DOCUMENT RESUME

ED 360 947 IR 016 175

AUTHOR Kahn, Peter H., Jr.; Friedman, Batya
TITLE Control and Power in Educational Computing.
PUB DATE Apr 93
NOTE 35p.; Paper presented at the Annual Meeting of the American Educational Research Association (Atlanta, GA, April 12-16, 1993).
PUB TYPE Viewpoints (Opinion/Position Papers, Essays, etc.) (120) -- Reports - Evaluative/Feasibility (142) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Computer Assisted Instruction; Computer Centers; Computer Science Education; Computer Software Development; *Constructivism (Learning); Educational Practices; Educational Technology; Elementary Secondary Education; *Epistemology; Higher Education; Knowledge Level; Man Machine Systems; *Personal Autonomy; *Power Structure; Users (Information)
IDENTIFIERS *Control (Social Behavior); Postmodernism

ABSTRACT

Educational computing based on the primacy of human agency is explored, considering ways in which ower can be apportioned and exercised in order to enhance educational computing. Ideas about power and control are situated epistemologically. A first consideration is educating for human control of computer technology. Research suggests that children and some adults attribute agency to computer technology, and much research in computer science comes from the belief that computer technology can mimic, if not duplicate, human agency. It is argued that people should control technology and take responsibility for the consequences of computer-mediated action, and control should be minimized if not eliminated from power relations between people. To promote this, a constructivist theory of education allows for the epistemological stance that the authors articulate. The principles of constructivist education call for educational software that supports students' construction of knowledge and foster interest, autonomy, and cooperation. Technologies and approaches in this direction are explored. A constructivist approach will avoid one of the pitfalls of deconstructionism, the emphasis on power that is often the only recourse of the deconstructivist when injustice occurs. It is suggested that postmodern epistemology be set aside, but that the best of what remains of postmodern thinking and practice be used to promote democracy in educational computing. (Contains 103 references.) (SLD)

Control and Power in Educational Computing

Peter H. Kahn, Jr. and Batya Friedman

Colby College

RUNNING HEAD: CONTROL AND POWER IN EDUCATIONAL COMPUTING

Authors' addresses. Peter H. Kahn, Jr., Department of Education and Human Development, Colby College, Waterville, ME 04901. Electronic mail: phkahn@colby.edu. Batya Friedman, Department of Mathematics and Computer Science, Colby College, Waterville, ME 04901. Electronic mail: b_friedm@colby.edu.

Control and Power in Educational Computing

> The computer center is overseen by a student group, the Computer Users Society. Members have keys to the room, and may gain access at any time, day or night, weekends. One Sunday morning, Harvey [the high school teacher] recalls, he came in around 10 a.m. thinking he would get some work done while the place was quiet. He could not find a free terminal. (Dormer, 1981)

> The Computer Center User Society...holds regular meetings to decide the policies and rules of the computer program. Managing the computer system are six students, called "superusers," who oversee the use of the computers... [One student] explained that although superusers have access to everything, by the time a student gets to the point of being a superuser, immoral actions such as looking at grades are "the last thing we'd have to worry about. This place is so special...it would be such a downfall that no one would want to risk it." (Appel, 1985)

These educational computing practices former part of the computer education that Harvey (1980, 1983) initiated in an otherwise traditional high school. In this paper, we want to say more about the direction Harvey points. It is toward student participation, cooperation, and interest, and a school culture that is imbued with using computer technology to foster democratization. It is toward educational computing where students act on and control the technology, and through their involvement build ethical community.

In addition, we seek to place such educational computing practice within a robust theoretical framework. Toward this end, we distinguish between control and power. In our use of these terms, control occurs when people direct coercively another agent, or direct the action or process of a non-agent (e.g., of a machine). The key here is that within social relations control is usually unethical because it undermines others' autonomy. Power, in turn, occurs only in social relations. It can be coercive, and in such cases power refers to how people control others. But power can also refer to non-coercive relationships within which people influence, organize, and lead others.

This distinction between control and power helps us to advance several ideas. One is that people should control objects, not other people. Another is that objects -- specifically computers -- cannot control people, and yet all too often inappropriately appear as if they do. More broadly, we stake out an account of educational computing based on the primacy of human agency. Our account presumes that there is a self (a morally responsible agent) who makes meaning in the world, and whose knowledge and values are shaped by, and in turn shape the larger society. Within this account, institutions and larger social forces exert power over individuals: teachers exert power over students, principals over teachers, school districts over schools, and so it goes through various interwoven hierarchies on the state and national level. Thus we aim to provide means by which power can be apportioned and exercised in ways that enhance educational computing. Finally, we situate our ideas about control and power epistemologically. Suffice for now, our ideas are constructive rather than deconstructive, and modern as opposed to postmodern.

Section I. Educating for Human Control of Computer Technology

Consider the following real-life story. A professor sought to make an airline reservation, and called the airline company to charge her ticket. In the process of writing the ticket, the reservationist requested the professor's home phone number. Now, it happened that the professor was between academic jobs, in transit for the summer, and without any established phone. Though the reservationist listened sympathetically to this explanation, he insisted that he have a phone number: "Our computer program has to have a phone number or it won't work. I can't even get to the next screen. It's the computer's fault." How many times have we heard that? "It's the computer's fault." Well is it? Can computers be at fault?

Such questions arise because computer technology often appears volitional and intelligent. Computer systems, for example, can "track" credit histories, and "decide" to reject loan applications. Computer guided missiles, as noted by Dawkins (1976), often appear to "search actively" for their target, and "predict" and "anticipate" the target's evasive moves. Medical expert systems "diagnose" illness and "recommend" cures. In terms of educational software, intelligent tutors "decide" when

problems a student will work on, and "correct" students when they are wrong, and "judge" when a student is ready for more advanced problem sets. Tracking. Deciding. Searching. Anticipating. Recommending. Tutoring. Correcting. Judging. Such terms would seem to imply that we believe computers are not so different from humans, and that they have, to varying degrees, intentional states such as thoughts, desires, consciousness, free will, and the capability to make intelligent decisions. How did we come to believe such things about computers?

One response is that we do not believe such things. Another response is that we believe such things because they are true. Both responses are worth our attention. In the first, it might be granted that in conversations we readily talk about "intelligent tutors" or "smart missiles". Or, like the airline reservationist, we might say "it is the computer's fault". But, it could be argued, what appears as sincere instantiations of attributing agency to computational systems are nothing more than superficial verbal responses, and that people do not really conceive of computers as human or human-like.

What little research data exists, however, suggests otherwise. For example, Turkle (1984) interviewed children about their experiences with an interactive computer game called Merlin that plays Tic-Tac-Toe. Her findings indicated that children attributed psychological characteristics to Merlin: that Merlin, for example, was capable of cheating. Rumelhart and Norman (1981) report on research with adults that similarly show some attribution of agency to computational systems (cf. Weizenbaum, 1976). More recently, Friedman (in preparation) has sought to examine directly whether people reason about computers as moral agents. Computer science majors from a research university were interviewed in three general areas: (1) Students' views of computer agency. (2) Students' assessments of computer capabilities. And (3) students' judgments of moral responsibility in two situations that involved delegation of decision-making to a complex computer system. One situation involved an automated computer system that evaluates the employability of job seekers, and rejects a qualified worker. The second situation involved an automated computer system that administers medical radiation treatment, and over-radiates a cancer patient. Preliminary results show that while virtually all students (98%) judged one or more human agents (e.g., computer programmers, computer operators, administrators) to

be morally responsible for computer error, less than a fifth of the students (17%) blamed the computer. It was also the case, however, that most students (83%) viewed computers as capable of making genuine decisions or of having intentions, albeit all students (100%) viewed such decision-making and intentions as unequal to that of humans. From in-depth qualitative analyses of students' reasoning for their views, two overarching orientations emerged. In the first orientation (Quantitative), an appeal is made to the magnitude of a particular characteristic found in both computers and humans (e.g., to the amount of knowledge a human or computer possesses or is capable of possessing). In the second orientation (Qualitative) an appeal is made to the presence or absence of qualities possessed by computers and humans (e.g., to emotions, desires, beliefs, consciousness, originality, and free will). Thus, initial evidence suggests that children and adults in certain respects do attribute agency to computer technology, and that such attributions go far beyond superficial use of language.

