

Part 1  
Introduction to Computer Ethics



# Computers, Ethics, and Social Responsibility

—Terry Winograd, 1991

It is a great honor to have been invited to give the computer science keynote address at this landmark conference: the first in what will surely be a long series of important events and activities produced by the Research Center on Computing and Society. It is a tremendous opportunity to meet with a community of people who have shared their concerns about computing and values over the years, and to join together in developing new directions and new possibilities for future collaboration.

I am especially pleased to be able to share in the honoring of Joseph Weizenbaum for his pioneering work and continuing efforts to bring considerations of values into the work of computing. Long before most of us here were even aware that such a topic was possible, he was bringing it into the heart of the technological world, and getting people to listen. Several years ago at the annual meeting of Computer Professionals for Social Responsibility I was honored to be able to present him with the Norbert Wiener Award for Social and Professional Responsibility, and it is a pleasure to be with him as he again receives well-deserved recognition.

## **What can a computer scientist say?**

Let me begin by admitting that it wasn't exactly clear to me just what the content should be for a "computer science keynote." The path seems clear for my colleagues who will present keynotes on philosophy and sociology, since those fields include ethics and values in their core subject matter. As a philosopher, one can develop theories of ethics. As a sociologist one can study the ways that people ream, change, and exhibit values. But as a computer scientist I don't study ethics and values; I study computers and computing. As we are all well aware, "ethics" and "values" aren't the kind of things to be addressed with the theories and techniques of computer science. For the computer scientist, they are not an object of study, but a domain in which we interpret and assess our actions as professionals. My role, then, is not as an observer, but as a subject. As a computer professional and a teacher of future computer professionals, my concern is very personal. The questions aren't academic, but practical: not "What is done?" but "What should we do?"

So I will enjoy the liberty in this paper of not having to precisely define the difficult concepts we speak about or having to argue the logical merits of a particular theory. Instead I will talk about how issues of computing and values show up in the work of our profession. You might think of it as being ichthyology from the point of view of the fish.

Also, in talking about these issues I will not try to draw a careful line between terms such as "ethics," "morals," "values," and "social responsibility." These distinctions can be important for some purposes, but I will proceed for the moment to interchange them

---

This article was delivered as the keynote address at the first National Conference on Computing and Values, held in New Haven, Connecticut, in August 1991. It appears in Terrell Ward Bynum, Walter Maner, and John L. Fodor (eds.), *Computing and Human Values: Proceedings of the 1991 Conference*, New Haven: Research Center on Computing and Society, 1992.

freely with more of a concern for the ring of the sentence than for the precise differentiation of the concepts.

### **The personal connection**

When I speak of my own work, I include more than the narrow pursuit of research and development in computer science. For almost ten years now I have been a participant in the work of Computer Professionals for Social Responsibility, an organization that has brought together people from around the country (in fact, around the world) to share understandings and to act collectively in many of the areas that are being discussed at this conference. That activity is not a diversion but a critical part of the work of a computer professional. One of the things I want to highlight is the way in which organizations like CPSR and NCCV play a central role in ethical conduct for computer professionals.

In addition, over the past three years, Helen Nissenbaum (now at Princeton) and I have developed and taught a course on “Computers, Ethics and Social Responsibility” for undergraduate computer science majors at Stanford. As all of us in academia know well, there is no better way to expand your own understanding than to throw yourself into a room full of bright undergraduates who want to master a difficult topic and expect you to help. Much of my understanding has grown from the generative interaction that comes in teaching, and that too is a central part of my work as a computer scientist. It has forced me into some hard and productive thinking about the questions being raised at a conference on Computing and Values.

### **What I will say**

In this talk I will present and contrast some common views of how ethics and values are related to computing and see what these views imply for the activities we can undertake to promote ethical behavior and social responsibility. My emphasis will be on the fundamentally social nature of ethical concerns: with looking beyond the role of the individual to the larger context of discourse and action that generates the world in which individuals make choices and to act. Rather than focusing on the isolated individual faced with an ethical dilemma, I want to direct our gaze to the larger swirl of human discourse, which is the source of the interpretations, values, and possibilities that make ethical choice meaningful.

The announcement for the NCCV conference declared a vision:

To integrate computer technology and human values in such a way that the technology advances and protects those values rather than doing damage to them.

This will require acts of individual moral courage, and it will be based on a lot more. We need to create an environment in which the consideration of human values in our technological work is not a brave act, but a professional norm. We need to produce a background of understanding in which it is simply taken for granted by all computer professionals that value considerations are foremost. We need to forge everyday practices and ways of teaching that reinforce that understanding.

In that spirit, I will argue that the kind of inquiry and discussion that motivate the conference, and that have been at the heart of CPSR's ten years of work, are a primary form of ethical behavior.

### **Being a “good” computer professional**

First, let us go back to the basic question of what values, ethics, and social responsibility have to do with computing. I said above that ethics and values constitute “a domain in which we, and others, view and assess our actions as computer professionals.” What do we mean by “assess our actions”?

### **Assessments and competence**

In every area of purposeful endeavor, there exist communities of assessment within which it is possible to meaningfully describe, compare, and evaluate action. As a computer scientist I am part of a scientific community with standards of practice, and practices of assessment. There may be no straightforward quantitative measure of whether I am a “good computer scientist,” but there are ways in which all of us measure the achievements of others and of ourselves. In the academic world these include publication records, peer review, awards, election to various professional and honorary societies, and the like. They also include less tangible but still consensual domains of reputation, status, and in the longer run your “place in the history of the field.” I identify myself as belonging to a scientific community and I participate in its consensual processes for assessment. For example, my kids may think I'm a fantastic computer scientist because I could get Tetris running on our Macintosh, but I don't value this in the same dimension as the judgment of colleagues whom I consider part of the community.

When we look to the computer science community in general we see a notable lack of concern for many of the values addressed at the NCCV conference. There is an implicit definition of “good computer scientist” that dismisses people like Joe Weizenbaum as bothersome troublemakers, and accepts without qualms people who are oblivious of the value consequences of their actions. One of my colleagues, in a note rejecting my questioning the sources of research funding said he feared I would describe him, as “Having the moral fiber of a styrofoam cup.” In judging whether people are “good computer scientists” the professional norms are strongly attuned to particular concerns and kinds of action and not to others.

But in our common sense assignment of “good” and “bad,” we take a broader view. A “good baseball player” isn't just one who hits home runs, but one who contributes to the efforts of the team as a whole. He may be a great source of spirit and enthusiasm, a kind mentor to younger players, and a contributor in many other ways to the success of the team.

But in talking of the success of the team we're back to assessments again: What constitutes “success” for the computer science team? Again there is a gap between what we see in many of our work settings and what we are striving for here. Our measure of success needs to be the one quoted above:

To integrate computer technology and human values in such a way that the technology advances and protects those values rather than doing damage to them.

With this as our measure, we are ready to look for “good computer science.”

### **What is the domain of ethical action?**

If our goal is to “advance and protect human values,” then what kind of actions will further it? Ethics isn’t an immediately obvious domain of actions. If you ask what competence is being developed in a cooking class, it is evidently “cooking.” We can identify people at specific times as being engaged in cooking. But we are never “ethicking” in that simple sense. We may be performing an engineering job, making a living, doing scientific research, (or, for that matter, cooking) and find ourselves in situations where our actions raise some kind of ethical question. How do we identify those situations?

In some sense this is an “academic question.” We all grow up with a tacit understanding that there can be things we do that are “right” and others that are “wrong,” and that as autonomous individuals we have responsibility for choosing between them. We all have a sense that we should do what is right, even though that isn’t always what we end up doing. We also grow up in today’s global pluralistic society with an awareness that although everyone has a sense that there is a difference between right and wrong, there is no agreement on just what actions should fall under which category. There are tremendous disputes between different cultural, religious and political groups, which have led to arguments, wars, and disagreements at all levels of society throughout history. But, on the other hand, there seem to be commonalities. Nearly everyone would agree that it is wrong to simply walk up to someone on a whim and inflict pain, and that it is right to help others in need.

