacy litigation. *SIGCHI Bulletin*, 9.
- Ackerman, M. S., & Cranor, L. (1999). Privacy critics: UI components to safeguard users' privacy. In *Extended Abstracts Conference on Human Factors in Computing Systems (CHI 99)* (pp. 258-259). New York: Association for Computing Machinery Press.
- Adler, P. S., & Winograd, T. (Eds.). (1992). *Usability: turning technologies into tools*. Oxford: Oxford University Press.
- Agre, P. E. (1998). Introduction. In P. E. Agre & M. Rotenberg (Eds.), *Technology and privacy: The new landscape* (pp. 1-28). Cambridge, MA: The MIT Press.
- Agre, P. E., & Mailloux, C. A. (1997). Social choice about privacy: Intelligent vehicle-highway systems in the United States. In B. Friedman (Ed.), *Human values and the design of computer systems* (pp. 289-310). New York: Cambridge University Press.
- Agre, P. E., & Rotenberg, M. (1998). *Technology and privacy: The new landscape*. Cambridge, MA: The MIT Press.
- Anderson, R. E., Johnson, D. G., Gotterbarn, D., & Perrolle, J. (1993). Using the ACM code of ethics for decision making. *Communications of the ACM*, 36(2), 98-107.
- Appadurai, A. (Ed.). (1988). *The social life of things*. New York: Cambridge University Press.
- Akrich, M. (1992). The description of technical objects. In W. E. Bijker & J. Law (Eds.), *Shaping technology/building society: Studies in sociotechnical change* (pp. 205-224). Cambridge, MA: The MIT Press.
- Attwell, P. (1987). The deskilling controversy. *Work and Occupation*, 14(3), 323-346.
- Baier, A. (1986). Trust and antitrust. *Ethics*, 96, 231-260.

- Bayles, M. D. (1981). *Professional ethics*. Belmont, CA: Wadsworth.
- Beard, M., & Korn, P. (2001). What I need is what I get: Downloadable user interfaces vis Jini and Java. In *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2001)* (pp. 15-16). New York: Association for Computing Machinery Press.
- Beck, A., & Katcher, A. (1996). *Between pets and people*. West Lafayette, IN: Purdue University Press.
- Becker, L. C. (1977). *Property rights: Philosophical foundations*. London: Routledge & Kegan Paul.
- The Belmont Report*. (1978). *Ethical principles and guidelines for the protection of human subjects of research*. Washington, DC: The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research.
- Bellotti, V. (1998). Design for privacy in multimedia computing and communications environments. In P. E. Agre & M. Rotenberg (Eds.), *Technology and privacy: The new landscape* (pp. 63-98). Cambridge, MA: The MIT Press.
- Bennett, C. J. (1998). Convergence revisited: Toward a global policy for the protection of personal data? In P. E. Agre & M. Rotenberg (Eds.), *Technology and privacy: The new landscape* (pp. 99-123). Cambridge, MA: The MIT Press.
- Bennion, F. A. R. (1969). *Professional ethics*. London: Charles Knight & Co.
- Berleur, J., & Brunnstein, K. (Eds.). (1997). *Ethics of computing: Codes, spaces for discussion, and law*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Bers, M. U., Gonzalez-Heydrich, J., & DeMaso, D. R. (2001). Identity construction environments: Supporting a virtual therapeutic community of pediatric patients undergoing dialysis. In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 380-387). New York: Association for Computing Machinery.
- Bjerknes, G., & Bratteteig, T. (1995). User participation and democracy: A discussion of Scandinavian research on system development. *Scandinavian Journal of Information Systems*, 7(1), 73-97.
- Bodker, S. (1990). *Through the interface—A human activity approach to user interface design*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Borgman, C. L. (2000). *From Gutenberg to the global information infrastructure: Access to information in the networked world*. Cambridge, MA: The MIT Press.
- Bos, N., Gergle, D., Olson, J. S., & Olson, G. M. (2001). Being there versus seeing there: Trust via video. In *Extended Abstracts of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 291-292). New York: Association for Computing Machinery.
- Boyd, R. A. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 181-228). Ithaca, NY: Cornell University Press.