The second response comes from a different direction. Here it would be claimed that the reason computer technologies appear to have aspects of human agency is quite simply because they do. Computers, for example, appear smart because they are smart, and are capable of having beliefs and intentions, and of making decisions. Historically, the military has been committed to this view, philosophically and financially. The military has sought to develop "smart" and "autonomous" computer-based weaponry, and to use computers to train military personnel. In turn, the military has had an enormous impact on shaping the fields of computer science and educational computing. For example, according to Thomborson (cited in Nobel, 1991, p. 13) 70 per cent of all academic research in computer science has been funded by the Department of Defense. Similarly, according to Noble (1991) "military agencies have provided three-fourths of all funding for educational technology research over the last three decades, and within government agencies, the military spends seven dollars for every civilian dollar spent on educational technology research" (p. 2) Thus, most of the research in computer science and most of our nation's computer-based educational applications comes out of a tradition that believes that computer technology can mimic, if not duplicate, human agency. Is such a belief warranted?

Perhaps the most sustained and incisive critique of the idea that computers can be intelligent in the sense of having intentional states has been advanced by Searle (1981, 1984. 1990, 1992). While this is not the place to go in depth into Searle's thinking, it is worthwhile to review the skeleton of his well known Chinese room argument. Through it, we stake out our philosophical position toward technology, which, in turn, provides a cornerstone for our educational practice. Searle's thought experiment goes something like this: Imagine that you do not understand Chinese writing or speaking, and you are put in a room with a basketful of Chinese symbols. You are then given a rule-book in English that tells you how to match up certain symbols with other symbols, based only on their visual configuration. For example, a double-bent squiggle might get matched with a half-bent squoggle. Now, imagine a Chinese person from outside the room writes you a question to try to determine if you know Chinese. You receive the written question, which to you are only meaningless Chinese symbols, and according to the rules of your English rule book, you match up the symbols with other appropriate symbols, and then send out the response. With a good rule book, the Chinese person will not be able to tell the difference between your answers and the answers given by a Chinese-speaking person in that room. The outside questioner might ask with Chinese squiggles, "Would you like a Big Mac or a Whopper?" You might respond with Chinese squoggles, "A Whopper, for I find Big Macs just a tad too greasy." But now we ask the question, would you understand what you just said? Or, more generally, when you correctly manipulate the symbols, do you understand Chinese? Searle argues, absolutely not. After all, when you send out an answer you do not know what you are sending out. All you are doing is matching symbols with other symbols based on the rule book. That is, there is nothing in the formal symbol manipulation (the syntax) that provides the understanding of Chinese (the semantics).

So, too, with computers. Because computational systems are purely formal (syntax), and because purely formal systems have no means to generate semantics, Searle argues that computational systems do not have the properties that are central to human agency. This is not to say that human-like agency might not someday be realized in material or structures other than biological brains. It is to say that computers as we can conceive of them today are not such material or structures. Thus while

humans may (sometimes with valid justification) delegate decision-making to computer technology, fundamentally it is humans who control technology, and are responsible for the consequences of its use.

In the last handful of years, some members of the Artificial-Intelligence (AI) community appear to have backed off from their strong philosophical claims. (See, for example, the editorial by Chandrasekaran, 1992, Editor-in-Chief of the artificial intelligence journal of the IEEE.) Their current position goes something like: Regardless of whether or not computers actually have intentionality, computers act as if they do, and thus we can rightly design and interact with them as if they do. In other words, on the one hand these AI researchers give up a good deal of philosophical ground to those like Searle. On the other hand, they keep their research agenda fully in place -- to build intelligent-like computer systems. But if computers continue to be designed as if they have agency, and if in their interactions with computers people act as if computers have agency, then it seems likely that people will think so as well. Thought and action cannot be so easily compartmentalized. Thus we have argued elsewhere (Friedman & Kahn, 1992) that increasingly sophisticated computer systems should be designed not to mimic human agency, but to support it. For example, computer interfaces should not be designed to intercede in the guise of another "agent" between human users and the computational system, but to "disappear" such that the user is freed to attend directly to, and take responsibility for, the tasks at hand (see, also, Shneiderman, 1987, 1989; Winograd and Flores, 1986).

Here is where we have been heading. We suggest that children and adults, in various ways, mistakenly attribute agency to computer technology. And we say mistakenly because it is our claim that computer technology does not have agency. Moreover, it is our claim that (a) people should control technology, and take responsibility for the consequences of computer-mediated action, and (b) control should be minimized if not eliminated from power relations between people.

To develop this distinction between control and power a little more fully, consider two ways people can exert power through computers. Following our definitions, one way is coercive, and thus reduces to control. For example, drawing on Arnstine (1973), imagine the hijacker who points a gun at an airline pilot and says: "Fly me to Havana or I'll blow your brains out." The gun could be said to

coerce the pilot. After all, without the gun the threat would have little force. However, contrary to some recent trends that view objects as having a social life (Appadurai, 1988), we maintain that the gun is fundamentally an inanimate tool mediated by a person (see also Scheffler, 1991, chap. 8): in this case a tool used for the purpose to establish control over another person. In this explicit psychological sense, we agree with the statement, "guns don't kill people, people do." It is also the case, however, that once people build tools, the tools have built-in features such that people usually use the tools in certain ways, toward certain ends (cf. Bromley, 1990). Thus with a gun, or with a computer-driven weapon, people often exert power to control other people.

But people can also exert non-coercive power over others by using computer technology. Consider what is involved in controlling air traffic at airports (at least when the method of hijacking is shunned). We speak of a control tower, where incoming and outgoing airplanes are regulated through use of sophisticated computer technology. Of course, the physical tower does not control the planes. But neither do the computers housed within it. Rather, through computers, people control the planes, and thus indirectly exert non-coercive power over passengers. Moreover, because in all such situations computers exert no control and have no power, the responsibility for computer generated mistakes lies, ultimately, with people: perhaps with the programmers who wrote the software, or with the administrators or other individuals who chose to implement the particular software, or with the airport controllers who operate the computer technology.

In terms of educational computing, this distinction between control and power leaves us with several questions. How do we teach students that humans, and often the students themselves, control computer technology, and are responsible for the consequences of computer-mediated action? How do we teach future computer scientists to design systems that foster such understandings? Finally, how do we teach students that by controlling computer technology -- and choosing wisely -- people have power to effect meaningful and ethical change, in educational settings and beyond?