People have debated for thousands of years what moral and ethical standards should be. Is there a universal ethics that applies to all people in all ages? Or is ethics a purely relative matter in that what is considered a fundamental moral principle by one people at one time may be equally validly rejected in another culture?

Now if I were a philosopher I would feel compelled to try to make sense of all this: to come up with a coherent moral philosophy that could serve as a basis for understanding what we see in the historical discourse about morality and for making decisions about our own actions. But, as I said at the beginning, I am taking the easy way out. I will leave the philosophical analysis for our colleagues who are much more skilled and knowledgeable, and will appeal to a rather commonsensical basis of agreement. I think we can all accept that in at least some interesting range of cases it makes sense to talk about doing “good” and “bad,” and furthermore that we all, to some degree, accept the value of “doing good.”

Further, we seem to have some general understanding of what kind of “doing good” constitutes an ethical or moral act. If you take a course on programming languages, you may learn that it is “good” to have a grammar that can be parsed by an LR(1) parser, and “bad” to have ambiguous constructs. But this domain of assessment, which is proper to

the computer professional, doesn't seem to have much to do with the kind of human values we are discussing here. Something is missing in the equation.

Taking the naive view again, it seems obvious that the missing element has to do with a regard for the interests of others. There is a popular refrain about acts that are "illegal, immoral or fattening." The distinction between law and morality is an important one we will not go into here. But it is clear that there is something different about "fattening." It may be stupid or unhealthy or unwise to fill myself up on chocolate bars and potato chips, but few of us would consider it unethical or immoral. In general we take moral questions to involve a potential for conflict of interests. In the case of religious morality, the "other" may be a deity. For secular ethics, it is among people (and perhaps other life forms or embodiments of intelligence).

To be fair, this is a very complicated issue, but again in the spirit of simplification, we can accept that for most of the issues that attract our interest, our actions have consequences of value to others. Consider, for example the four clusters of values that are the focus of this conference: Privacy, security, ownership and fair access. In each case it is easy to identify the different parties and potential conflicts of interest, and we do so as a matter of course in teaching about these topics.

Finally in completing this background discussion, I want to make a key point about the role of intentions. Putting it generally, the domain in which an action is assessed is not necessarily the same as the domain in which the actor interprets it. If I ask what some person is doing, you may say that she is busy "establishing an image of authority" even though she is not consciously acting in that realm. Someone can be assessed as a "great teacher" when what she sees herself doing is having an argument or commenting on a talk she is listening to.

Similarly, acts can be observed in the domain of ethics with respect to standard practices, independent of whether the person characterizes them that way. The fact that someone didn't think about the consequences of an act doesn't remove them from being subjected to moral judgment about it. In fact, we can take ethical obliviousness as a key sign of "bad" behavior.

But thinking a moment further, it also doesn't seem appropriate to assess an act as wrong if there was no background of understanding in which it could show up as such. If we now see harmful consequences of the farming practices of primitive tribes, it doesn't mean that they were acting unethically to do them.

Again, we must look to the social context. A person does not exist in a vacuum, but as part of one or more social collectivities, with their shared interpretations of actions, values and assessments. A person cannot be held responsible for considerations that lie completely outside the range of vision opened up within these backgrounds. There may often be cases where an individual rejects the current consensus of society and appeals to a larger context of human meanings and values. But in doing this he or she is responsible for participating in the social discourse and not simply ignoring the concerns of his or her co-denizens. This means that a key component of moral action is the development of

understanding within a social background, which is what provides the relevant field of choice for individuals.

### Three cartoons for how to “ethic” well

So far we have been taking the view of the observer: one who interprets and assesses the acts that have been done by someone (who may be him/herself) in the past. Let us shift to the view of the doer: the person who is engaged in action that can have consequences in areas of values and morals. Faced with a particular range of possibilities, how does one “ethic” well?


One of the things that becomes clear in teaching this material to students is that people come to this question with a variety of tacit pre-understandings of what we are trying to do. They draw on images that are deeply embedded in our culture, and I want to present some of them in the form of cartoons that exaggerate, but also point out some of the key features. For each of the images, we need to ask several questions:

- What are the assumptions that lie behind its perspective?
- What problems does it raise?
- How from that perspective do we develop people’s competence to act?

### The angel/devil debate

The cartoon shown in Figure 1 is the familiar angel/devil debate you’ve all seen on Saturday morning TV. A character is faced with an ethical choice and is obviously having trouble deciding what to do. Sitting on one of his shoulders is a little pointy-tailed demon whispering into his ear “C’mon, take it, he’ll never know you did.” and on the other shoulder a haloed cherub, sweetly whispering “You know you shouldn’t steal.” In the end one of them is brushed away with a flick of the fingers and the other dances

Figure 1. The angel/devil debate



gleefully in victory.

There are several assumptions implicit in this view of morality:

1. You know what is right or wrong in the particular case.
2. Some part of you wants to do the thing that is wrong.
3. You need to exert moral strength to overcome this impulse and do what is right.

If, in fact, this view were the whole story, then the teaching of computer ethics would be a very different matter from what we see here today. Education directed to this kind of ethical competence might include sermons, examples (stories of sinners sizzling in Hell), and practices such as self-denial. In fact much of the resistance to the teaching of computer ethics within computer science departments comes from the impression that this is what it will consist of, and a skepticism as to whether such moralizing has any positive effects.

The fallacy in the angel/devil view is obvious if we look at Bynum's characterization of the goals for teaching computer ethics. (Bynum, 1992)

1. To sensitize students to computer ethics issues
2. To provide tools and methods for analyzing cases
3. To provide practice in applying the tools and methods to actual or realistic cases
4. To develop in the student "good judgment" and "helpful intuitions" for spur-of-the-moment decision-making as computer professionals and computer users

Faced with an ethically problematic situation we must first recognize it as such. It doesn't appear with angels and devils drawn in the corners, but must be seen through a background of interpretation in which ethical issues have been distinguished and made a part of our everyday discourse.

Of course I know I shouldn't kill other people. I walk around every day doing the right thing hundreds of times by not killing someone. But it isn't an ethical issue for me; it's part of the taken-for-granted background. But when I need to decide whether to build a worker-monitoring system or take research funding from a military agency, I am in the situation of debating what is in fact right and what the underlying issues are. It isn't a simple matter of steeling myself to be righteous. We must be able to recognize our specific situation as it fits into the context of issues and other cases that has shown up historically. These activities are skills to be learned and development, not character traits like "moral will."

### **The morality computer**

Having shifted the question from moral character to understanding and analysis, we find ourselves closer to the cartoon shown in Figure 2, the "morality computer."

Faced with a decision, you don't know which is the right action to take. Should you steal a horse to chase the bandit? Should you cut off life support to relieve suffering? Should you work on nuclear physics knowing the results may produce economic prosperity and also may lead to a weapon of mass destruction?

**Figure 2. The morality computer**



So you type the information into your morality computer, which has been programmed with the correct moral rules. It sifts the facts, weighs alternatives, makes judgments and pops out after a few microseconds with “Here’s the right thing to do. . . .”

What are the assumptions lying behind this picture?

1. If you can figure out the right thing to do, then you will do it.
2. There is a basic set of moral rules from which to deduce the rightness of action in any given case.
3. You may know the rules but not know how to apply them in this case, and more data, knowledge, or computation is needed.