- Boyle, M., Edwards, C., & Greenberg, S. (2000). The effects of filtered video on awareness and privacy. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 1-10). New York: Association for Computing Machinery.
- Bulliet, R. (1994). Determinism and pre-industrial technology. In M. R. Smith & L. Marx (Eds.), *Does technology drive history? The dilemma of technological determinism* (pp. 201-216). Cambridge, MA: The MIT Press.
- Bynum, T. W. (Ed.). (1985). [Entire issue]. *Metaphilosophy*, 16(4).
- Campbell, R. L., & Christopher, J. C. (1996). Moral development theory: A critique of its Kantian presuppositions. *Developmental Review*, 16, 1-47.
- Cohen, A. L., Cash, D., & Muller, M. J. (2000). Designing to support adversarial collaboration. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 31-39). New York: Association for Computing Machinery Press.
- Cole, M. (1991, April). Discussant to B. Friedman & P. H. Kahn, Jr. (Presenters) "Who is responsible for what? And can what be responsible? The psychological boundaries of moral responsibility." Paper session at the biennial meeting of the Society for Research in Child Development, Seattle, WA.
- Cooper, M., & Rejmer, P. (2001). Case study: Localization of an accessibility evaluation. In *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2001)* (pp. 141-142). New York: Association for Computing Machinery Press.
- Cranor, L., & Resnick, P. (2000). Protocols for automated negotiations with buyer anonymity and seller reputations. *Netnomics* 2(1), 1-23.
- Dewan, P., & Shen, H. (1998). Flexible meta access-control for collaborative applications. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 1998)* (pp. 247-256). New York: Association for Computing Machinery Press.
- Dieberger, A., Hook, K., Svensson, M., & Lonnqvist, P. (2001). Social navigation research agenda. In *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2001)* (pp. 107-108). New York: Association of Computing Machinery Press.
- Dourish, P., & Kiesler, S. (2000). From the papers co-chairs. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (p. v). New York: Association for Computing Machinery Press.
- Dworkin, R. (1978). *Taking rights seriously*. Cambridge, MA: Harvard University Press.
- Egger, F. N. (2000). "Trust me, I'm an online vendor": Towards a model of trust for e-commerce system design. In *Extended Abstracts of the Conference of Human Factors in Computing Systems (CHI 2000)* (pp. 101-102). New York: Association for Computing Machinery.
- Ehn, P. (1989). *Work-oriented design of computer artifacts*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ehn, P., & Kyng, M. (1991). Cardboard computers: Mocking-it-up or hands-on the future. In J. Greenbaum & M. Kyng (Eds.), *Design at work: Cooperative design of computer systems* (pp. 169-195). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Faden, R., & Beauchamp, T. (1986). *A history and theory of informed consent*. New York: Oxford University Press.
- Farrell, S. (2001). Social and informational proxies in a fishtank. In *Extended Abstracts of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 365-366). New York: Association for Computing Machinery.
- Federal Trade Commission. (2000, May). *Privacy online: Fair information practices in the electronic marketplace* (A Report to Congress). Washington, DC: Author.
- Fishkin, J. S. (1982). *The limits of obligation*. New Haven, CT: Yale University Press.
- Fogg, B. J., & Tseng, H. (1999). The elements of computer credibility. In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 1999)* (pp. 80-87). New York: Association for Computing Machinery.
- Foot, P. (1978). *Virtues and vices*. Berkeley, CA: University of California Press.
- Frankel, M. S., & Siang, S. (1999, November). *Ethical and legal aspects of human subjects research on the internet*. Washington, DC: American Association for the Advancement of Science. Workshop report [On-line]. Available: <http://www.aaas.org/spp/dssp/sf1/projects/inters/main.htm>
- Friedman, B. (Ed.). (1997a). *Human values and the design of computer technology*. New York: Cambridge University Press.
- Friedman, B. (1997b). Social judgments and technological innovation: Adolescents' understanding of property, privacy, and electronic information. *Computers in Human Behavior*, 13(3), 327-351.