Our answers build from a constructivist account of education. Thus before highlighting specific computing activities, we should like to say a few words about this perspective by elaborating on four

constructivist principles proposed and practiced by DeVries (1988; DeVries & Kohlberg, 1990). These principles are part of a larger social-cognitive research tradition (e.g., Arsenio, 1988; Damon, 1977; Kahn, 1992; Killen, 1990; Kohlberg, 1969; Laupa, 1991; Nucci, 1981; Piaget, 1932/1969; Selman, 1980; Smetana, 1982; Thorkildsen, 1989; Tisak, 1986; Turiel, 1983; Wainryb, 1991; Youniss, 1980). These principles are also, in some respects, compatible with other educational theories, for example, those that are progressive (Dewey, 1916/1966) and experiential (Wigginton, 1986). We prefer, however, the constructivist label for two reasons. First, it highlights that -- no matter how social the discourse, elaborate the organizational structure, or co-constructive the reciprocal interaction -- fundamentally it is a self (a morally responsible agent) who makes meaning in the world. Second, for reasons that will become clear in the third section, constructivism allows best for the epistemological stance we would wish to articulate. The four constructivist principles can be summarized as follows:

From Instruction to Construction. Many people believe that for students to learn, teachers must instruct, by which it is meant that learning depends on a teacher who correctly sequences curriculum content, drills students on correct performance, corrects mistakes, and then tests for achievement. Granted, one might note a few sidewise embraces of critical thinking and cognition. But, if push comes to shove -- if, for example, test scores go down -- the call is clear. Back to basics. Instruction in the 3-R's. In contrast, in the move from instruction to construction, learning involves neither simply the replacement of one view (the incorrect one) with another (the presumed correct one), nor simply the stacking, like building blocks, of new knowledge on top of old knowledge, but rather transformations of knowledge. Transformations, in turn, occur not through the child's passivity, but through active, original thinking. As Baldwin (1897/1973) says, a child's knowledge "at each new plane is also a real invention....He makes it; he gets it for himself by his own action; he achieves, invents it" (p. 106). Think of it this way. On a daily level, children encounter problems, of all sorts: logical, mathematical, physical, social, ethical. Problems require solutions. The disequilibrated state is not a comfortable one. Thus the child strives toward a more comprehensive, more adequate, means of resolving problems, of synthesizing disparate ideas, of making sense of the world. Constructivist education, therefore,

centrally involves experimentation and problem solving, and student confusion and mistakes are not antithetical to learning, but a basis for it.

From Reinforcement to Interest. Traditional educators often seek to shape student behavior through four types of reinforcement procedures. Perhaps the most effective is positive reinforcement, wherein a positive stimulus (e.g., a high grade, or for younger children a gold star or sticker) is administered to increase a behavior (time spent on task). Through punishment, an aversive stimulus (being sent to the principal's office) is administered to a decrease a behavior (hitting another child). Through response cost, a positive stimulus (free time to play) is decreased to decrease a behavior (yelling in class). Finally, through negative reinforcement, an aversive stimulus (time spent on a spelling assignment) is decreased (by an accomplished student finishing the assignment quickly, and then having free time) to increase a behavior (to get the student to complete the spelling assignment). All of these procedures share a common feature. They build on a conception that children learn through stimulus-response conditioning, and that for effective instruction the teacher needs to strengthen, weaken, extinguish, or maintain learned behaviors through such reinforcement procedures. In contrast, as we have argued, humans have not just syntax, which leads to a conception of the child as a computer, to be programmed, but semantics, meaning. It follows that a child constructs meaning more fully when engaged with problems and issues that centrally captivate his or her interest. Thus, from a constructivist perspective, teachers find out what interests their students, and build curriculum to support and extend those interests. They give students some power to shape curriculum, and freedom to explore, to take risks, to make mistakes. Indeed, it can be argued that many of the behavior problems that traditional teachers try so hard to control arise precisely because students find the curriculum drudgery.

From Obedience to Autonomy: Construction and interest do not thrive in an environment in which the teacher is the authority demanding obedience. Moreover, obedience leads to conformity, and to the acceptance of ideas without understanding. Thus, from a constructivist perspective, teachers should move away from demanding obedience, and toward fostering the child's autonomy. Now, by autonomy we mean in part something like independence from others. For it is only through being an

independent thinker and actor that a person can refrain from being unduly influenced by others (e.g., by Neo-Nazis, youth gangs, political movements, and advertising). Autonomy in this sense is a prerequisite for agency, and thus necessary for an individual to control technology. But by autonomy we do not mean a divisive individualism, as constructivist autonomy is often said to be (Hogan, 1975; Shweder, 1986). Rather, within a constructivist framework (Baldwin, 1897/1973; Kohlberg, 1969, 1984), autonomy is highly social, developed through reciprocal interactions on a microgenetic level, and evidenced structurally in incorporating and coordinating considerations of self, others, and society. In other words, the social constrains or bounds the individual(ism), and vice-versa.

From Coercion to Cooperation: In some sense, the movement from coercion to cooperation reflects the flip side of obedience to autonomy, but more from the student's and not the teacher's standpoint. Like for autonomy, cooperation entails incorporating and coordinating one's own feelings, values, and perspectives with those of others. Given that the adult's relationship to children are laden (often necessarily so) with coercive interactions, peer relationships are centrally important. Through them, concepts of equality, justice, and democracy flourish (Piaget, 1932/1969), and academic learning is advanced (Vygotsky, 1978).

Section II: Educational Computing in the Classroom and Beyond

Based on these four constructivist principles, much is possible in terms of educational computing. To begin, educational computing practice can be enhanced by well designed educational software. But to be well designed from a constructivist perspective, such software needs to steer clear of authoritatively directing student learning, as occurs in a good deal of computer-assisted instruction. Instead, educational software should support students' construction of knowledge, and foster interest, autonomy, and cooperation.

Computer simulations provide one means. Simulations allow students to explore hypothetical situations from multiple perspectives, and to follow different courses of action of each. For example, in the computer simulation Our Town Meeting (Snyder 1985a), students working in groups are assigned to

one of three town agencies, each with its own agenda. Through the simulation, students find that cooperative strategies often provide the most effective means for resolving disputes and achieving agency-specific goals. Cooperation is also encouraged in another computer simulation, The Other Side (Snyder, 1985b), where students work in teams to build the economy of their own country while working toward building a bridge between their own and a rival country.

Computer construction kits provide another means for advancing constructivist computing. Such a kit presents a student with the opportunity to design parts of the simulation. For example, one of the early kits for the popular market, Pinball Construction Set (Budge, 1983), provides the user with a generic pinball game where the user defines and assembles many of the components. The user chooses, for example, how many flippers and stoppers, and where they should be placed on the simulated pin ball machine, the elasticity of the walls, and so forth. The user can then play the game based on each customized game board, and assess the merits of the design. Such kits offer a great deal of educational promise as they begin to be created with more academic content. For example, Interactive Physics II (Baszucki, 1992) allows the user to access a large number of physical components (springs, ropes, blocks, and disks) and properties (forces and gravity) to model and investigate two-dimensional physical systems. Through such construction kits, a student engages actively in a design, receives immediate data by which to judge its success, and then is positioned to rework a better design, and thus continue the generative process.

A constructivist-oriented teacher can also engage students in any number of student-directed collaborative projects that make use of computer technology. For example, consider a student initiated and run newspaper. In this (potentially school-wide) endeavor, computing tools support freedom to generate, share, and argue about ideas. Word processors and desk-top publishing tools support students' writing, spelling checkers its presentation. Spreadsheets can help students to manage the financial aspects of running the newspaper, databases to organize and access relevant information. Thus, through publishing a newspaper, students learn how to use a wide range of computer applications. In addition,

students learn that they control computer technology, and are responsible for the consequences of their computer-mediated action: for their published ideas.