In the “morality computer” view, the problem is determining what is the ethical action. Competence consists in knowing the moral axioms and having deductive skill in applying them to cases where the question “What is the right thing to do?” comes up. This approach appeals strongly to people with a background in science and engineering. When they encounter ambivalence and ambiguity they see it as a symptom that the problem has not been well formulated, or that we do not have enough knowledge. The fix is the kind of fix that works in technical domains: get the rules right, find the correct methods of applying them, and the right answers will come out. It offers the possibility of a “technical fix to the ethics problem.”

Education in this perspective, as in other forms of education in science and engineering, is a matter of giving students the right principles and giving them practice in applying them to cases. At times we hear frustration from some of the students who take

our course because we are not providing them with this kind of structure: we aren't able to give them the precise rules and methods, so they can learn to plug in the data and come up with answers.

But, of course, it doesn't take sophisticated philosophical reading to recognize that despite millennia of debate, humankind doesn't seem close to reaching agreement as to the general grounding for moral reasoning at all, much less the specific rules. Within any moderately diverse group of people you will find a wide range of beliefs: Some will believe that morality is grounded in some form of divine intention, others that it is a feature of human psychology, and others that it is some kind of "optimization" principle concerned with the welfare of the species. Some will base their moral reasoning on a structure of absolute principles—do's and don'ts—while others see it as some kind of calculus of costs and benefits.

Regardless of which approach you take, no matter how certain you are about the basic principles and rules, you find yourself puzzled by individual cases. One of the things we have become painfully aware of in nearly a half century of work on artificial intelligence is that there is a huge distance from abstract rules to real situations. Before rules can be applied, there must be interpretation as to how the terms in them actually fit the situation) and in doing so there is a wide-open field of human judgment and implicit understanding that has not yielded to logical analysis.

This all may be painfully obvious to those who have been working in the field of computer ethics, but let me give an example to clarify what I am pointing at.

Assume that you accept some form of the rule "Thou shalt not steal." Then in order to apply it, you need to know when an act is "stealing." We may define it as something like "taking property that belongs to someone else, without their consent." That's a good start, but what is "property?" There are clear examples such as someone's wallet or car, but what about their "ideas." Are those the kind of things that can be property at all? Now we are in a complicated realm of definitions, which has occupied philosophers (and lawyers) for centuries. The apparently simple notion of "property" has different interpretations in different cultures, legal systems, and traditions. Further, what do we mean by someone's "consent." What kind of consent is implied by opening a box that has a label on it saying "By opening this box you hereby agree to. . . .?"

In reading the literature on computer ethics<sup>1</sup> we encounter many more such examples and become painfully aware of how difficult it is to come up with consistent principles and standards for applying them to cases we encounter. The point should be obvious: we are not able to provide the kind of rules and methods that work in normal science and engineering, to come up with answers to problems. The morality computer is an idealized fantasy, and can mislead us if it shifts attention to a quest for the "right answer," away from the questioning activity that is required of each of us and involves us in a never ending dialog with others.

## A troupe of jugglers

And this leads us to the cartoon of Figure 3, a troupe of jugglers. It may seem that juggling is too frivolous an activity to be a relevant analogy, but let us look more closely at several key features of the situation.


### 1. Engaged activity

First we note that the jugglers are constantly engaged in action. The first two cartoons directed our attention to conscious moments of decision, and put the locus of ethical action in determining the outcome. This cartoon suggests that we are always “thrown” into acting and that the assessments of ethics apply to these actions, not just those where we stop to ponder. This is suggested by the fourth of Bynum’s descriptions of what we are doing in teaching:

To develop in the student ”good judgment” and “helpful intuitions” for spur-of-the-moment decision making as computer professionals and computer users.

This “spur of the moment decision making” is the basic condition of acting in the world. In fact, it often does not show up to the actor as decision making at all. We all remember the interviews with someone who has jumped into a river to save a drowning child, when the interviewer asks “What made you decide to do it?” and the hero or heroine says “Decide? I didn’t decide, I just jumped in?” In order to be skillful at “ethicking” we need to develop the kind of continuing judgment in action that a juggler exhibits, not just the kind of careful argument that a logician applies in constructing a proof.

Figure 3. A troupe of jugglers



## *2. Social context*

Second, the focus is not on the isolated actions of an individual, but on the coordinated actions of the troupe as a whole. What I do makes sense—is “right” or “wrong”—in the context of what others are doing. When I look to alternatives, I need to consider not just what else I might do, but what we all might do through some kind of agreement and coordination. As I suggested earlier, this is a key feature of ethical action. If we wait until someone is put into a true moral dilemma, we may get exciting drama, but we will not further the overall pursuit of values as much as if we develop standard practices that make it natural rather than heroic to do the “right thing.”

One of the most powerful ethical acts we can each do is to participate in creating a social context in which the future actions of ourselves and others are consistently in line with our values. This includes educating our colleagues and students, working to develop professional standards, exploring new technologies and identifying their consequences for values. Even though we may not face hard individual ethical decisions as part of that work, we are actively engaged in the juggling process.

## *3. Evolving understandings, practices, and standards*

Finally, we recognize in a juggling troupe the eternal need to learn and change. There is no ultimate “right” way to juggle. Clearly, any form of juggling will have to conform to the laws of gravity and physical motion. Less obviously, but plausibly, there may be perceptions of what is “good” that are grounded in the nature of the human animal and will be true across cultures and times. But within this, the community evolves practices and standards in which its members are trained, and by which their actions are assessed. Part of the skill we recognize in a community of jugglers is their ability over time to recognize new possibilities, develop skills in areas that hadn’t been previously explored, and be sensitive to the changing environment in which they perform.

In some cases, this may require focus on detail: evolving a new concept of just what constitutes property and ownership in a new domain such as software, interfaces and algorithms. At other times, we can make major leaps. When Gandhi proposed nonviolent civil disobedience as a way of furthering the human values he cared about, he created a new “clearing”—a new way of looking at possibilities and taking actions, which could have meaning and power in the world of the late 20th century.

Now it should be clear that I favor this third cartoon, and to be fair, we should apply the same questions as we did to the other two. First, what are the assumptions?

1. There is a social activity in which we are engaged, in which characterizations and assessments in an ethical domain can be made.
2. There is no formal system that determines what is right, but there is an ongoing structure of discourse within a community, in which rightness is the issue, and in which there are stable areas of agreement.
3. An individual is never fully aware of what is possible to do, what effects an act will have, or how it will be assessed, and nevertheless will continue to act.

In a way this is comforting and in a way it is challenging. It is comforting because it does not posit some unachievable ideal: either the ideal of always having the strength to clothe the right thing, or the ideal of being able to determine just what is the right thing. Instead it puts the emphasis on being committed to entering into discussion with others and to taking seriously their concerns and understandings.

At the same time, for the same reason, it is challenging. It says that we will never have the satisfaction of knowing exactly when to apply our social and ethical concerns, or being confident that “Now we have it figured out.” We are always being thrown into activity which may, in unanticipated ways, have implications for values, and we are part of a community that is always responsible for evolving new understandings and ways of “juggling” to maintain those values. It’s exciting, but at times can be a little shaky.

### **What does it mean to *do* ethics and social responsibility?**

In the final section, I want to look at what all this implies for the kinds of activities engaged in by individuals and organizations committed to making connections between computing and values.

All along I have been emphasizing the “doing” side of ethics: the way in which our actions more than our deliberations speak to our values. There are three key components in “doing” ethics and social responsibility:

1. Identifying social/ethical issues
2. Entering into serious discourse about the possibilities, -with yourself and with others
3. Taking actions

Each of these has both an individual and a social component. There are cases where one person alone is faced with recognizing a problem, considering what to do, and doing it. Many of the most powerful pieces of literature in our culture grip us because they let us feel what it is like to wrestle with this ultimate responsibility of the individual.

But in this paper I want to focus more on the ways in which each of these components is situated in the actions of larger groups, and in particular the kinds of organizations represented by NCCV and by CPSR.