- Friedman, B. (1999). *Value-sensitive design: A research agenda for information technology* (Contract No.: SBR-9729633). Report to the National Science Foundation [On-line]. Available: <http://ischool.washington.edu/vsd>
- Friedman, B., & Borning, A. (2001). Value-sensitive design: Cultivating research and community (Contract No.: IIS-0000567). Unpublished report to the National Science Foundation.
- Friedman, B., Howe, D. C., & Felten, E. (2002). Informed consent in the Mozilla browser: Implementing value-sensitive design. *Proceedings of the Thirty-Fifth Annual Hawai'i International Conference on Systems Science*. Abstract, p. 247; CD-ROM of full-paper, OSPE101. Los Alamitos, CA: IEEE Computer Society.
- Friedman, B., & Kahn, P. H., Jr. (1992). Human agency and responsible computing: Implications for computer system design. *Journal of Systems Software*, 17, 7-14.
- Friedman, B., & Kahn, P. H., Jr. (in press). A value-sensitive design approach to augmented reality. In W. Mackay (Ed.), *Design of augmented reality environments*. Cambridge, MA: The MIT Press.
- Friedman, B., Kahn, P. H., Jr., & Howe, D. C. (2000). Trust online. *Communications of the ACM*, 43(12), 34-40.
- Friedman, B., & Millett, L. (1995). "It's the computer's fault"—Reasoning about computers as moral agents. In *Conference Companion of the Conference on Human Factors in Computing Systems (CHI 95)* (pp. 226-227). New York: Association for Computing Machinery Press.
- Friedman, B., Millett, L., & Felten, E. (2000). *Informed consent online: A conceptual model and design principles*. UW-CSE Technical Report 00-12-2. University of Washington, Department of Computer Science & Engineering. Seattle, WA.
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330-347.
- Friedman, B., & Nissenbaum, H. (1997). Software agents and user autonomy. *Proceedings of the First International Conference on Autonomous Agents* (pp. 466-469). New York: Association for Computing Machinery Press.
- Fuchs, L. (1999). AREA: A cross-application notification service for groupware. In S. Bodker, M. Kyng, & K. Schmidt (Eds.), *Proceedings of the Sixth European Conference on Computer Supported Cooperative Work* (pp. 61-80). Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Galegher, J., Kraut, R. E., & Egidio, C. (Eds.). (1990). *Intellectual teamwork: Social and technological foundations of cooperative work*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gewirth, A. (1978). *Reason and morality*. Chicago, IL: University of Chicago Press.
- Godefroid, P., Herbsleb, J. D., Jagadeesan, L., & Li, D. (2000). Ensuring privacy in presence awareness systems: Automated verification approach. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 59-68). New York: Association for Computing Machinery Press.
- Goldberg, K. (Ed.). (2000). *The robot in the garden: Telerobotics and telepresence on the Internet*. Cambridge, MA: The MIT Press.
- Gotterbarn, D. (1999, November/December). How the new software engineering code of ethics affects you. *IEEE Software*, 16(6), 58-64.
- Greenbaum, J. (1996). Back to labor: Returning to labor process discussions in the study of work. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 1996)* (pp. 315-324). New York: Association for Computing Machinery Press.
- Greenbaum, J., & Kyng, M. (Eds.). (1991). *Design at work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Greenspan, S., Goldberg, D., Weimer, D., & Basso, A. (2000). Interpersonal trust and common ground in electronically mediated communication. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 251-260). New York: Association for Computing Machinery Press.
- Grief, I. (Ed.). (1988). *Computer-supported cooperative work: A book of readings*. San Mateo, CA: Morgan Kaufmann.
- Gross, M. D., Parker, L., & Elliot, A. M. (1997). MUD: Exploring tradeoffs in urban design. In *Proceedings of the 7th International Conference on Computer Aided Architectural Design Futures* (pp. 373-387). Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Grudin, J. (1988). Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '88)* (pp. 85-93). New York: Association for Computing Machinery Press.
- Harris Poll/*Business Week*. (2000). *A growing threat*. http://www.businessweek.com/2000/00_12/b3673010.htm
- Harrison, S., & Dourish, P. (1996). Re-placing space: The roles of place and space in collaborative systems. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 96)* (pp. 67-76). New York: Association for Computing Machinery Press.