Electronic mail provides an even wider community for discourse (cf. Cummins, 1988; Horowitz, 1984). Imagine, for instance, students in an inner-city class conducting a science experiment, in consort with students from a rural setting. Using electronic mail, they could perhaps decide to conduct tests for local water or air quality. Then, through using electronic mail to share their data and interpretations, they would be able to construct a wider understanding of how environmental problems often transcend their geographical locations. Or, building from the idea of a school newspaper, students from two geographically distant areas might write editorials on the same national issues. These editorials, which may well reflect students' regional views, could be printed side by side in each newspaper. In ways like these, computer technology can be used to enhance communication and understanding between students who might differ on the basis of geography, culture, race, and economic standing.

It is precisely this role of computer technology that is at work globally, in a political context. For example, due to repressive political forces and political upheaval during 1991 in the former Soviet Union, electronic networks served as one of the few means for individuals to communicate with others in outside counties (Goodman, 1992). Or note the story that Alexander Randall tells of his experience walking the streets of places like Warsaw, Moscow, and Kiev during their recent struggles for independence (Lewis, 1992). Randall talked with people in various organizations, with dissidents, and with mainstream journalists. He asked, "What is the next step to make sure that freedom of speech and free government survive?" Their answers surprised him. They all said they wanted desktop publishing equipment. It dawned on him. "What is freedom of the press if you don't have the press?" For such reasons, Randall founded the East-West Foundation. This non-profit organization solicits used (and by many standards outdated) personal computers from Westerners and provides them free of charge to dissidents and journalists in the Eastern European countries. In Randall's words: "Our old XT's and AT's and Mac 512's are awesome tools to people who don't have them...The fall of Communism did not

axiomatically mean the rise of democracy. This is our people's response to insure that the forces trying to make democracy happen will ultimately succeed" (p. F12).

In these sort of ways, people can control computer technology to enhance political freedom. If we accept, however, that children develop understandings about political freedom, and of how to create societies in which such freedoms thrive, then it follows that such understandings and practice need to be made an integral part of children's education. Toward this end, we suggest that schools be organized to increase student self-governance, wherein students determine many of the policies that regulate their own classroom and school activities, and thus gain experience with democratic and consensus decision-making processes. Self-governance within entire schools has been extensively discussed and successfully practiced by others (e.g., Kohlberg, 1980, 1984, 1985; Power, Higgins, & Kohlberg, 1989). Moreover, Friedman (1986, 1991) provides many suggestions for how self governance can readily carry over into the computer classroom, and how student decisions about the uses of computer technology can be bounded by larger societal and ethical considerations. A similar focus is found in the educational practices quoted at the outset of this paper that Harvey initiated in his computer center. Recall that students held regular meetings to decide the policies and rules of their computer center, and checked themselves from unethical action (e.g., from using the computer to access other students' grades). Students also chose what computer languages they wanted to learn, and largely held themselves accountable to high academic standards. Harvey (1980) writes that educational freedom means first of all that students can make significant choices, not trivial ones invented by a teacher. Self-governance allows for such significance, and thus embraces an educational freedom that seems to us but part and parcel of political freedom.

Investigating the impact of technology on individuals within organizations provides another means to broaden the scope of educational computing. Here is an idea for a semester project. Have students choose an organization -- McDonald's, a bank, a manufacturing plant, the school attendance office -- most any will do, and have them research how the computer technology was brought into the organization, and how it supports the organization's goals, and shapes social interactions. Students can

collect "data" through observations, surveys, and extended interviews. Their data can inform on such questions as: Who decided that computer technology would be good to have in the organization? Who decided on the specific hardware and software? Were workers consulted? What sort of needs was the computer technology to meet? How well have such needs been met? How did workers get training with the technology? How do management and workers feel about the technology's current use? Were there any unanticipated consequences from using the technology?

Such projects were successfully conducted by undergraduate students at our own institution. One group of students, for example, in an administrative science course, elected to study computer use in a small accounting firm. The firm used a computerized tax preparation program chosen for its ease of use and compatibility with existing hardware. In addition to these features (and unnoticed by the firm members at the time of purchase), the software kept a running tally of the number of tax forms completed by each accountant. This information on each accountant was revealed at the end of the tax season. Office discord followed. In response, the firm made a collective decision to "hide" this unsolicited information in subsequent years. Through reflecting on findings like these, students can identify situations where workers feel either controlled or supported by technology, and be better positioned to propose changes for the work place.

Many of our ideas for educational computing carry over to computer science education: to educating those students who in the coming years will design and build the computational systems that shape our computing environments. One type of activity, for example, would be to give a group of computer science students the goal of writing one large collaborative computer program. In this context, which is akin to that of a business environment, students must decide what program modules to create, who will take major responsibility for each module, how to fit the modules together, and, once together, how to remove the errors. Thus, to be successful, students must be not only capable intellectually, and able to assert themselves on intellectual ground, but able to discuss, cooperate, and build with others.

Educating computer science students poses the additional challenge of teaching these students about the social responsibilities that come with their technical expertise. Teaching in this area has

received increasing attention in the literature (Bynum, Maner, & Fodor, 1992; Denning, 1992; Dunlop & King, 1991; Ermann, Williams, & Gutierrez, 1990; Friedman & Winograd, 1990; Gotterbarn, 1992; Johnson, 1985; Miller, 1988; Parker, Swope, & Baker, 1990; Perrolle, 1987; Winograd, 1992). It has also become part of the core curriculum recommended by leading computing organizations ("A Summary of the ACM/IEEE," 1991), and by which computer science departments are judged for national accreditation. Thus we wish to describe several ways of integrating social concerns into students' computer design experiences, such that when students define and implement computer systems, standard issues include not only technical ones (e.g., "What data structure should I use?"), but social ones (e.g., "How does my system impact the intended users?").

Consider one example of a programming assignment to highlight issues of bias -- systematic unfairness to individuals or groups. Following recent analyses by Friedman and Nissenbaum (1992, in preparation) on bias in computer system design, three overarching categories of bias were identified: preexisting social bias, technical bias, and emergent social bias. Preexisting social bias occurs when computer systems embody biases that exist independently of, and usually prior to, the creation of the software. For instance, a legal "expert" computer program was written to offer advice to immigrants seeking citizenship in Britain. Some have argued (Berlins & Hodges, 1981), however, that the British immigration laws are themselves biased against certain nationalities and people of color. To the extent they are, such biases also became embedded in the expert system. In contrast to preexisting social bias, technical bias occurs in the resolution of technical design problems that often arise due to limitations of the programming language or algorithm. For instance, in the above legal expert system, the programming language, Prolog, sometimes went into infinite loops and thus failed to prove theorems that were logically implied by the axioms. Due to this technical limitation of the programming language, it follows that the system systematically would fail to identify individuals who were otherwise entitled to British citizenship. Finally, 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. For instance, since the early 1970's, the computerized National Resident Medical Match Program

has placed most medical students in their first jobs. In the system's design, programmers assumed that only one individual in a family would be looking for a residency, and programmers thereby inadvertently discriminated against married residents. At the time, such an assumption was perhaps not out of line since there were few women residents. But as women have increasingly made their way into the medical profession, marriages between residents are now not uncommon. Such bias against married residents only emerged when the social conditions changed.