The activities with which I am most familiar are the work over the last ten years of Computer Professionals for Social Responsibility. Initially we were motivated by what appeared to be a mad rush towards nuclear war on the part of our government. As with man, groups that emerged in the early 1980s with the words “social responsibility” in their names, we felt that the only responsible thing to do in that climate was to work wherever we had the most possibility of influence, in order to avert catastrophe.

In fact, there were many clear connections between computing and nuclear war, and in particular we came to focus heavily on the proposed Strategic Defense Initiative, or “Star wars” system. As computer professionals, we were sensitized to the problems of reliability and complexity that made the plans unrealistic, and could lead both to

tremendous wasted resources and to a false sense of security then could dangerously destabilize the nuclear situation.

We approached this problem in a number of ways. CPSR members, individually and together, wrote papers analyzing the problems and bringing them to the attention of policy makers, both directly and through publication in newspapers and magazines. A number of us were active in the movement to have scientists pledge not to take research money from the SDI office. It was clear that the promise of research funds was being used as a lure to get tacit approval from the scientific community for the project. One person in particular who recognized this was David Parnas, who resigned from the panel that was convened to develop an analysis of the computing requirements, and later made public his analysis of the problems and of the ways in which the Pentagon was trying to influence scientists' assessments.

I mention all this not just because it is an important piece of CPSR's history, but to illustrate the range of activities that constituted "ethicking" for people involved with the issue. Some of the actions, like Parnas' resignation, can be viewed as explicit and difficult ethical choices made by individuals, and are noteworthy as such. But the impact they had was magnified by the fact that they were part of a coordinated campaign, in the context of organizations that could bring them to public attention and connect them the basic issues being fought. If one scientist quietly decides not to do SDI research, the impact is on that person and his other work. If a whole community is involved, not only is the political impact greater, but the thinking of the entire profession is moved. The ability to recognize the potential of issues with consequences for values and ethics is increased for all those who become party to the discussion, even if they don't take direct action at the moment.

As the imminent danger of nuclear war appeared to subside, CPSR was able to take into fuller consideration the recognition that "social responsibility" really does cover more than just preventing annihilation (even though that is certainly a good place to start). Our concerns have overlapped strongly with those that form the core of the NCCV conference: Privacy, security, ownership and fair access. For example, we recently took part in a successful nationwide campaign to block the sale of a product proposed by Lotus that would have made it possible to find out private information about millions of consumers. In that case computers played a helping role as well, with much of the education and awareness about the issue being transmitted by electronic mail.

There isn't sufficient space to go into detail here on all of the CPSR program areas, but I have recounted this history to serve as an example as we look back at the three components of "ethicking" listed above:

### *1. Identify social/ethical issues*

The first is to identify social and ethical issues to which computers and computing are relevant. Each of us needs to ask how our actions as a computer professional might have ethical and social consequences, and there are a variety of answers. In some cases, as with every profession, the consequences grow directly from specifics of conduct. If I am asked to write a piece of life-critical software then I must proceed in a way that is

responsible: taking care as best I am able given current software practices, and being honest with the clients about the risks and limitations. If I do less, I am cheating.

In some cases, harmful consequences come not from the quality of the work itself, but from the uses to which it will be put. If I am asked to build a program that makes it possible for employers to invisibly monitor the details of a worker's activities, I need to be aware of the consequences such programs can have in the workplace. Often, there is no clear boundary between beneficial and harmful technologies. A data base system used by the FBI to track drug dealers may have a positive effect on reducing the drug traffic, but can also be used to keep track of people with unpopular political beliefs. Even then, if I see the danger as great enough (what if those beliefs are racist and violent?) it may be a net gain to society if I produce such a system. But where is the boundary?

To go a step farther, there will be cases where the work itself is positive, but there is a larger context in which it can play a harmful role. Much of the research sponsored by the Strategic Defense Initiative is of a general kind that most of us would assess as having positive applications. The developing of computer networking, although originally sponsored by military agencies, has had a tremendous affect on our ability to function as computer professionals, and is rapidly becoming available to the entire population.

But what about the larger context? When a General testifies before congress that the scientific community is in favor of the SDI plans, as evidenced by the number of them who are actively working on research for it, what consequences has our research had? When a tremendous proportion of the research in computing in general is directed by the military, what long-term effect does this have on the kind of problems that are posed, or on the role of military thinking on the direction of our national economy?<sup>2</sup>

There are rarely easy answers to such questions. In order to make responsible decisions about values, an individual needs a broad understanding of the consequences his or her actions might have in this overall situation. Such an understanding develops only through extended open discussion that brings in people from outside the computer profession as well as within it. It also extends beyond those who engage directly in it. The "styrofoam professor" I mentioned earlier has become conscious of issues of research funding through having interacted with me about my own rejection of military sponsorship. Even if he disagrees, the fact that the discussion exists (and has engaged his students as well) gives it a new standing in his "moral calculation."

## *2. Enter into serious discourse about the possibilities*

The second step in "ethicking" was to enter into serious discourse about the possibilities. I use the word "discourse" here instead of "thinking" to emphasize the social construction that is at the heart of decision-making even when a person does not directly enter into conversation with others. In a real, if extended sense I am in discourse not only with the people I speak with but with those who have written the things that have influenced me, and those I have talked with, and in turn those in the future who will be influenced by what I say and write.

This includes people within the computer profession, and also in the larger society within which we work. The job of “public education” is a key part of creating the background of expectations that constitutes the fabric of ethical and social responsibility.

It should be obvious in looking at this conference, both at the participants and the materials that have been prepared and will be produced, that the weaving of the discourse is a function of groups of people who gather together (literally or through communications media) to think things through (or should we say “talk things through”?). This is a key role played by institutions and organizations devoted to issues of computers and ethics.

### *3. Take actions*

Finally, the bottom line is the actions we take, both individually and collectively. It would be futile to try to catalog all of the different kinds of actions that have ethical implications. There are obvious individual acts such as whistleblowing, in which a clear value statement is being made in spite of some potential damage or loss to the actor. There are many other acts, such as choosing whether or not to work on a particular project or to take a job with a company that pursues projects of social concern, in which the decision is more subtle and the ethical factor may be one of many, which cannot be untangled in looking back at why the decision was made.

As members of a profession and its professional organizations, we also take acts that are intended to affect the direction and activity of the profession as a whole. They often don't have the visceral quality of whistleblowing or rejection of funding, but they contribute to creating the atmosphere in which those acts can be given sense. These include teaching, public education, working with professional organizations, developing standards and many other forms of everyday “ethicking.”

Some of these activities have an overt political objective, such as lobbying for legislation or providing expertise to lawmaking and judicial bodies. Others operate at a broader cultural level, helping people both in the profession and in the public learn to see the issues, understand their consequences and apply human values to technical decisions. Although it is easier to get an individual to consider “Is it ethical (or socially responsible) to work on a bomb project?” It is equally important and more frequently relevant to ask “Is it ethical not to contribute my part to being responsible for how the public and the profession guide the ways in which computers will be used?” At the CPSR annual meeting last year, the slogan on the posters was “Technology is driving the future. It's time to find out who's steering.” In the end, we all have a hand on the wheel.

### **Conclusion**

Looking back at the three cartoons we might come up with different views of what we are engaged in when we participate in a conference on Computing and Values. From the angel/devil perspective, we could see it as a “revival meeting” at which we encourage each other to act in accord with our values, and tell stories that will help us to be resolute and remain steadfast. From the morality computer perspective, the conference is a think

tank: it is our job to come up with the right rules and descriptions that will form the knowledge base that computer professionals can use to figure out what they should do.

These both have some particle of truth, but I much prefer to see our activity as a working session in which we are engaged in juggling the issues, ideas and discussions that generate the world of possibilities in which we and our colleagues live and work. We are creating those possibilities, increasing our own understanding and commitment to their value, and building a community that can continue to create and learn in the future.

