Forecasting the Future or Shaping it?

Herbert A. Simon
Department of Psychology
Carnegie Mellon University

Complex Information Processing
Working Paper #550

Forecasting the Future or Shaping it?

Herbert A. Simon
19 October 2000*

In today's sessions, we have been hearing a number of well-informed views of how the world's future will be shaped by the burgeoning information technologies: a kind of menu of alternative possible worlds. These predictions and forecasts have generally been accompanied with some evaluative statements, about the good and bad consequences for humankind that these futures would bring. I propose to approach the same topic, but, echoing some of the remarks of David Gelernter, Alexander Singer, and Arthur Clarke, I will proceed in a slightly different way.

In an important sense, predicting the future is not really the task that faces us. After all, we, or at least the younger ones among us, are going to be a part of that future. Our task is not to *predict* the future; our task is to *design* a future for a sustainable and acceptable world, and then to devote our efforts to bringing that future about. We are not observers of the future; we are actors who, whether we wish to or not, by our actions and our very existence, will determine the future's shape.

Some of you may regard my proposed goal of a sustainable and acceptable world as unreachable. That is the counsel of despair, and serves no purpose. Others will think that my goal of mere acceptability is too pessimistic. If you think that, then the best way to show that I am wrong is to help the world find an attainable path to higher goals.

But what are the conditions of acceptability for the future we wish to design? First, we must find a way for living at peace with all of nature, not destroying the bases for the survival of all of us: we humans and all the other life with which we share this planet. Meeting this condition almost surely requires finding

---

* This paper was presented at the Carnegie Mellon University Earthware Symposium sponsored by the Department of Computer Science.

acceptable ways for limiting total demands on the planet's resources even as we find more efficient ways of meeting these demands. It also requires us to give up the false pride that views us, the human species, as separate from Nature, and to recognize that we are a part of Nature. We must come to see that our human worth lies not in any uniqueness or distinctness from the rest of Nature, but in our abilities, in common with the other creatures inhabiting our Earth, to contribute to the cooperative venture of mutual survival.

Second, we must find a way for sharing broadly and fairly the outputs of our productive efforts of all kinds, no matter how ample or limited these outputs are in a sustainable world. Without that fair sharing, our talk of supporting diversity is meaningless. Contemporary theories that oversimplify the distributive system to a set of competitive markets with their associated prices are a caricature of the actual complexities of distribution, basically ignoring the moral issues of fair allocation and sharing of the Earth's products and resources.

Third, we must, as a necessary condition for maintaining any standard of fairness, find some way of greatly mitigating, and if possible eliminating, the innumerable and passionate divisions of "we" from "they" that continue to make the human world a blood-stained collection of warring tribes, continually engaging in fluctuating patterns of mutual hostility and collective mayhem, of which traditional war is only one of the more obvious forms, perhaps not even the one most devastating or difficult to root out.

Probably most of us would agree that if we could even come halfway toward meeting these three conditions, we would regard our world as pretty well designed.

The Uses of the Knowledge that is Technology

We have been hearing a great deal today about the future of technology, especially electronic and communications technology, its promise and its dangers. Oddly enough, you have observed that my specification of a design for the future didn't mention technology at all, but went back to simple, age-old human desires that must have been felt by our earliest ancestors. As David Gelernter said, we want enough to eat,

warm relations with some other people, pleasant and intellectually challenging tasks and surroundings, freedom from pain, fear, and hate. Technology must be evaluated by its ability to help us or hinder us in pursuing these goals; not by the flashing lights it enables us to produce. Perhaps e-mail and network discussion groups are examples of recent technological events that come closest to impinging on these basic everyday values, closer even than supercomputers and the prospective circuits of molecular size.

What about the technology we have been discussing today? Does this technology give us the means to reach the kind of future we would like, or is it at the root of the problems we must solve? Surely the answer is: "Some of both." Today, I think that we understand that technology is not metal and glass and plastic molded to our human purposes; technology is knowledge, and knowledge provides the capability for doing new things, things we haven't been able to do before — good things and bad things.

We have long had two myths about technology. We have had the myth of Prometheus, who stole fire from the gods for the use of mankind. He is the hero of modern science, demonstrating the power of knowledge embedded in technology to improve the human condition.