In the programming assignment to highlight bias, students in a course on data structures were asked to design and implement a computer dating program. The technical material focused on the use of linked lists. The issue of bias arose when students determined issues like who would be included in the database, and how individuals in the database would be searched. Some students, for example, assumed only heterosexual users, and their programs were critiqued by other students on the basis that the program resulted in the unfair exclusion of homosexuals, due either to oversight or to the programmer's preexisting bias against homosexuals. Other programs searched for matches with a first-entered first-searched strategy, and thus unfairly favored those individuals who joined the database earlier over those who joined later. This instance of bias commonly arose because students used a linear linked list. To remove the bias, some students redesigned the program with a circular linked list. Our point here is that the social import of system design, and the designer's role and responsibility for the design, can emerge compellingly from students' own design experience.

Moreover, from a constructivist perspective, the social dimensions of computing become more sensible to students when they design computer software for use by real people in real settings. Such designing makes salient the need to take seriously debugging, prototyping, field testing, interface issues, and user satisfaction. Based on our experience, students have successfully pursued a variety of design projects. For example, one student wrote a customized program for helping a baseball coach keep track of his team statistics. Another student built an interactive videodisk of selected artwork and performances for an art history teacher. Yet another student designed a program for the college's radio station

that would identify the daily records played, and the station's top ten hits. Other students using HyperCard worked cooperatively to design a guided tour of the college's library.

Overall we have offered activities that are broadly constructivist. They seek to move education from instruction to construction, reinforcement to interest, obedience to autonomy, and coercion to cooperation. Moreover, the activities point to two broad ways by which people can exert control over computer technology. One way is as the user, wherein people control technology by means of what they choose to use, and how they choose to use what they use. Thus we described the use of computer simulations, electronic mail, computer applications, and so forth. A second way people control technology is as the designer, by means of what people create technologically. Thus we focused as well on the education of computer science students. Many of our activities, of course, encourage users to think about design considerations, and designers to consider the social impact of their work. Our goals here fit within a broader conception of participatory design (Bodker, 1991; Greenbaum & Kyng, 1991; Namioka & Schuler, 1990) which, in its robust form, seeks not only to create better designs sensitive to the users' needs but to enfranchise and validate users in the design process. Through such participation, the idea of the "technocrat" as keeper and definer of the technology gives way to a more embracing conception of the designer and user as not only complementing one another, but essentially linked to help create the environment within which we live and work.

We are well aware that educational computing as we have described can be hindered by the coercive power wielded by organizations beyond the immediate scope of any single individual or school. When this occurs, other organizations need to respond. Here we just want to suggest that effective organizational response often depends on being able to defend sound educational principles. For example, in choosing the standard by which to measure academic computing achievement among high school students, the Educational Testing Service (ETS) selected a single programming language, Pascal, for the Advanced Placement Exam in Computer Science. There is undoubtedly an interesting story about why ETS chose specifically Pascal. But what we find especially remarkable is that ETS chose only one programming language. After all, in the field of computer science there is no standard

programming language, in the way English is the standard spoken and written language in the United States. For example, on the college level, computer science departments often choose C, Scheme, or Modula-2 instead of Pascal. In turn, in advanced high school computer study, teachers and students with interest in artificial intelligence would find most useful the languages of LOGO or LISP, while those interested in operating systems with access to UNIX would find most useful C. Thus why only Pascal on the Advanced Placement Exam? We suspect that part of the answer hinges on the increased costs that arise from creating and administering exams in multiple languages. But, whatever the reason, the end is the same. ETS coerces high schools into teaching Pascal, and precipitates an impoverished view of what constitutes advanced computing at the high school level.

How might the educational community address this problem? Minimally, those in secondary and higher education need to discuss these problems with ETS, such that a choice of programming languages is provided on the Advanced Placement Exam. But, perhaps this process, even if it is successful, misses the larger point for it still compels high schools to teach to standardized tests if their students are to be competitive in their college applications. Another solution is possible, which would also solve the problem that arises if ETS is unable or unwilling to add additional programming languages to the Advanced Placement Exam. Namely, in their admissions process, colleges and universities can go beyond standardized testing to more robust and valid ways of assessing the depth, scope, and integrity of student's high school computer science education. In this way, standards still matter, as they should. Excellence counts. Indeed, excellence thrives because high school educators and students gain increased freedom to make significant choices to pursue advanced computer science as best suited to their interests, talents, and resources.

It is sometimes said that computer technology depersonalizes social life, mechanizes it, and reduces the voice and importance of each individual. If what is said is true, we hope our constructivist activities can go some distance toward changing the course: to give people control over the technology in and around their lives, and the power to effect meaningful and ethical change.

Section III: Control, Power, and Modern Epistemology

Our ideas in this paper have been advanced by focusing on the constructs of control and power. But, at this point, something needs to be said about these constructs from an epistemic standpoint, for both terms often get appropriated by postmodern thinkers. Such appropriation leads to problems that trouble us, and if left unattended can undermine educational computing. Allow us then some liberty as we first sketch what is at stake with postmodern theory. Then we shall be able to turn the discussion back to that of educational computing.

Postmodern theory, of course, is a broad term, and cuts across many disciplines, and means different things to different people. Yet amidst such broad territory, two types of postmodern theories can be characterized on the basis of their epistemic claims: deconstructivist postmodernism and affirmative postmodernism. Deconstructivism (e.g., Culler, 1982; Derrida, 1978; Foucault, 1983; Morss, 1992; Norris, 1982; Scholes, 1989) sharply calls into question traditional scientific and western assumptions. The theory holds that since all knowledge is a product of a particular person, in a particular culture, at a particular point in time, statements about knowledge cannot transcend their particularistic origins. What is true or ethical for one culture (or subculture) may not be true or ethical for another. Thus deconstructivism limits knowledge claims to first person statements. One could say, for example, "I believe it is raining today," or "I believe it is immoral for an employer to discriminate against people of color." But in all such judgments, one cannot appeal to a third-party, independent, or privileged perspective or process. We cannot appeal, for example, to rationality. We cannot say, "if one is rational, morality demands that we not discriminate on the basis of color." We cannot, because rationality -- like objectivity, necessity, logic, causality, responsibility, and truth -- is a human construct, a product of particular people, in a particular culture, at a particular point in time.

The impetus for deconstructivist theory is understandable. Great injustice and tragic harm has resulted from people and cultures who have thought themselves privy to truth. But does deconstructivism solve the problem? Any positive answer would first need to find its way through at least three internal contradictions embedded in deconstructivist theory. First, deconstructivists argue against theory

building, and yet themselves advance a theoretical position. Second, deconstructivists seek to deconstruct the tools of logic, reason, and rationality, and yet they seek to do so with those very tools. Third, deconstructivists argue against privileging any position. Yet, if their theory (that holds that no theory can be true for everyone) holds for everyone, even for the person who mistakenly believes it false, then the theory does what it says cannot be done. It privileges itself. It establishes some basis for truth that transcends its own confines. (For a discussion of these and related issues, see, e.g., Crews, 1986, 1989; Hoy, 1985; Kahn 1991; Rosenau, 1992; Searle, 1983; Turiel, 1989, in press; Williams, 1986.)

More to our point. In deconstructivist theory, what recourse does an oppressed person have? Consider, for a moment, an imaginary dialogue between a deconstructivist job-seeker (D) and a racist employer (R):

D: You are wrong to discriminate again me.

R: I think I'm right.

D: Let me tell you the reasons why I think you are wrong. Be mindful, though, that these are reasons I find compelling. You may not find them so because you are a different person, in a different place. But I do hope you find my reasons compelling.

R: In my view, reason doesn't matter.

D: I think that is too bad, and mistaken, but I respect your judgment that reason doesn't matter. But please hear me out. When you discriminate against me, you are hurting me materially and psychologically. And you know I have a family to feed. Moreover, the discrimination is unjust, for as you yourself have said, I am more qualified for the job than anyone else.