### **Notes**

1. For a wide variety of cases, see the papers in Ermann and Williams, Dunlop and Kling, 1991, Parker, Swope and Baker, 1990, and Johnson and Snapper, 1985. For a list of syllabi covering a broad range of topics related to computing and ethics, see Friedman and Winograd, 1989.
2. For further discussion of the issues of military funding, see Winograd, 1989 and the papers in Mitcham and Siekevitz, 1989.

### **References**

1. Bynum, Terrell, Human values and the computer science curriculum, in T. Bynum, W. Maner and J. Fodor (eds.) *Computing and Human Values*, RCCS: 1992.
2. Bynum, Terrell Ward, Walter Maner, and John L. Fodor (eds.), *Computing and Human Values: Proceedings of the 1991 Conference*, New Haven: Research Center on Computing and Society, 1992.
3. Dunlop, Charles and Rob Kling (eds.), *Computerization and Controversy*, Boston: Academic Press, 1991
4. Friedman, Batya and Terry Winograd (eds.), *Computing and Social Responsibility: a Collection of Course Syllabi*, Palo Alto: Computer Professionals for Social Responsibility, 1989.
5. Johnson, Deborah G. and John W. Snapper, *Ethical Issues In the Use of Computers*, Belmont, CA: Wadsworth Publishing Company, 1985.
6. Mitcham, Carl and Philip Siekevitz (eds.), *Ethical Issues Associated with Scientific and Technological Research for the Military*, New York, NY, New York Academy of Sciences, 1989.
7. Parker, Donn, Susan Swope and Bruce Baker, *Ethical Conflicts in Information & Computer Science, Technology and Business*, Wellesley, MA: QED Information Sciences, 1990.
8. Winograd, Terry, Strategic computing research and the universities, in J. Jacky and D. Schuler, eds., *Directions and Implications of Advanced Computing, Volume I*. Norwood NJ: Ablex, 1989,18-32. Reprinted in Dunlop and Kling, 1991, pp. 700-716.

Part 2  
Scientific Ethics and the Military



# Letter to President Roosevelt

—Albert Einstein, 1939

Albert Einstein  
Old Grove Rd.  
Nassau Point  
Peconic, Long Island

August 2nd, 1939

F.D. Roosevelt,  
President of the United States,  
White House  
Washington, D.C.

Sir:

Some recent work by E. Fermi and L. Szilard, which has been communicated to me in manuscript, leads me to expect that the element uranium may be turned into a new and important source of energy in the immediate future. Certain aspects of the situation which has arisen seem to call for watchfulness and, if necessary, quick action on the part of the Administration. I believe therefore that it is my duty to bring to your attention the following facts and recommendations:

In the course of the last four months it has been made probable - through the work of Joliot in France as well as Fermi and Szilard in America - that it may become possible to set up a nuclear chain reaction in a large mass of uranium, by which vast amounts of power and large quantities of new radium-like elements would be generated. Now it appears almost certain that this could be achieved in the immediate future.

This new phenomenon would also lead to the construction of bombs, and it is conceivable - though much less certain - that extremely powerful bombs of a new type may thus be constructed. A single bomb of this type, carried by boat and exploded in a port, might very well destroy the whole port together with some of the surrounding territory. However, such bombs might very well prove to be too heavy for transportation by air.

The United States has only very poor ores of uranium in moderate quantities. There is some good ore in Canada and the former Czechoslovakia, while the most important source of uranium is Belgian Congo.

In view of this situation you may think it desirable to have some permanent contact maintained between the Administration and the group of physicists working on chain reactions in America. One possible way of achieving this might be for you to entrust with this task a person who has your confidence and who could perhaps serve in an unofficial capacity. His task might comprise the following:

a) to approach Government Departments, keep them informed of the further development, and put forward recommendations for Government action, giving particular attention to the problem of securing a supply of uranium ore for the United States;

b) to speed up the experimental work, which is at present being carried on within the limits of the budgets of University laboratories, by providing funds, if such funds be required, through his contacts with private persons who are willing to make contributions for this cause, and perhaps also by obtaining the co-operation of industrial laboratories which have the necessary equipment.

I understand that Germany has actually stopped the sale of uranium from the Czechoslovakian mines which she has taken over. That she should have taken such early action might perhaps be understood on the ground that the son of the German Under-Secretary of State, von Weizsäcker, is attached to the Kaiser-Wilhelm-Institut in Berlin where some of the American work on uranium is now being repeated.

Yours very truly,

*A. Einstein*  
(Albert Einstein)

# Reply to Albert Einstein

—President Roosevelt, 1939

C O P Y

October 19, 1939

My dear Professor:

I want to thank you for your recent letter and  
and the most interesting and important enclosure.

I found this data of such import that I have  
convened a Board consisting of the head of the Bureau  
of Standards and a chosen representative of the Army  
and Navy to thoroughly investigate the possibilities  
of your suggestion regarding the element of uranium.

I am glad to say that Dr. Sachs will cooperate  
and work with this Committee and I feel this is the  
most practical and effective method of dealing with  
the subject.

Please accept my sincere thanks.

Very sincerely yours,

(signed) Franklin D. Roosevelt

Dr. Albert Einstein,  
Old Grove Road,  
Nassau Point,  
Peconic, Long Island,  
New York.

C O P Y

# A Petition to the President of the United States

—Leo Szilard *et al.*, July 17, 1945

Discoveries of which the people of the United States are not aware may affect the welfare of this nation in the near future. The liberation of atomic power which has been achieved places atomic bombs in the hands of the Army. It places in your hands, as Commander-in-Chief, the fateful decision whether or not to sanction the use of such bombs in the present phase of the war against Japan.

We, the undersigned scientists, have been working in the field of atomic power. Until recently, we have had to fear that the United States might be attacked by atomic bombs during this war and that her only defense might lie in a counterattack by the same means. Today, with the defeat of Germany, this danger is averted and we feel impelled to say what follows:

The war has to be brought speedily to a successful conclusion and attacks by atomic bombs may very well be an effective method of warfare. We feel, however, that such attacks on Japan could not be justified, at least not unless the terms which will be imposed after the war on Japan were made public in detail and Japan were given an opportunity to surrender.

If such public announcement gave assurance to the Japanese that they could look forward to a life devoted to peaceful pursuits in their homeland and if Japan still refused to surrender our nation might then, in certain circumstances, find itself forced to resort to the use of atomic bombs. Such a step, however, ought not to be made at any time without seriously considering the moral responsibilities which are involved.

The development of atomic power will provide the nations with new means of destruction. The atomic bombs at our disposal represent only the first step in this direction, and there is almost no limit to the destructive power which will become available in the course of their future development. Thus a nation which sets the precedent of using these newly liberated forces of nature for purposes of destruction may have to bear the responsibility of opening the door to an era of devastation on an unimaginable scale.

If after this war a situation is allowed to develop in the world which permits rival powers to be in uncontrolled possession of these new means of destruction, the cities of the United States as well as the cities of other nations will be in continuous danger of sudden annihilation. All the resources of the United States, moral and material, may have to be mobilized to prevent the advent of such a world situation. Its prevention is at present the solemn responsibility of the United States—singled out by virtue of her lead in the field of atomic power.

The added material strength which this lead gives to the United States brings with it the obligation of restraint and if we were to violate this obligation our moral position would be weakened in the eyes of the world and in our own eyes. It would then be more difficult for us to live up to our responsibility of bringing the unloosened forces of destruction under control.

In view of the foregoing, we, the undersigned, respectfully petition: first, that you exercise your power as Commander-in-Chief, to rule that the United States shall not resort to the use of atomic bombs in this war unless the terms which will be imposed upon Japan have been made public in detail and Japan knowing these terms has refused to surrender; second, that in such an event the question whether or not to use atomic bombs be decided by you in light of the considerations presented in this petition as well as all the other moral responsibilities which are involved.