- Hatch, E. (1983). *Culture and morality*. New York: Columbia University Press.
- Herskovits, M. J. (1952). *Economic anthropology: A study of comparative economics*. New York: Alfred A. Knopf.
- Hert, C., Marchionini, G., Liddy, L., & Shneiderman, B. (2000). Extending understanding of federal statistics in tables: Integrating technology and user behavior in support of university usability [On-line]. Available: <http://istweb.syr.edu/~tables>
- Houston, J. (1995). *Confessions of an igloo dweller*. New York: Houghton Mifflin.
- Hudson, S. E., & Smith, I. (1996). Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work (CSCW 96)* (pp. 248-257). New York: Association for Computing Machinery Press.
- Hughes, T. P. (1994). Technological momentum. In M. R. Smith & L. Marx (Eds.), *Does technology drive history? The dilemma of technological determinism* (pp. 101-113). Cambridge, MA: The MIT Press.
- Iacono, S., & Kling, R. (1987). Changing office technologies and the transformation of clerical jobs. In R. Kraut (Ed.), *Technology and the transformation of white collar work*. Hillsdale, NJ: Lawrence Erlbaum.
- Isaacs, E. A., Tang, J. C., & Morris, T. (1996). Piazza: A desktop environment supporting impromptu and planned interactions. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 96)* (pp. 315-324). New York: Association for Computing Machinery Press.
- Jacko, J. A., Dixon, M. A., Rosa, R. H., Jr., Scott, I. U., & Pappas, C. J. (1999). Visual profiles: A critical component of universal access. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 99)* (pp. 330-337). New York: Association for Computing Machinery Press.
- Jancke, G., Venolia, G. D., Grudin, J., Cadiz, J. J., & Gupta, A. (2001). Linking public spaces: Technical and social issues. In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 530-537). New York: Association for Computing Machinery.
- Johnson, D. G. (1985). *Computer ethics* (1st ed.). Englewood Cliffs, NJ: Prentice Hall.
- Johnson, D. G. (2001). *Computer ethics* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.
- Johnson, D. G., & Miller, K. (1997). Ethical issues for computer scientists and engineers. In A. B. Tucker, Jr. (Ed.), *The computer science and engineering handbook* (pp. 16-26). Boca Raton, FL: CRC Press.

- Johnson, E. H. (2000). Getting beyond the simple assumptions of organization impact [social informatics]. *Bulletin of the American Society for Information Science*, 26(3), 18-19.
- Jones, S., Wilkens, M., Morris, P., & Masera, M. (2000). Trust requirements in e-business. *Communications of the ACM*, 43(12), 80-87.
- Kahn, P. H., Jr. (1991). Bounding the controversies: Foundational issues in the study of moral development. *Human Development*, 34, 325-340.
- Kahn, P. H., Jr. (1992). Children's obligatory and discretionary moral judgments. *Child Development*, 63, 416-430.
- Kahn, P. H., Jr. (1999). *The human relationship with nature: Development and culture*. Cambridge, MA: MIT Press.
- Kahn, P. H., Jr., & Friedman, B. (in preparation). Augmented reality of the natural world and its psychological effects: A value-sensitive design approach.
- Kahn, P. H., Jr., & Kellert, S. R. (Eds.). (2002). *Children and nature: Psychological, sociocultural, and evolutionary investigations*. Cambridge, MA: MIT Press.
- Kahn, P. H., Jr., & Lourenço, O. (1999). Reinstating modernity in social science research—or—The status of Bullwinkle in a post-postmodern era. *Human Development*, 42, 92-108.
- Kahn, P. H., Jr., & Turiel, E. (1988). Children's conceptions of trust in the context of social expectations. *Merrill-Palmer Quarterly*, 34, 403-419.
- Kant, I. (1785/1964). *Groundwork of the metaphysics of morals* (H. J. Paton, Trans.). New York: Harper Torchbooks. (Original work published 1785.)
- Kensing, F., & Madsen, K. H. (1991). Generating visions: Future workshops and metaphorical design. In J. Greenbaum & M. Kyng (Eds.), *Design at work: Cooperative design of computer systems* (pp. 155-168). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kensing, F., Simonsen, J., & Bodker, K. (1998). MUST: A method for participatory design. *Human-Computer Interaction*, 13(2), 167-198.