R: That's interesting. But you are asking me to take seriously your welfare and rights. But such constructs don't have any universal or objective standing. The way I was brought up, people of your color don't merit moral consideration. Have a nice day.

Deconstructivist theory seems to have little to say in such situations. True, after an appeal to the oppressor fails, the person can appeal to various constituencies or governmental bodies. In the United States, for example, one could appeal to Congress to pass laws against the said oppression, or appeal to

the Supreme Court to rule on constitutional grounds. But such appeals from deconstructivists are somewhat deceptive in that they are not based on the belief that, for example, the Supreme Court tries to hold true to -- or in concept could hold true to -- fundamental principles of justice. Rather, from a deconstructivist perspective, the appeal is made because in our culture that is the way to gain power. Judicial decisions are secondary, power primary. Our point then is this. Deconstructivists ultimately have only one recourse when injustices occur: to gain power to stop the injustice. This is the reason that deconstructivists so often emphasize power in their analyses.

Now, the added twist to this scenario is that once deconstructivists gain power, it is very easy for them to fell prey to perpetrating the same injustices that they rebelled against. After all, other groups are the "other," are different, and thus potentially not deserving of the same moral considerations as those of one's own group. This is especially so given that the very construct of morality is but a product of person, place, and time. Thus, for example, if computer education has emphasized boys at the expense of girls, and there are reasons to believe it has (cf. Huff & Cooper, 1987), deconstructivist feminists can come in and argue that computer education should emphasize girls, and so what if it is at the expense of boys. The "so what" occurs because there is fundamentally division between the genders, rather than a common humanity that is an umbrella over differences. Thus it is our contention, and worry, that deconstruction as a theory is open to become totalitarian, disregarding of human rights and dignity.

Such worries are not unfounded. For example, a feminist scholar at the University of Illinois, Chicago was recently removed from teaching courses in sociology and women's studies (Magner, 1992). What appears to have happened is that a male student had disagreed with many of the teacher's feminist positions. As a result, according to university investigators, it followed that the teacher did not accord this male student the same classroom talking rights as female students, and pressured him to drop the course. And there is our beef. Because female students have been and often still are unjustly silenced in the classroom, this teacher saw no wrong with silencing a male student. Such exclusionary orientations, in the name of authenticating the female students' voice, abound in feminist scholarship

(see, e.g., Daly, 1980). Albeit, it has been noted by some women scholars (Patai, 1992; Spelman, 1988) that such exclusionary practices usually end up excluding not only men, but many women, particularly those who are not white, middle class, and privileged.

Educational computing that takes seriously deconstructivist theory can fall prey to similar abuses of power. Consider again the computer education that Harvey implemented. Traditionally, adults control what students learn, how they learn it, and when. Harvey sought to change that. Students had access to the class computers every hour of the day, every day of the week. Students chose what programming languages to learn, and indeed if they wanted to learn one at all. Students established computer policy. But, from a deconstructivist perspective, what would or could bound these choices? What, for example, was to prevent students from violating property rights by copying commercial software for home use? or from violating students' privacy rights by accessing personal files? or from disenfranchising other students who might want to join the group? or, indeed, from destroying computer equipment if they so chose? The deconstructivist answer, again, would be something like: students are checked by the power other people have over them: by, for example, the power exerted by teachers, principals, and police. Thus it should be clear that deconstructivist theory is at odds with constructivist theory. Constructivists seek to move students from obedience to autonomy, to a freedom that is bounded ethically by respect for person, care, human dignity, human rights. Deconstructivists, while they speak as if they promote autonomy -- since anything goes -- in fact promote obedience, since one becomes checked not by a regulated self, but by outside regulators.

The lack of establishing principled positions limits, as well, ethical computing practice on a societal level. Consider the case of a recent marketing venture proposed by Equifax and Lotus Development Corporation. Together, they proposed to develop a commercial "database containing profiles on 120 million people in the United States, obtained from census records, public files, and mailing lists. Included were the age, gender, marital status, household income, and buying habits of each of these Americans" (Rothfeder, 1992, p. 100). Equifax planned to sell the database to small business on a CD-ROM. This venture was picked up by the national press, and Equifax received more

than thirty thousand telephone calls in protest. Over and over again, citizens objected to what they viewed would be an invasion of their privacy. In response to this protest, Equifax and Lotus dropped their proposal. On a commercial basis, it made no sense for them to garner the ill-will of the consumer public. But while consumer power won one battle, in the process no principled position on computer privacy was articulated, principled in the sense that judgments could transcend that particular proposed venture, in that particular context, at that particular time. Thus, as deconstructivists might well argue, it follows that each new similar venture must be fought against again, with equal outrage and in equal numbers. Recent history suggests that will not likely happen. As documented by Rothfeder (1992), much worse invasions of privacy by means of computerization occur daily by such large corporations as TRW, but without censure.

Many postmodern theorists have been troubled by at least some of the above concerns about deconstructivism, in theory and practice. In response, they have attempted to put forth modified positions which Rosenau (1992) and others have labeled as "affirmative" postmodern theories. Affirmative theories (e.g., Giroux, 1990; Hammer & MacLaren, 1991; Hassan, 1985; Murphy, 1987; Richardson, 1988; Weiler & Mitchell, 1992; Wyschogrod, 1990) still argue for the plurality of value systems but do not maintain that such plurality necessarily leads to the relativism that is so troubling in deconstructivism. As noted by Rosenau (1992) "[a]ffirmative post-modernists frequently employ terms such as oppression, exploitation, domination, liberation, freedom, insubordination, and resistance -- all of which imply judgment or at least a normative frame of reference in which some definitive preferences are expressed" (p. 136). Moreover, in contrast to nihilism that often pervades deconstructivist political theory, affirmatives often favor forms of democracy that empower individuals and especially underrepresented groups. As noted by Apple (1992),"[a]ll too often, 'legitimate' knowledge does not include the historical experiences and cultural expressions of labor, women, people of color and others who have been less powerful" (p. 7). Thus Apple favors the "growth of more democratically run schools, of practices and policies that give community groups and teachers considerably more authority in text selection and curriculum determination, in teaching strategy...." (p. 5). At the same time, affirmatives

usually embrace a deconstructivist-like epistemology wherein it is maintained that all knowledge is socially constructed, and that "what counts as legitimate knowledge is the result of complex power relations, struggles, and compromises among identifiable class, race, gender, and religious groups" (Apple, 1992, p. 4).

We applaud the affirmative's focus on democracy; and such representation and empowerment are central to the educational computing practices we have described. But can affirmatives maintain their non-relativistic views in light of their deconstructivist-like epistemology? Affirmatives think they can, though are often circumspect in articulating exactly how. As we understand their position, however, the skeleton of their response looks something like this. They maintain that knowledge is not objective. At the same time they maintain that neither is knowledge subjective because knowledge is grounded in socially constituted relations, bounded by community. As Murphy (1988) says: "[A]narchy is not necessarily the outcome of postmodernism, because public discourse can culminate in the promulgation of social rules" (pp. 181-182). Thus like deconstructivists they deconstruct the objective/subjective polarity; but as affirmatives they maintain that not anything goes. QED: postmodernism without relativism.