R. S. S. Rupp  
D. D. Mulliken  
E. P. Wigner  
Georges. Menon  
Leo Szilard

J. G. Nilsen  
J. H. Zacherian  
Francis R. Shorke  
John A. Simpson  
Walter Bartley  
John R. Howe

Frankly Forte

# The Pugwash Manifesto

—Bertrand Russell, Albert Einstein, *et al.*, 1955

IN the tragic situation which confronts humanity, we feel that scientists should assemble in conference to appraise the perils that have arisen as a result of the development of weapons of mass destruction, and to discuss a resolution in the spirit of the appended draft.

We are speaking on this occasion, not as members of this or that nation, continent, or creed, but as human beings, members of the species Man, whose continued existence is in doubt. The world is full of conflicts; and, overshadowing all minor conflicts, the titanic struggle between Communism and anti-Communism.

Almost everybody who is politically conscious has strong feelings about one or more of these issues; but we want you, if you can, to set aside such feelings and consider yourselves only as members of a biological species which has had a remarkable history, and whose disappearance none of us can desire.

We shall try to say no single word which should appeal to one group rather than to another. All, equally, are in peril, and, if the peril is understood, there is hope that they may collectively avert it.

We have to learn to think in a new way. We have to learn to ask ourselves, not what steps can be taken to give military victory to whatever group we prefer, for there no longer are such steps; the question we have to ask ourselves is: what steps can be taken to prevent a military contest of which the issue must be disastrous to all parties?

The general public, and even many men in positions of authority, have not realized what would be involved in a war with nuclear bombs. The general public still thinks in terms of the obliteration of cities. It is understood that the new bombs are more powerful than the old, and that, while one A-bomb could obliterate Hiroshima, one H-bomb could obliterate the largest cities, such as London, New York, and Moscow.

No doubt in an H-bomb war great cities would be obliterated. But this is one of the minor disasters that would have to be faced. If everybody in London, New York, and Moscow were exterminated, the world might, in the course of a few centuries, recover from the blow. But we now know, especially since the Bikini test, that nuclear bombs can gradually spread destruction over a very much wider area than had been supposed.

It is stated on very good authority that a bomb can now be manufactured which will be 2,500 times as powerful as that which destroyed Hiroshima. Such a bomb, if exploded near the ground or under water, sends radio-active particles into the upper air. They sink gradually and reach the surface of the earth in the form of a deadly dust or rain. It was this dust which infected the Japanese fishermen and their catch of fish. No one knows

---

This manifesto was issued in July of 1955. The group of signers went on to found the Pugwash Conferences on Science and World Affairs.

how widely such lethal radio-active particles might be diffused, but the best authorities are unanimous in saying that a war with H-bombs might possibly put an end to the human race. It is feared that if many H-bombs are used there will be universal death, sudden only for a minority, but for the majority a slow torture of disease and disintegration.

Many warnings have been uttered by eminent men of science and by authorities in military strategy. None of them will say that the worst results are certain. What they do say is that these results are possible, and no one can be sure that they will not be realized. We have not yet found that the views of experts on this question depend in any degree upon their politics or prejudices. They depend only, so far as our researches have revealed, upon the extent of the particular expert's knowledge. We have found that the men who know most are the most gloomy.

Here, then, is the problem which we present to you, stark and dreadful and inescapable: Shall we put an end to the human race; or shall mankind renounce war? People will not face this alternative because it is so difficult to abolish war.

The abolition of war will demand distasteful limitations of national sovereignty. But what perhaps impedes understanding of the situation more than anything else is that the term "mankind" feels vague and abstract. People scarcely realize in imagination that the danger is to themselves and their children and their grandchildren, and not only to a dimly apprehended humanity. They can scarcely bring themselves to grasp that they, individually, and those whom they love are in imminent danger of perishing agonizingly. And so they hope that perhaps war may be allowed to continue provided modern weapons are prohibited.

This hope is illusory. Whatever agreements not to use H-bombs had been reached in time of peace, they would no longer be considered binding in time of war, and both sides would set to work to manufacture H-bombs as soon as war broke out, for, if one side manufactured the bombs and the other did not, the side that manufactured them would inevitably be victorious.

Although an agreement to renounce nuclear weapons as part of a general reduction of armaments would not afford an ultimate solution, it would serve certain important purposes. First, any agreement between East and West is to the good in so far as it tends to diminish tension. Second, the abolition of thermo-nuclear weapons, if each side believed that the other had carried it out sincerely, would lessen the fear of a sudden attack in the style of Pearl Harbour, which at present keeps both sides in a state of nervous apprehension. We should, therefore, welcome such an agreement though only as a first step.

Most of us are not neutral in feeling, but, as human beings, we have to remember that, if the issues between East and West are to be decided in any manner that can give any possible satisfaction to anybody, whether Communist or anti-Communist, whether Asian or European or American, whether White or Black, then these issues must not be decided by war. We should wish this to be understood, both in the East and in the West.

There lies before us, if we choose, continual progress in happiness, knowledge, and wisdom. Shall we, instead, choose death, because we cannot forget our quarrels? We appeal as human beings to human beings: Remember your humanity, and forget the rest. If you can do so, the way lies open to a new Paradise; if you cannot, there lies before you the risk of universal death.

**Resolution:**

WE invite this Congress, and through it the scientists of the world and the general public, to subscribe to the following resolution:

“In view of the fact that in any future world war nuclear weapons will certainly be employed, and that such weapons threaten the continued existence of mankind, we urge the governments of the world to realize, and to acknowledge publicly, that their purpose cannot be furthered by a world war, and we urge them, consequently, to find peaceful means for the settlement of all matters of dispute between them.”

Max Born  
Percy W. Bridgman  
Albert Einstein  
Leopold Infeld

Frederic Joliot-Curie  
Herman J. Muller  
Linus Pauling  
Cecil F. Powell

Joseph Rotblat  
Bertrand Russell  
Hideki Yukawa

## Remember Your Humanity

—Joseph Rotblat, 1995

At this momentous event in my life—the acceptance of the Nobel Peace Prize—I want to speak as a scientist, but also as a human being. From my earliest days I had a passion for science. But science, the exercise of the supreme power of the human intellect, was always linked in my mind with benefit to people. I saw science as being in harmony with humanity. I did not imagine that the second half of my life would be spent on efforts to avert a mortal danger to humanity created by science.

The practical release of nuclear energy was the outcome of many years of experimental and theoretical research. It had great potential for the common good. But the first the general public learned about this discovery was the news of the destruction of Hiroshima by the atom bomb. A splendid achievement of science and technology had turned malign. Science became identified with death and destruction.

It is painful to me to admit that this depiction of science was deserved. The decision to use the atom bomb on Japanese cities, and the consequent buildup of enormous nuclear arsenals, was made by governments on the basis of political and military perceptions. But scientists on both sides of the Iron Curtain played a very significant role in maintaining the momentum of the nuclear arms race throughout the four decades of the Cold War.

The role of scientists in the nuclear arms race was expressed by Lord [Solly] Zuckerman, for many years chief scientific adviser to the British Government: “When it comes to nuclear weapons . . . it is the man in the laboratory who at the start proposes that for this or that arcane reason it would be useful to improve an old or to devise a new nuclear warhead. It is he, the technician, not the commander in the field, who is at the heart of the arms race.”

Long before the terrifying potential of the arms race was recognized, there was a widespread instinctive abhorrence of nuclear weapons, and a strong desire to get rid of them. Indeed, the very first resolution of the General Assembly of the United Nations—adopted unanimously—called for the elimination of nuclear weapons. But the world was then polarized by the bitter ideological struggle between East and West. There was no chance to meet this call. The chief task was to stop the arms race before it brought utter disaster. However, after the collapse of communism and the disintegration of the Soviet Union, any rationale for having nuclear weapons disappeared. The quest for their total elimination could be resumed. But the nuclear powers still cling tenaciously to their weapons.