- Kiesler, S. (Ed.). (1997). *The culture of the Internet*. Mahwah, NJ: Lawrence Erlbaum Associates.
- King, J. L. (1983). Centralized versus decentralized computing: Organizational considerations and management options. *Computing Surveys*, 15(4), 320-349.
- Kling, R. (1980). Social analyses of computing: Theoretical perspectives in recent empirical research. *Computing Surveys*, 12(1), 61-110.
- Kling, R. (1999). What is social informatics and why does it matter? *D-Lib Magazine*, 5(1) [On-line]. Available: <http://www.dlib.org/dlib/january99/kling/oikling.html>
- Kling, R., Rosenbaum, H., & Hert, C. (1998). Social informatics in information science: An introduction. *Journal of the American Society for Information Science*, 49(12), 1047-1052.
- Kling, R., & Star, S. L. (1998). Human centered systems in the perspective of organizational and social informatics. *Computers and Society*, 28(1), 22-29.
- Korpela, M., Soriyan, H. A., Olufokunbi, K. C., Onayade, A. A., Davies-Adetugbo, A., & Adesanmi, D. (1998). Community participation in health informatics in Africa: An experiment in tripartite partnership in Ile-Ife, Nigeria. *CSCW Journal*, 7(3-4), 339-358.
- Kuhn, S., & Winograd, T. (1996). Profile 14: Participatory design. In T. Winograd (Ed.), *Bringing design to software*. Reading, MA: Addison-Wesley.
- Kyng, M., & Mathiassen, L. (Eds.). (1997). *Computers and design in context*. Cambridge, MA: The MIT Press.
- Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts. In W. E. Bijker & J. Law (Eds.), *Shaping technology/building society: Studies in sociotechnical change* (pp. 225-258). Cambridge, MA: The MIT Press.
- Leveson, N. G. (1991). Software safety in embedded computer systems. *Communications of the ACM*, 34(2), 34-46.
- Leveson, N. G., & Turner, C. S. (1993). An investigation of the Therac-25 accidents. *IEEE Computer*, 26(7), 18-41.
- Lipinski, T. A., & Britz, J. J. (2000). Rethinking the ownership of information in the 21st century: Ethical implications. *Ethics and Information Technology*, 2(1), 49-71.
- Louden, R. B. (1984). On some vices of virtue ethics. *American Philosophical Quarterly*, 21, 227-235.
- MacIntyre, A. (1984). *After virtue*. Notre Dame, IN: University of Notre Dame Press.
- Mackay, W. E. (1995). Ethics, lies and videotape. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 95)* (pp. 138-145). New York: Association for Computing Machinery Press.
- Mackay, W. E., & Fayard, A. L. (1999). Designing interactive paper: Lessons from three augmented reality projects. In *Proceedings of IWAR '98* (pp. 81-90). Natick, MA: A. K. Peters.
- Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. New York: Penguin Books.
- Malone, T. W. (1994). Commentary on Suchman article and Winograd response. *CSCW Journal*, 3, 37-38.
- McCarthy, J. F., & Anagnost, T. D. (1998). MusicFX: An arbiter of group preferences for computer supported collaborative workouts. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 98)* (pp. 363-372). New York: Association for Computing Machinery Press.
- Milewski, A. E., & Smith, T. M. (2000). Providing presence cues to telephone users. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 89-96). New York: Association for Computing Machinery Press.
- Milliet, L., Friedman, B., & Felten, E. (2001). Cookies and Web browser design: Toward realizing informed consent online. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2001)* (pp. 46-52). New York: Association for Computing Machinery Press.
- Molich, R. (2001). Ethics in HCI. In *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2001)* (pp. 231-232). New York: Association for Computing Machinery Press.
- Moor, J. H. (1985). What is computer ethics? *Metaphilosophy*, 16(4), 266-275.
- Muller, M. J. (1997). Ethnocritical heuristics for reflecting on work with users and other interested parties. In M. Kyng & L. Mathiassen (Eds.), *Computers and design in context* (pp. 349-380). Cambridge, MA: The MIT Press.
