Not Without Us*

Joseph Weizenbaum
Professor of Computer Science
Massachusetts Institute of Technology
Cambridge, MA.

Whenever I am in West Germany, I am amazed by the apparent normality of everyday life. As only an occasional visitor to Germany I see strange things that must by now appear routine, even natural to Germans. For example, holes in the streets that are intended to be filled with nuclear land lines or the closeness of every German citizen to nuclear weapons storage facilities. I notice, in other words, the Germans' physical, but even more their psychological proximity to the final catastrophe.

We in America are no more distant from the catastrophe than the Germans. In case of war, regardless of whether unintentionally initiated by technology allegedly designed to avert war, or by so called statesmen or women who thought it their duty to push the button, Germans may die ten minutes earlier than we in fortress America, but we shall all die.

We have no holes in our streets for atomic land mines. We see our missile silos only now and then, that is, only whenever it pleases someone to show them to us on television. No matter how passionately our government tries to convince us that the nasty Soviets are effectively as near to us as to the Europeans, that they threaten us from, for example Cuba or Nicaragua, Americans are, on the whole, unconvinced and therefore untroubled by such efforts. It would therefore be more astounding were the average American aware of the danger that confronts us all, than that he worries so little about it. The American experience of war allows a "it can't happen here" attitude to grow rather than a concrete fear of what appears to be far removed from the immediate concerns of daily life.

I am aware that it is emotionally impossible for people to live for very long in the face of immediate threats to their very existence without bringing to bear psychological mechanisms that serve to exclude those dangers from their consciousness. But when repression necessitates systematically misdirected efforts or excludes potentially life-saving behavior, then it is time to replace it by a deep look into the threat itself.

This time has come for computer professionals. We now have the power to alter the state of the world fundamentally and in a way conducive to life.

It is a prosaic truth that none of the weapon systems which today threaten murder on a genocidal scale, and whose design, manufacture and sale condemns countless people, especially children, to poverty and starvation, that none of these devices could be developed without the earnest, even enthusiastic, cooperation of computer professionals. It cannot go on without us! Without us the arms race, especially the qualitative arms race, could not advance another step.

Does this plain, simple and obvious fact say anything to us as computer professionals? I think so:

First those among us who, perhaps without being aware of it, exercise our talents in the service of

---

*An abbreviated English translation of a talk given to the Gesellschaft fur Informatik, at Karlsruhe, West Germany, on 17. July, 1986. Translated by the author.

death rather than that of life have little right to curse politicians, statesmen and women for not bringing us peace. Without our devoted help they could no longer endanger the peoples of our earth. All of us must therefore consider whether our daily work contributes to the insanity of further armament or to genuine possibilities for peace.

In this context, Artificial Intelligence (AI) comes especially to mind. Many of the technical tasks and problems in this subdiscipline of computer science stimulate the imagination and creativity of technically oriented workers particularly strongly. Goals like making a thinking being out of the computer, giving the computer the ability to understand spoken language, making it possible for the computer to see, goals like these offer nearly irresistible temptations to those among us who have not fully sublimated our playful sandbox fantasies or who mean to satisfy our delusions of omnipotence on the computer stage, i.e., in terms of computer systems. Such tasks are extraordinarily demanding and interesting. Robert Oppenheimer called them sweet. Besides, research projects in these areas are generously funded. The required monies usually come out of the coffers of the military - at least in America.

It is enormously tempting and, especially in Artificial Intelligence work, seductively simple to lose or hide oneself in details, in subproblems and their subproblems, and so on. The actual problems on which one works - and which are so generously supported - are disguised and transformed until their representations are mere fables, harmless, innocent, lovely fairy tales.

> An example: A doctoral student characterized his projected dissertation task as follows.
> A child, perhaps six or seven years old, sits in front of a computer display on which one can see a kitten and a bear - all this in full color of course. The kitten is playing with a ball. The child speaks to the computer system: "The bear should say 'thank you' when someone gives him something". The system responds in a synthetic but nevertheless pleasing voice: "Thank you, I understand." Then the child again: "Kitty, give your ball to your friend." Immediately we see the kitten on the computer display throw the ball to the bear. Then we hear the bear say: "Thank you my dear kitten."

This is the kernel of what the system, whose development is to constitute the student's doctoral work, is to accomplish. Seen from a technical point of view, the system is to understand spoken instructions - that alone is not simple - and translate them into a computer program which it is then to integrate seamlessly into its own computational structure. Not at all trivial, and beyond that, quite touching.