Let me remind you that nuclear disarmament is not just an ardent desire of the people, as expressed in many resolutions of the United Nations. It is a legal commitment by the five official nuclear states, entered into when they signed the Nuclear Non-Proliferation Treaty (NPT). Only a few months ago, when the indefinite extension of the treaty was agreed, the nuclear powers committed themselves again to complete nuclear

---

Joseph Rotblat delivered this address in Stockholm when he received, together with the Pugwash Conferences on Science and World Affairs, the Nobel Peace Prize for 1995.

disarmament. This is still their declared goal. But the declarations are not matched by their policies, and this divergence seems to be intrinsic.

Since the end of the Cold War the two main nuclear powers have begun to make big reductions in their nuclear arsenals. Each of them is dismantling about 2,000 nuclear warheads a year. If this program continued, all nuclear warheads could be dismantled in little over ten years from now. We have the technical means to create a nuclear-weapon-free world in about a decade. Alas, the present program does not provide for this. When the START II Treaty has been implemented—and remember it has not yet been ratified—we will be left with some 15,000 nuclear warheads, active and in reserve. Fifteen thousand weapons with an average yield of 20 Hiroshima bombs.

Unless there is a change in basic philosophy, we will not see a reduction of nuclear arsenals to zero for a very long time, if ever. The present basic philosophy is nuclear deterrence. This was stated clearly in the U.S. Nuclear Posture Review which concluded that the “post-Cold War environment requires nuclear deterrence,” and this is echoed by other nuclear states. Nuclear weapons are kept as a hedge against some unspecified dangers.

This policy is simply an inertial continuation from the Cold War era. The Cold War is over, but Cold War thinking survives. Then, we were told that a world war was prevented by the existence of nuclear weapons. Now, we are told that nuclear weapons prevent all kinds of war. These are arguments that purport to prove a negative. I am reminded of a story told in my boyhood, at the time when radio communication began.

“Two wise men were arguing about the ancient civilization in their respective countries. One said: ‘My country has a long history of technological development: We have carried out deep excavations and found a wire, which shows that already in the old days we had the telegraph.’ The other man retorted: ‘We too made excavations; we dug much deeper than you and found . . . nothing, which proves that already in those days we had wireless communication!’”

There is no direct evidence that nuclear weapons prevented a world war. Conversely, it is known that they nearly caused one. The most terrifying moment in my life was October 1962, during the Cuban missile crisis. I did not know all the facts—we have learned only recently how close we were to war—but I knew enough to make me tremble. The lives of millions of people were about to end abruptly; millions of others were to suffer a lingering death; much of our civilization was to be destroyed. It all hung on the decision of one man, Nikita Khrushchev: would he or would he not yield to the U.S. ultimatum? This is the reality of nuclear weapons: they may trigger a world war; a war which, unlike previous ones, destroys all of civilization.

As for the assertion that nuclear weapons prevent wars, how many more wars are needed to refute this argument? Tens of millions have died in the many wars that have taken place since 1945. In a number of them nuclear states were directly involved. In two they were actually defeated. Having nuclear weapons was of no use to them.

To sum up, there is no evidence that a world without nuclear weapons would be a dangerous world. On the contrary, it would be a safer world, as I will show later.

We are told that the possession of nuclear weapons—in some cases even the testing of these weapons—is essential for national security. But this argument can be made by other countries as well. If the militarily most powerful—and least threatened—states need nuclear weapons for their security, how can one deny such security to countries that are truly insecure? The present nuclear policy is a recipe for proliferation. It is a policy for disaster.

To prevent this disaster—for the sake of humanity—we must get rid of all nuclear weapons.

Achieving this goal will take time, but it will never happen unless we make a start. Some essential steps towards it can be taken now. Several studies, and a number of public statements by senior military and political personalities, testify that—except for disputes between the present nuclear states—all military conflicts, as well as threats to peace, can be dealt with using conventional weapons. This means that the only function of nuclear weapons, while they exist, is to deter a nuclear attack.

All nuclear weapon states should now recognize that this is so, and declare—in treaty form—that they will never be the first to use nuclear weapons. This would open the way to the gradual, mutual reduction of nuclear arsenals, down to zero. It would also open the way for a nuclear weapons convention. This would be universal—it would prohibit all possession of nuclear weapons.

We will need to work out the necessary verification system to safeguard the convention. A Pugwash study produced suggestions on these matters. The mechanism for negotiating such a convention already exists. Entering into negotiations does not commit the parties. There is no reason why they should not begin now. If not now, when?

So I ask the nuclear powers to abandon the out-of-date thinking of the Cold War period and take a fresh look. Above all, I appeal to them to bear in mind the long-term threat that nuclear weapons pose to humankind and to begin action towards their elimination. Remember your duty to humanity

My second appeal is to my fellow scientists. I described earlier the disgraceful role played by a few scientists, caricatured as “Dr. Strangeloves,” in fueling the arms race. They did great damage to the image of science.

On the other side there are the scientists, in Pugwash and other bodies, who devote much of their time and ingenuity to averting the dangers created by advances in science and technology. However, they embrace only a small part of the scientific community. I want to address the scientific community as a whole.

You are doing fundamental work, pushing forward the frontiers of knowledge, but often you do it without giving much thought to the impact of your work on society. Precepts such as “science is neutral” or “science has nothing to do with politics,” still prevail. They are remnants of the ivory tower mentality, although the ivory tower was finally demolished by the Hiroshima bomb.

Here, for instance, is a question: Should any scientist work on the development of weapons of mass destruction? A clear “no” was the answer recently given by Hans Bethe.

Professor Bethe, a Nobel laureate, is the most senior of the surviving members of the Manhattan Project. On the occasion of the fiftieth anniversary of Hiroshima, he issued a statement that I will quote in full:

As the director of the Theoretical Division of Los Alamos, I participated at the most senior level in the World War II project that produced the first atomic weapons.

Now, at age 88, I am one of the few remaining such senior persons alive. Looking back at the half century since that time, I feel the most intense relief that these weapons have not been used since World War II, mixed with the horror that tens of thousands of such weapons have been built since that time—one hundred times more than any of us at Los Alamos could ever have imagined.

Today we are rightly in an era of disarmament and dismantlement of nuclear weapons. But in some countries nuclear weapons development still continues. Whether and when the various nations of the world can agree to stop this is uncertain. But individual scientists can still influence this process by withholding their skills.

Accordingly, I call on all scientists in all countries to cease and desist from work creating, developing, improving and manufacturing further nuclear weapons—and, for that matter, other weapons of potential mass destruction such as chemical and biological weapons.

If all scientists heeded this call there would be no more new nuclear warheads; no French scientists at Moruroa; no new chemical and biological poisons. The arms race would be over.

But there are other areas of scientific research that may directly or indirectly lead to harm to society. This calls for constant vigilance. The purpose of some government or industrial research is sometimes concealed, and misleading information is presented to the public. It should be the duty of scientists to expose such malfeasance. “Whistle-blowing” should become part of the scientist’s ethos. This may bring reprisals; a price to be paid for one’s convictions. The price may be very heavy, as illustrated by the disproportionately severe punishment of Mordechai Vanunu. I believe he has suffered enough.

The time has come to formulate guidelines for the ethical conduct of scientists, perhaps in the form of a voluntary Hippocratic Oath. This would be particularly valuable for young scientists when they embark on a scientific career. The U.S. Student Pugwash Group has taken up this idea—and that is very heartening.

At a time when science plays such a powerful role in the life of society, when the destiny of the whole of mankind may hinge on the results of scientific research, it is incumbent on all scientists to be fully conscious of that role, and conduct themselves accordingly. I appeal to my fellow scientists to remember their responsibility to humanity.