- Mynatt, E. D., Adler, A., Ito, M., Linde, C., & O'Day, V. L. (1999). The network communities of SeniorNet. In *Proceedings of the Sixth European Conference on Computer Supported Cooperative Work* (pp. 219-238). Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Nagel, T. (1986). *The view from nowhere*. Oxford, England: Oxford University Press.
- Nass, C., & Gong, L. (2000). Speech interfaces from an evolutionary perspective. *Communications of the ACM*, 43(9), 36-43.
- Neumann, P. G. (1991). Inside risks: Certifying professionals. *Communications of the ACM*, 34(2), 130.
- Neumann, P. G. (1995). *Computer related risks*. New York: Association for Computing Machinery Press.
- Nielsen, J. (1993). *Usability engineering*. Boston, MA: AP Professional.
- Nissenbaum, H. (1999). Can trust be secured online? A theoretical perspective. *Etica e Politica* [On-line], 2. Available: http://www.univ.trieste.it/~dipfilo/etica_e_politica/199_2/homepage.html

- Noble, D. D. (1991). *The classroom arsenal: Military research, information technology and public education*. London, England: Falmer Press.
- Norman, D. A. (1988). *The psychology of everyday things*. New York: Basic Books.
- Noth, M., Borning, A., & Waddell, P. (2000). An extensible, modular architecture for simulating urban development, transportation, and environmental impacts (UW CSE TR 2000-12-01) [On-line]. Available: <http://www.urbansim.org>
- Olson, J. S., & Olson, G. M. (2000). i2i trust in e-commerce. *Communications of the ACM*, 43(12), 41-44.
- Olson, J. S., & Teasley, S. (1996). Groupware in the wild: Lessons learned from a year of virtual collaboration. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 96)* (pp. 419-427). New York: Association for Computing Machinery Press.
- O'Neill, D. K., & Gomez, L. M. (1998). Sustaining mentoring relationships on-line. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 98)* (pp. 325-334). New York: Association for Computing Machinery Press.
- Orlikowski, W. J. (1993). Learning from Notes: Organizational issues in groupware implementation. *The Information Society*, 9(3), 237-250.
- Orlikowski, W. J. (2000). Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization Science*, 11(4), 404-428.
- Oviatt, S. (1999). Mutual disambiguation of recognition errors in a multimodal architecture. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 99)* (pp. 576-583). New York: Association for Computing Machinery Press.
- Palen, L., Salzman, M., & Youngs, E. (2000). Going wireless: Behavior & practice of new mobile phone users. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 2000)* (pp. 201-210). New York: Association for Computing Machinery Press.
- Parker, D. (1979). *Ethical conflicts in computer science and technology*. Arlington, VA: AFIPS Press.
- Parker, D., Swope, S., & Baker, B. (1990). *Ethical conflicts in information and computer science, technology, and business*. Wellesley, MA: QED Information Sciences.
- Pelto, P. J. (1973). *The snowmobile revolution: Technology and social change in the arctic*. Menlo Park, CA: Cummings Publishing Company.
- Perdue, P. C. (1994). Technological determinism in agrarian societies. In M. R. Smith & L. Marx (Eds.), *Does technology drive history? The dilemma of technological determinism* (pp. 169-200). Cambridge, MA: The MIT Press.
- Perry, J., Macken, E., Scott, N., & McKinley, J. L. (1997). Disability, inability and cyberspace. In B. Friedman (Ed.), *Human values and the design of computer technology* (pp. 65-89). New York: Cambridge University Press.
- Phillips, D. J. (1998). Cryptography, secrets, and structuring of trust. In P. E. Agre & M. Rotenberg (Eds.), *Technology and privacy: The new landscape* (pp. 243-276). Cambridge, MA: The MIT Press.
- Poltrock, S. E., & Grudin, J. (1994). Organizational obstacles to interface design and development. *ACM Transactions on Computer-Human Interaction*, 1, 52-80.