> Now a translation to reality:
> A fighter pilot is addressed by his pilot's associate system: "Sir, I see an enemy tank column below. Your orders please." The pilot: "When you see something like that, don't bother me, destroy the bastards and record the action. That's all." The system answers: "Yes sir!" and the plane's rockets fly earthward.

This pilot's associate system is one of three weapons which are expressly described, mainly as a problem for artificial intelligence, in the Strategic Computing Initiative, a new major research and development program of the American military. Over six hundred million dollars are to be spent on this program in the next four or five years.

It isn't my intention to assail or revile military systems. I intend this example from the actual

practice of academic artificial intelligence research in America to illustrate the euphemistic linguistic dissimulation whose effect it is to hinder thought and, ultimately, to still conscience.

I don't quite know whether it is especially computer science or its subdiscipline Artificial Intelligence that has such an enormous affection for euphemism. We speak so spectacularly and so readily of computer systems that understand, that see, decide, make judgments, and so on, without ourselves recognizing our own superficiality and immeasurable naivete with respect to these concepts. And, in the process of so speaking, we anesthetise our ability to evaluate the quality of our work and, what is more important, to identify and become conscious of its end use.

The student I mentioned above imagines his work to be about computer games for children, involving perhaps toy kittens, bears and balls. Its actual end use will likely mean that some day a young man, quite like the student himself and who has parents and possibly a girl friend, will be set afire by an exploding missile which was sent his way by a pilot's associate system shaped by the student's research. The psychological distance between the student's conception of his work and its actual implications is astronomic. It is precisely this enormous distance which makes it possible not to know and not to ask if one is doing sensible work or contributing to the greater efficiency of murderous devices.

One can't escape this state without asking, again and again: "What do I actually do? What is the final application and use of the products of my work?" and
ultimately, "am I content or ashamed to have contributed to this use?"

I am reminded in this context of a well-known American journalist who, during a Middle East highjacking, suggested that, under certain circumstances, the Israelies shoot ten Arab prisoners, selected from the many prisoners they were at the time holding, and, should the circumstances not change, shoot ten more the next day, and so on. He should not have made this suggestion unless he was prepared to go personally among the prisoners, to look with his own eyes into the eyes of the men to some of whom he will say, "you, you will die today," and then hold the pistol to the heads of those selected for murder and command his own finger to pull the trigger.

Just so should we, once we have abandoned the prettyfying of our language, begin to speak realistically and in earnest about our work as computer professionals. We should, for example, ask questions with respect to attempts to make it possible for computer systems to see. Progress in this domain will, with absolute certainty, be used to steer missiles like the Cruise and the Pershing ever more precisely to their targets. And at their targets, mass murder will be committed.

Such statements are often countered with the assertion that the computer is merely a tool. As such it can be used for good or for evil. In and of itself, it is value free. Furthermore, scientists and technicians cannot know how the products of their work will be applied, whether they will find a good or an evil use. Hence scientists and technicians cannot be held responsible for the final application of their work.

I see this argument concretely manifested in the building next to the one in which I work, the world-famous Draper Laboratory. This institution is devoted almost entirely to missile guidance and submarine navigation. [It was once, by the way, part of the Massachusetts Institute of Technology.] Many of the scientists employed there adopt the argument just stated as their own. They say that the systems on which they work can take men to the moon and bring them back just as these same systems can guarantee that missiles aimed at Moscow will actually hit Moscow when fired. They cannot know in

advance, they say, which of these two or still other goals their work will serve in the end. How then can they be held responsible for whatever consequences their work may entail? So it is, on the whole, with computer professionals. The doctoral student I mentioned, who wishes to be able to converse with his computer display, does in fact believe that future applications of his work will be exclusively in innocent applications like, for examples, childrens' games. Perhaps his research is not sponsored by the Pentagon's Strategic Computing Initiative, perhaps he never even heard of SCI. How then can he be assigned any responsibility for anti human use of which his results might be put?

Here we come to the essence of the matter: Today we know with virtual certainty that every scientific and technical result will, if at all possible, be put to use in military systems. The computer, together with the history of its development, is perhaps the key example. In these circumstances, scientific and technical workers cannot escape their responsibility to inquire about the end use of their work. They must then decide, once they know to what end it will be used, whether or not they would serve these ends with their own hands, that is, with the psychological distance between themselves and the final consequences of their work reduced to zero.

I think it important to say that I don't believe the military, in and of itself, to be an evil. Nor would I assert that the fact that a specific technology that has been adopted by the military is, on that ground alone, an evil. In the present state of the evolution of the sovereign nation-state, each state needs a military just as every city needs a fire department. (On the other hand, no one pleads for a fire station on every corner, and no one wishes for a city fire department that makes a side business out of committing prophylactic arson in the villages adjacent to the city.)