My third appeal is to my fellow citizens in all countries: Help us to establish lasting peace in the world.

I have to bring to your notice a terrifying reality: With the development of nuclear weapons man has acquired, for the first time in history, the technical means to destroy the whole of civilization in a single act. Indeed, the whole human species is endangered, by nuclear weapons or by other means of wholesale destruction which further advances in science are likely to produce.

I have argued that we must eliminate nuclear weapons. While this would remove the immediate threat, it will not provide permanent security. Nuclear weapons cannot be disinvented. The knowledge of how to make them cannot be erased. Even in a nuclear-weapon-free world, should any of the great powers become involved in a military confrontation, they would be tempted to rebuild their nuclear arsenals. That would still be a better situation than the one we have now, because the rebuilding would take a considerable time, and in that time the dispute might be settled. A nuclear-weapon-free world would be safer than the present one. But the danger of the ultimate catastrophe would still be there.

The only way to prevent it is to abolish war altogether. War must cease to be an admissible social institution. We must learn to resolve our disputes by means other than military confrontation.

This need was recognized 40 years ago when we said in the Russell-Einstein Manifesto: "Here then is the problem which we present to you, stark and dreadful, and inescapable: shall we put an end to the human race: or shall mankind renounce war?"

The abolition of war is also the commitment of the nuclear weapon states: Article VI of the NPT calls for a treaty on general and complete disarmament under strict and effective international control.

Any international treaty entails some surrender of national sovereignty, and is generally unpopular. As we said in the Russell-Einstein Manifesto: "The abolition of war will demand distasteful limitations of national sovereignty." Whatever system of governance is eventually adopted, it is important that it carries the people with it. We need to convey the message that safeguarding our common property, humankind, will require developing in each of us a new loyalty: a loyalty to mankind. It calls for the nurturing of a feeling of belonging to the human race. We have to become world citizens.

Notwithstanding the fractionation that has occurred since the end of the Cold War, and the many wars for recognition of national or ethnic identities, I believe that the prospects for the acceptance of this new loyalty are now better than at the time of the Russell-Einstein Manifesto. This is so largely because of the enormous progress made by science and technology during these 40 years. The fantastic advances in communication and transportation have shrunk our globe. All nations of the world have become close neighbors. Modern information techniques enable us to learn instantly about every event in every part of the globe. We can talk to each other via the various networks. This facility will improve enormously with time, because the achievements so far have only scratched the surface. Technology is driving us together. In many ways we are becoming like one family.

In advocating the new loyalty to mankind I am not suggesting that we give up national loyalties. Each of us has loyalties to several groups—from the smallest, the family, to the largest, at present, the nation. Many of these groups provide protection for their members. With the global threats resulting from science and technology, the whole of humankind now needs protection.

We have to extend our loyalty to the whole of the human race.

What we are advocating in Pugwash, a war-free world, will be seen by many as a utopian dream. It is not utopian. There already exist in the world large regions, for example, the European Union, within which war is inconceivable. What is needed is to extend these to cover the world's major powers.

In any case, we have no choice. The alternative is unacceptable. Let me quote the last sentence of the Russell-Einstein Manifesto: "We appeal, as human beings, to human beings: Remember your humanity and forget the rest. If you can do so, the way lies open for a new paradise; if you cannot, there lies before you the risk of universal death."

The quest for a war-free world has a basic purpose: survival. But if in the process we learn how to achieve it by love rather than by fear, by kindness rather than by compulsion; if in the process we learn to combine the essential with the enjoyable, the expedient with the benevolent, the practical with the beautiful, this will be an extra incentive to embark on this great task.

And above all, remember your humanity.



# UNIVERSITY OF VICTORIA

PO BOX 1700, VICTORIA. BRITISH COLUMBIA. CANADA V8W 2Y2  
TELEPHONE (604) 721-7211, TELEX 049-7222

*Department of Computer Science*  
721-7209

28 June, 1985

Mr. James H. Offut  
Assistant Director, BM/C3  
Strategic Defense Initiative Organization Office of the Secretary of Defense  
Washington, D.C.  
20301

Dear Mr. Offut:

Thank you for your letter of 5 June 1985 appointing me a member of the SDIO Panel on Computing in Support of Battle Management. I appreciate the recognition implicit in being chosen as one of your expert advisors on computer science.

After attending the first meeting of the panel and giving the problem considerable thought, I am resigning my membership in the panel. I do not believe that further work by the panel will be useful and I cannot, in good conscience, accept further payment for useless effort.

The panel's work will not be useful for two reasons.

- 1) The goals stated for the Strategic Defense System cannot be attained by the class of systems that you are considering.
- 2) The SDIO is not the appropriate organization to fund and administer the research it is supporting. Most of the money spent will be wasted. The panel on which you have asked me to serve, is not appropriately constituted, clearly chartered, and adequately informed. There are better ways to select and manage research.

My conclusions are not based on political or policy judgements. Unlike many other academic critics of the SDI effort, I have not, in the past, objected to defense efforts or defense sponsored research. I have been deeply involved in such research and have consulted extensively on defense projects. My conclusions are based on more than 20 years of research on software engineering including more than 8 years of work on real-time software used in military aircraft. They are based on familiarity with both operational military software and computer science research. My conclusions are based on characteristics peculiar to this particular effort, not objections to weapons development in general.

Before making my decision and writing this letter I have carefully reconsidered what I have learned in my own research area and I have reviewed reports of work in related fields. These reviews lead inevitably to the judgements stated above. I am willing to stake my professional reputation on my conclusions.

Enclosed with this letter are several brief papers (1 - 2 pages each) summarizing my observations and substantiating the conclusions stated above. Their purpose is to explain my decision.

These papers explain:

- 1) The fundamental technological differences between software engineering and other areas of engineering and why software is unreliable,
- 2) The properties of the proposed SDI software that make it unattainable,
- 3) Why the techniques commonly used to build military software are inadequate for this job,
- 4) The nature of research in Software Engineering, and why the improvements that it can effect will not be sufficient to allow construction of a truly reliable strategic defense system,
- 5) The nature of research in Artificial Intelligence, and why I do not expect it to help in building reliable military software,
- 6) The history of research in Automatic Programming, and why I do not expect it to bring about the substantial improvements that are needed,
- 7) Why Program Verification cannot give us a reliable strategic defense battle management software system,
- 8) My opinions on the management of applied research, why I consider this panel and the SDIO in general to be an inappropriate vehicle for funding research, and what I would do instead.

I am quite certain that you will be able to find software experts who disagree with my conclusions. For many, the project offers a source of funding, funding that will enrich some personally, while offering others new and generous support for their personal research projects. During the first sittings of our panel, I could see the dollar figures dazzling everyone involved. Almost everyone that I know within the military industrial complex sees in the SDI a new "pot of gold" just waiting to be tapped.

For others, the project offers an unending set of technological puzzles that are fun to work on; such problems are exciting and challenging whether or not the work ever produces useful results. Almost every software expert that I know, entered the field because they enjoy this kind of challenge. Several of the speakers at the first meeting of our panel could not hide their delight at the unbounded set of technical challenges implicit in the unattainable goals of the project.

I can tell you, as one who likes both money and technical challenges, that these temptations are very hard to resist. You will find it very hard to find unbiased expert opinions on this issue.

In March 1983 the President asked us, as members of the scientific community, to provide the means of rendering nuclear weapons impotent and obsolete. I believe that it is our duty, as scientists and engineers, to reply that we have no technological magic that will accomplish that. The short term applied research and focussed development that SDI is now funding is not going to solve the problem; the President and the public should know that.

Yours truly,

David L. Parnas Lansdowne Professor  
DLP:jcs

Cc: S. Wilson, panel members