- President's Information Technology Advisory Committee (PITAC). (1999, February 24). *Information technology research: Investing in our future* (Advance Copy). National Coordination Office for Computing, Information and Communications. Washington, DC [On-line]. Available: <http://www.itrd.gov/ac/report/>
- Putnam, R., Leonardi, R., & Nanetti, R. (1993). *Making democracy work: Civic traditions in modern Italy*. Princeton, NJ: Princeton University Press.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press.
- Rocco, E. (1998). Trust breaks down in electronic contexts but can be repaired by some initial face-to-face contact. In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 98)* (pp. 496-502). New York: Association for Computing Machinery.
- Rorty, R. (1982). *Consequences of pragmatism*. Minneapolis, MN: University of Minnesota Press.
- Rosenberg, S. (1977). Multiplicity of selves. In R. D. Ashmore & L. Jussim (Eds.), *Self and identity* (pp. 23-45). New York: Oxford University Press.
- Rotenberg, M. (2000). *The privacy law sourcebook 2000: United States law, international law, and recent developments*. Washington, DC: Electronic Privacy Information Center.
- Roth, A. E. (1990). New physicians: A natural experiment in market organization. *Science*, 250, 1524-1528.
- Sawyer, S., & Rosenbaum, H. (2000). Social informatics in the information sciences: Current activities and emerging direction. *Informing Science*, 3(2), 89-95.
- Scheffler, S. (1982). *The rejection of consequentialism*. Oxford, England: Oxford University Press.
- Schiano, D. J., & White, S. (1998). The first noble truth of cyberspace: People are people (even when they MOO). In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 98)* (pp. 352-359). New York: Association for Computing Machinery.
- Schneider, F. B. (Ed.). (1999). *Trust in cyberspace*. Washington, DC: National Academy Press.
- Schoeman, F. D. (Ed.). (1984). *Philosophical dimensions of privacy: An anthology*. Cambridge, England: Cambridge University Press.
- Scranton, P. (1994). Determinism and indeterminacy in the history of technology. In M. R. Smith & L. Marx (Eds.), *Does technology drive history? The dilemma of technological determinism* (pp. 143-168). Cambridge, MA: The MIT Press.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417-458.
- Searle, J. R. (1983, October 27). The word turned upside down. *New York Review of Books*, 74-79.
- Sharp, L. (1952/1980). Steel axes for stone-age Australians. In J. P. Spradley & D. W. McCurdy (Eds.), *Conformity and conflict* (pp. 345-359). Boston: Little, Brown, & Company. (Reprinted from *Human Organization*, 1952, 11, 17-22.)
- Shneiderman, B. (1999). Universal usability: Pushing human-computer interaction research to empower every citizen. ISR Technical Report 99-72. University of Maryland, Institute for Systems Research. College Park, MD.
- Shneiderman, B. (2000a). Designing trust into online experience. *Communications of the ACM*, 43(12), 57-59.
- Shneiderman, B. (2000b). Universal usability. *Communications of the ACM*, 43(5), 84-91.
- Shneiderman, B., & Rose, A. (1997). Social impact statements: Engaging public participation in information technology design. In B. Friedman (Ed.), *Human values and the design of computer systems* (pp. 117-133). New York: Cambridge University Press.
- Shoemaker, G. B. D., & Inkpen, K. M. (2001). Single display privacyware: Augmenting public displays with private information. In *Proceedings of the Conference of Human Factors in Computing Systems*

- (CHI 2001) (pp. 522-529). New York: Association for Computing Machinery.
- Shweder, R. A., Mahapatra, M., & Miller, J. B. (1987). Culture and moral development. In J. Kagan & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 1-82). Chicago: University of Chicago Press.
- Smith, M. R. (1994). Technological determinism in American culture. In M. R. Smith & L. Marx (Eds.), *Does technology drive history? The dilemma of technological determinism* (pp. 1-35). Cambridge, MA: The MIT Press.
- Smith, M. R., & Marx, L. (Eds.). (1994). *Does technology drive history? The dilemma of technological determinism*. Cambridge, MA: The MIT Press.
- Software engineering code of ethics and professional practice. (1998). [On-line]. Available <http://www.acm.org/serving/se/code.htm>
- Spinello, R. A., & Tavani, H. T. (2001). The Internet, ethical values, and conceptual frameworks: An introduction to cyberethics. *Computers and Society*, 31(2), 5-7.