But we see our entire world, particularly its universities and science and engineering facilities, being increasingly and ever more profoundly militarized every day. "Little" wars burn in almost every part of the earth. [They serve in part to test the high-tech weapons of the "more advanced nations."] More than half of all the earth's scientists and engineers work more or less directly in military institutions or in institutions supported in the main by the military.

It is only our already deeply internalized habit of prettifying our language that permits us to speak in terms of weapons and weapons delivery systems at all, when we are, in fact, discussing atomic explosives and hydrogen bombs. Those aren't weapons! They are mass murder machines and mass murder machine delivery systems - and that is how we should speak of them, clearly, distinctly and without evasion. When one once recognizes that a nuclear mass murder machine is nothing other than an Instant Auschwitz, an instant extermination camp, an Auschwitz without railroads or Eichmans or Drs. Mengele - but an Auschwitz just the same - can one then work on systems that steer devices of this kind toward living cities? That is what I ask my colleagues. They must earnestly ask themselves such questions and deeply consider whatever responses they find in themselves. Their answers will finally manifest themselves in their actions - concretely in what they do every day.

Probably the most pandemic mental illness of our time is the almost universally held belief that the individual is powerless. This (self-fulfilling) delusion will surely be offered as a counter argument to my thesis. I demand, do I not, that a whole professions refuse to participate in the murderous insanity of our time. "That cannot be effective," I can already hear it said, "Yes, if actually no one worked on such things ... but that is plainly impossible. After all, if I don't do it, someone else will."

First, and on the most elementary level, I must say that the rule: "If I don't do it, someone else will" cannot serve as a basis of moral behavior. Every crime imaginable can be justified on its basis. For example: If I don't steal the sleeping drunk's money, someone else will.

But it is not at all trivial to ask after the meaning of effectiveness in the present context. Surely, effectiveness is not a binary matter, an either/or matter. To be sure, if what I say here were to induce a strike on the part of all scientists with respect to weapons work, that would have to be counted as effective. But there are many much more modest degrees of effectiveness toward which I aim.

I think it was George Orwell who once wrote "The highest duty of intellectuals in these times is to speak the simplest truths in the simplest possible words." For me that means first of all the duty to articulate the absurdity of our world in my actions, my writings and with my voice. I hope thereby to stir my students, my colleagues, everyone to whom I can speak directly. I hope thereby to encourage those who have already begun to think similarly, and to be encouraged by them, and possibly rouse all others I can reach out of their slumber. Courage like fear is catching! Even the most modest success in such attempts has also to count as effectiveness. Beyond that, in speaking as I do, I put what I here discuss on the public agenda and contribute to its legitimation. These are modest goals that can surely be reached.

But, finally, I want to address such larger goals as for example

* Ridding the world of nuclear mass murder devices and perhaps also of nuclear power generators.

* So reordering the world that it becomes impossible ever again to convince workers of one country that it is a necessity of life that they feel their families on the flesh and the blood and the tears of people of other countries. (That is, unfortunately, the fate of many workers today - and not only of those who earn their daily bread in armaments factories, but equally that of those of us whose daily work is to sharpen high-tech weapons.)

* So reordering the world that every human being has available to him or herself all material goods necessary for living in dignity. (I have often heard well-meaning people say that, if we apply technology, especially computer and communications technology wisely, we may reach this goal in perhaps fifty to a hundred years. But we can reach it sooner, and without waiting for technological advances. For the obstacle is not the absence of technology, it is the absence of political will!)

I once heard Elie Wiesel say: "We must believe the impossible is possible." I understand that in two different ways.

* First, had we been able to believe that "the land of the poets and the thinkers" could give birth to human extermination factories which could compete in efficiency with the automobile factories of Detroit, we might not have had to experience Bergen Belsen. The impossible horror proved possible and became reality.

* Second, it was "impossible" in the America of only 150 years ago to abolish the slavery of the black people. After all, the entire economy of America's South was built on cotton. Cotton could neither be planted nor harvested without the unpaid toil of thousands of human beings out of whose wretchedness the plantation master could squeeze his profit. Nevertheless,

> at first only a few far seeing men and women, dreamers all, in Massachusetts, later many more citizens, realists among them, came to believe the impossible was possible, that the slaves could be freed and slavery ended. And it became possible. And it became reality.

The impossible goals I mentioned here are possible, just as it is possible that we will destroy the human race. None of us can alone achieve the one nor prevent the other. But each of us must believe "it cannot be done without me."