The problem here lies in believing that majority opinion or community beliefs solves the problem of relativism, when in fact it does little more than raise the problem from an individual to group level. A case in point: Imagine people inside a house without windows listening to a slight pitter patter on the roof. After much discussion and factional power struggles, they all agree that it is raining outside. Then a person from outside their community, and literally from outside their house, walks from the beautiful sunny day into their house, and asserts that it is sunny outside. Now, presumably there are real occurrences of "raining" and "not raining". Presumably in this case the people inside the house are simply mistaken in believing it is raining outside. Thus one can agree that the people inside the house have socially shared knowledge, and that that knowledge goes beyond mere subjectivity of each member. But to say that is not the same as to say that shared knowledge ipso-facto validates that knowledge. And the same holds true for ethical knowledge. A community can agree to discriminate

against (or torture or slaughter) members from outside their community, but such agreements do not establish ethical validity.

Affirmatives might respond by saying that for a community to have valid ethical knowledge, not only must members within its community agree to it (thus protecting their own members from oppression), but similarly any time norms are applied to outside members, then those outside members must agree as well. Perhaps affirmatives would thereby establish the following principle: Membership in a democratic community is accorded to those who are affected by its norms, and norms protect minority from majority oppression. A move like this then begins to bound the ethical by establishing universal criteria, and by a conception of what constitutes oppression in a principled and privileged, if not objective, sense. In so doing, affirmatives begin to cast aside their postmodern epistemology, as they must if they are to escape serious internal contradictions or ethical relativism.

And that is precisely what we suggest. Cast aside postmodern epistemology, and then move ahead with the best of what remains of postmodern thinking and practice, which is quite a lot, but which is no longer postmodern but a reinvigorated or perhaps transformed modern view. Yes, take seriously culture and context. Promote democracy. Argue for increasing inclusiveness of previously disenfranchised groups. Bring people in. Build community. Recognize that prevailing scientific theories are (as scientists themselves agree) not fully adequate. Do better science. Change our technological practices -- in the work place and schools -- to promote more humane and responsible social interactions. Recognize the limitations of our Western outlook, particularly as it divorces us from each other and the natural landscape with which we have a deep evolutionary connection, which we ignore at our peril. But do all of this, and more, without rejecting reason, science, logic, objectivity, truth, morality. It is within this reinvigorated "modern" view that much of the interesting work is being done across the disciplines, and it is within this context that we seek to place the constructivist educational computing practices and communities that we have described.

References

A Summary of the ACM/IEEE-CS Joint Curriculum Task force Report: Computing Curricula 1991. (1991). Communications of the ACM, 34(6), 69-84.

Appadurai, A. (Ed.). (1988). The social life of things. New York: Cambridge University Press.

Appel, C. (1985, May 16). L-S "superusers" run the computer room. The Sudbury Town Crier, pp. 3, 43.

Apple, M. W. (1992). The text and cultural politics. Educational Researcher, 21(7), 4-11, 19.

Arnstine, D. (1973). The use of coercion in changing the schools. Educational Theory, 23, 277-287.

Arsenio, W. F. (1988). Children's conceptions of the situational affective consequences of sociomoral events. Child Development, 59, 1611-1622.

Baldwin, J. M. (1973). Social and ethical interpretations in mental development. New York: Arno. (Original work published 1897)

Berlins, M., & Hodges, L. (1981, January 15). Nationality Bill sets out three new citizenship categories. The London Times, pp. 1, 15.

Bodker, S. (1991). Through the interface: A human activity approach to user interface design. Hillsdale, NJ: Lawrence Erlbaum Associates.

Budge, B. (1983). Pinball construction set. [Computer program]. San Mateo, CA: Electronic Arts.

Bromley, H. (1990). Culture, power and educational computing. Unpublished manuscript, University of Wisconsin, Madison.

Bynum, T. W., Maner, W., & Fodor, J. L. (1992). Teaching computer ethics. New Haven, CT: Research Center on Computing and Society.

Chandrasekaran, B. (1992). What's in a logo? IEEE Expert, 7(5), 2.

Crews, F. (1986, May 29). In the big house of theory. The New York Review of Books, pp. 36-42.

Crews, F. (1989, July 20). The parting of the Twains. The New York Review of Books, pp. 39-44.

Culler, J. (1982). On deconstruction: Theory and criticism after structuralism. Ithaca, NY: Cornell University Press.

Cummins, J. (1988). From the inner city to the global village: The microcomputer as a catalyst for collaborative learning and cultural interchange. Language, Culture, and Curriculum, 1(1), 1-13.

Daly, M. (1980). Gyn/ecology: The metaethics of radical feminism. Boston: Beacon Press.

Damon, W. (1977). The social world of the child. San Francisco: Jossey-Bass.

Dawkins, R. (1976). The selfish gene. New York: Oxford University Press.

Denning, P. J. (1992). Educating a new engineer. Communications of the ACM, 35(12), 83-97.

Derrida, J. (1978). Writing and difference. London: Routledge and Kegan Paul.

DeVries, R. (1988, October). Constructivist education. Paper presented to the Association for Constructivist Teaching, West Point Academy.

DeVries, R., & Kohlberg, L. (1990). Constructivist early education: Overview and comparison with other programs. Washington, D. C.: National Association for the Education of Young Children.

Dewey, J. (1966). Democracy and education. New York: Macmillan. (Original work published 1916)

Donmoyer, R. (1992). The Annual Meeting: Attention must be paid. Educational Research, 21(5), 38.

Dormer, J. (1981, June 18). Computer center is a novel program. Concord Journal, p. 19.

Dunlop, C., & Kling, R. (Eds.). (1991). Computerization and controversy: Value conflicts and social choices. San Diego, CA: Academic Press.

Ermann, M. D., & Williams, M. B., & Gutierrez, C. (Eds.). (1990). Computers, ethics, and society. New York: Oxford University Press.

Friedman, B. (1986, October). If I only had one more computer... Facing the sticky issues of resource allocation. Classroom Computer Learning, pp. 44-45.

Friedman, B. (1991). Social and moral development through computer use: A constructivist approach. Journal of Research on Computing in Education, 23, 560-567.

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., & Nissenbaum, H. (1992, August). Bias in computer system design: Pre-existing, technical, and contextual. Paper presented at the 4S/EASST Joint Conference, Goteborg, Sweden. Abstract published in Technoscience, 5(2), 25.

Friedman, B., & Winograd, T. (Eds.) (1990). Computing and social responsibility: A collection of course syllabi. Palo Alto, CA: Computer Professionals for Social Responsibility.

Giroux, H. A. (1990). Curriculum discourse as postmodernist critical practice. Victoria, British Columbia: Deakin University Press.

Giroux, H. A., & Simon, R. I. (1992). Schooling, popular culture, and a pedagogy of possibility. In K. Weiler & C. Mitchell (Eds.), What schools can do (pp. 217-236). Albany, NY: State University of New York Press.

Gotterbarn, D. (1992). The use and abuse of computer ethics. Journal of Systems Software, 17, 75-80.

Greenbaum, J., & Kyng, M. (Eds.). (1991). Design at work: Cooperative design of computer systems. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hammer, R., & MacLaren, P. (1991). Rethinking the dialectic: A social semiotic perspective for educators. Educational Theory, 41(1), 23-46.

Harvey, B. (1980, May). Using computers for educational freedom. Paper presented at the Second annual Computer Conference, Lesley College, Boston.

Harvey, B. (1983). The Lincoln-Sudbury Regional High School computer department: History, goals, strategies, and problems. Unpublished manuscript.

Hassan, I. (1985). The culture of postmodernism. Theory, culture, and society, 2, 119-131.

Hogan, R. (1975). Theoretical egocentricism and the problem of compliance. American Psychologist, 30, 533-539.