- Sproull, L., & Kiesler, S. (1991). *Connections: New ways of working in the networked organization*. Cambridge, MA: The MIT Press.
- Stephanidis, C. (Ed.) (2001). *User interfaces for all: Concepts, methods, and tools*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Sturgeon, N. L. (1988). Moral explanations. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 229-255). Ithaca, NY: Cornell University Press.
- Suchman, L. (1994). Do categories have politics? The language/action perspective reconsidered. *CSCW Journal*, 2(3), 177-190.
- Svensson, M., Hook, K., Laakolahti, J., & Waern, A. (2001). Social navigation of food recipes. In *Proceedings of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 341-348). New York: Association for Computing Machinery.
- Tang, J. C. (1997). Eliminating a hardware switch: Weighing economics and values in a design decision. In B. Friedman (Ed.), *Human values and the design of computer technology* (pp. 259-269). New York: Cambridge University Press.
- Thomas, J. C. (1997). Steps toward universal access within a communications company. In B. Friedman (Ed.), *Human values and the design of computer systems* (pp. 271-287). New York: Cambridge University Press.
- Turiel, E. (1983). *The development of social knowledge*. Cambridge, England: Cambridge University Press.
- Turiel, E. (1998). Moral development. In N. Eisenberg (Ed.), *Social, emotional, and personality development* (pp. 863-932). Volume 3 of W. Damon (Ed.), *Handbook of child psychology* (5th ed). New York, NY: Wiley.
- Turiel, E. (in press). *The Culture of Morality: Social Development and Social Opposition*. Cambridge, England: Cambridge University Press.
- Turiel, E., Killen, M., & Helwig, C. C. (1987). Morality: Its structure, functions and vagaries. In J. Kagan & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 155-244). Chicago: University of Chicago Press.
- Turkle, S. (1996). *Life on the screen: Identify in the age of the Internet*. New York: Simon and Schuster.
- Ulrich, R. S. (1984). View through a window may influence recovery from surgery. *Science*, 224, 420-421.
- Ulrich, R. S. (1993). Biophilia, biophobia, and natural landscapes. In S. R. Kellert & E. O. Wilson (Eds.), *The biophilia hypothesis* (pp. 73-137). Washington, DC: Island Press.
- Van House, N. A., Butler, M. H., & Schiff, L. R. (1998). Cooperative knowledge work and practices of trust: Sharing environmental planning data sets. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 98)* (pp. 335-343). New York: Association for Computing Machinery Press.
- Wainryb, C. (1995). Reasoning about social conflicts in different cultures: Druze and Jewish children in Israel. *Child Development*, 66, 390-401.
- Weckert, J., & Adeney, D. (1997). *Computer and information ethics*. Westport, CT: Greenwood Press.
- Weiser, M., & Brown, J. S. (1997). The coming age of calm technology. In P. Denning & B. Metcalfe (Eds.), *Beyond calculation: The next 50 years of computing* (pp. 75-85). New York: Springer-Verlag.
- Weizenbaum, J. (1972, May 12). On the impact of the computer on society: How does one insult a machine? *Science*, 178, 609-614.
- Wiener, N. (1953/1985). The machine as threat and promise. In P. Masani (Ed.), *Norbert Wiener: Collected works and commentaries* (Vol. IV, pp. 673-678). Cambridge, MA: The MIT Press. (Reprinted from *St. Louis Post Dispatch*, December 13, 1953.)
- Williams, B. (1985). *Ethics and the limits of philosophy*. Cambridge, MA: Harvard University Press.
- Winner, L. (1986). *The whale and the reactor: A search for limits in an age of high technology*. Chicago: University of Chicago Press.
- Winograd, T. (1994). Categories, disciplines, and social coordination. *CSCW Journal*, 2(3), 191-197.
- Zheng, J., Bos, N., Olson, J., & Olson, G. M. (2001). Trust without touch: Jump-start trust with social chat. In *Extended Abstracts of the Conference of Human Factors in Computing Systems (CHI 2001)* (pp. 293-294). New York: Association for Computing Machinery.