Horowitz, P. (1984, May, June). Telecommunications: An electronic, interstate school newspaper. Electronic Learning, p. 48.

Hoy, D. C. (1985). Interpreting the law: Hermeneutical and poststructuralist perspectives. Southern California Law Review, 58, 135-176.

Huff, C., & Cooper, J. (1987). Sex bias in educational software: The effect of designers' stereotypes on the software they design. Journal of Applied Social Psychology, 17, 519-532.

Johnson, D. G. (1985). Computer ethics. Englewood Cliffs, NJ: Prentice-Hall.

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.

Killen, M. (1990). Children's evaluations of morality in the context of peer, teacher-child, and familial relations. Journal of Genetic Psychology, 151, 395-410.

Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D. A. Goslin (Ed.), Handbook of socialization theory and research (pp. 347-480). New York: Rand McNally.

Kohlberg, L. (1980). High school democracy and educating for a just society. In R. L. Mosher (Ed.), Moral education: A first generation of research (pp. 20-57). New York: Praeger.

Kohlberg, L. (1984). Essays on moral development: The psychology of moral development. Vol. 2. San Francisco: Harper & Row.

Kohlberg, L. (1985). The just community approach to moral education in theory and practice. In M. W. Berkowitz & F. Oser (Eds.), Moral education: Theory and application (pp. 27-87). Hillsdale, NJ: Lawrence Erlbaum.

Laupa, M. (1991). Children's reasoning about three authority attributes: Adult status, knowledge, and social position. Developmental Psychology, 27, 321-329.

Lewis, P. H. (1992, November 8). Nurturing democracy through used computers. The New York Times, p. F12.

Magner, D. K. (1992, October 7). Feminist scholar at U. of Ill. -- Chicago rebuked after discrimination charge. The Chronicle of Higher Education, pp. A27-28.

Miller, K. (1988). Integrating computer ethics into the computer science curriculum. Computer Science Education, 1, 37-52.

Morss, J. R. (1992). Making waves: Deconstruction and developmental psychology. Theory and Psychology, 2, 445-465.

Murphy, J. W. (1987). Deconstruction, discourse, and liberation. Social Science Information, 26, 417-433.

Murphy, J. W. (1988). Computerization, postmodern epistemology, and reading in the postmodern era. Educational Theory, 38, 175-182.

Namioka, A., & Schuler, D. (Eds.). (1990). Proceedings of the 1990 Participatory Design Conference. Palo Alto, CA: Computer Professionals for Social Responsibility.

Noble, D. D. (1991). The classroom arsenal: Military research, information technology and public education. London: Falmer Press.

Norris, C. (1982). Deconstruction: Theory and practice. New York: Methuen.

Nucci, L. P. (1981). The development of personal concepts: A domain distinct from moral and societal concepts. Child Development, 52, 114-121.

Parker, D. B., Swope, S., & Baker, B. N. (1990). Ethical conflicts in information and computer science, technology, and business. Wellesley, MA: QED Informations Sciences.

Patai, D. (1992, February 5). The struggle for feminist purity threatens the goals of feminism. The Chronicle of Higher Education, pp. B1, B2.

Perrolle, J. A. (1987). Computers and social change: Information, property, and power. Belmont, CA: Wadsworth.

Piaget, J. (1969). The moral judgment of the child. Glencoe, IL: Free Press. (Original work published 1932)

Power, C., Higgins, A., & Kohlberg, L. (1989). The habit of the common life: Building character through democratic community schools. In L. Nucci (Ed.), Moral development and character education: A dialogue (pp. 125-143). Berkeley: McCutchan.

Richardson, L. (1988). The collective story: Postmodernism and the writing of sociology. Sociological Focus, 21, 199-207.

Rothfeder, J. (1992). Privacy for sale. New York: Simon & Schuster.

Scheffler, I. (1991). In praise of cognitive emotions and other essays in the philosophy of education. New York: Routledge.

Scholes, R. (1989). Protocols of reading. New Haven, CT: Yale University Press.

Searle, J. R. (1981). Minds, brains, and programs. In D. R. Hofstadter & D. C. Dennett (Eds.), The mind's I (pp. 353-382). New York: Basic Books.

Searle, J. R. (1983, October 27). The word turned upside down. New York Review of Books, pp. 74-79.

Searle, J. R. (1984). Minds, brains and science. Cambridge: Harvard University Press.

Searle, J. R. (1990). Is the brain's mind a computer program? Scientific American, 262, 26- 31.

Searle, J. R. (1992). The rediscovery of the mind. Cambridge: MIT Press.

Selman, R. L. (1980). The growth of interpersonal understanding. New York: Academic Press.

Shneiderman, B. (1987). Designing the user interface: Strategies for effective human-computer interaction. Reading, MA: Addison-Wesley.

Shneiderman, B. (1989). Designing the user interface. In T. Forester (Ed.), Computers in the human context: Information technology, productivity, and people. Cambridge, MA: The MIT Press.

Shweder, R. A. (1986). Divergent rationalities. In D. W. Fiske & R. A. Shweder (Eds.), Metatheory in social science: pluralisms and subjectivities (pp. 163-196). Chicago: University of Chicago Press.

Smetana, J. (1982). Concepts of self and morality: Women's reasoning about abortion. New York: Praeger.

Snyder, T. (1985a). Our town meeting. [Computer program]. Cambridge, MA: Tom Snyder Productions.

Snyder, T. (1985b). The other side. [Computer program]. Cambridge, MA: Tom Snyder Productions.

Spelman, E. V. (1988). Inessential woman: Problems of exclusion in feminist thought. Boston: Beacon Press.

Thorkildsen, T. A. (1989). Pluralism in children's reasoning about social justice. Child Development, 60, 965-972.

Tisak, M. S (1986). Children's conceptions of parental authority. Child Development, 57, 166-176.

Turkle, S. (1984). The second self: Computers and the human spirit. New York: Simon & Schuster.

Turiel, E. (1983). The development of social knowledge. Cambridge: Cambridge University Press.

Turiel, E. (1989). Multifaceted social reasoning and educating for character, culture, and development. In L. Nucci (Ed.), Moral development and character education: A dialogue (pp. 161-182). Berkeley: McCutchan.

Turiel, E. (in press). Morality, authoritarianism, and personal agency in cultural contexts. To appear in R. J. Sternberg & P. Ruzgis (Eds.), Intelligence and Personality, Cambridge: Cambridge University Press.

Vygotsky, L. S. (1978). Mind in society. Cambridge: Harvard University Press.

Wainryb, C. (1991). Understanding differences in moral judgments: The role of informational assumptions. Child Development 62, 840-851.

Weiler, K., & Mitchell, C. (Eds.). (1992). What schools can do. Albany, NY: State University of New York Press.

Weizenbaum, J. (1976). Computer power and human reason. New York: Freeman and Co.

Wigginton, E. (1986). Sometimes a shining moment: The Foxfire experience. Garden City, NY: Anchor Books.

Williams, B. (1985). Ethics and the limits of philosophy. Cambridge: Harvard University Press.

Winograd, T. (1992, Summer). Computers, ethics, and social responsibility. The CPSR Newsletter, 10(3), pp. 17-27.

Winograd, T., & Flores, F. (1986). Understanding computers and cognition: A new foundation for design. Reading, MA: Addison-Wesley.

Wyschogrod, E. (1990). Saints and postmodernism: Revisioning moral philosophy. Chicago: University of Chicago Press.

Youniss, J. (1980). Parents and peers in social development. Chicago: University of Chicago Press.