Our second myth is the myth of Pandora, created by Zeus to punish humankind for the presumption of Prometheus. On opening the basket of mischiefs that Zeus gave her, Pandora brought all kinds of ills to us. Our modern struggle to defeat infectious viruses re-enacts the morality play of the combat between Prometheus' technology and Pandora's mischiefs. As we create new knowledge in the form of anti-viral technology, mutation and natural selection create new knowledge in the form of new viruses. Good knowledge here creates bad knowledge (at least as judged from the human vantage-point; the viruses might reverse the rating).

Technology, at base, is knowledge, and it is we human beings who decide to what extent it will be used to advance the kinds of human goals I have proposed, and to what extent it will be used to defeat these goals. The fundamental moral responsibility of scientists is not to predict whether the uses of

the technology they create will, on balance, be good or bad: that is an impossible task. The important moral duty, which is shared by scientist and non-scientist alike, is to devote effort to fostering the beneficial uses of knowledge, old and new alike, and opposing its harmful uses. Participating in designing a sustainable and acceptable world is one way of discharging that duty.

I see no reason to take a pessimistic view of what balance the struggle between Prometheus and Pandora will produce. That is for us to choose, not to forecast. My own assessment of the past is that the accumulation of knowledge and technology over the ages has brought, on balance, more good than harm to humankind, but only through our present efforts do we determine how favorable that balance will be in the future.

The kind of electronic knowledge that we are acquiring today is unique in its direct relevance to the problems that face us. Of course, it is not really "electronic" knowledge; it is knowledge about the creation, transmission, and processing of information, and thus gets at the very roots of our own humanity. It is knowledge about ourselves, as thinking individuals and as communicating collectivities, that has given us a whole new picture of what human thought processes are all about and how they are carried on. To be sure, the picture we have been forming has largely focussed on thought, to the relative neglect of emotion and motives, as Mr. Singer so dramatically showed us, but as we progress in our studies, that gap is also beginning to be closed.

What I would look for in this new technology, first and foremost, are better ways to understand ourselves, and by understanding ourselves, to find solutions to the problems, the human problems, I have proposed as the central ones facing us. For these are problems in which we are the central actors; problems about our ability to govern our own conduct in using the world's resources and conserving it as a life maintaining planet; and most of all, problems of governing our relations with each other.

The Knowledge that Computers are Bringing to Us

Because computers can process every kind of pattern and symbol, they provide us with a powerful new instrument and a powerful new methodology for modeling individual human minds as well as interacting collectivities of minds. They constitute a tool that, along with the new tools of neurobiology, provide us with the scientific means for understanding these complex systems called people and societies. They must play a central role in designing a plan for that attainable and sustainable future.

But of course the computer extends far beyond this single domain of application, however important and focal it may be. The computer is an invention that has a power and generality comparable to the power and generality of the pencil, or even of written language in general. If it is not the most important of all human inventions, it ranks among the top three, along with language and the skill of organizing. It has already become, and will continue to become increasingly, a constant companion and partner of the human mind.

For this reason, any picture of the future of our world is inseparable from a picture of the computer's future. As we search for the physical resources that the world of our desires will need, and the balance between resources and requirements that we must work out, the computer will play a central role in our thinking. With its assistance, we can aspire to handle the formidable complexities of understanding and designing the future, and we must not let our human vanity and misplaced desire for uniqueness limit the role we assign to the computer in this partnership.

While engaged in the partnership, we must also remain sensitive to the need to keep the computer's goals attuned with our own. I am not convinced that this will be difficult. The reading of history persuades me that the most dangerous villains we will encounter along the way will rarely be the forces of nature, and in particular, that they are more likely to have human than computer form. When we observe computers misbehaving, our cry should be, not "Cherchez l'ordinateur," but "Cherchez l'homme." Remember that the next time you are told that something can't happen because "the computer is

down." A human being, careless, lazy, stupid, or even malignant, rather than a computer is likely to be at the root of the trouble.

Conclusion

Here around CMU, we have been amazed, amused, gratified, and instructed by the developments in robot soccer. For four years, and with rapidly increasing skill, computers have been playing a human game requiring skillful coordination of all of the senses and motor capabilities of each player, as well as communication and coordination between players on each team, and strategic responses to the moves of the opposing team. We have seen in the soccer games, an entire social drama, played out with far less skill (thus far) than professional human soccer, but with all the important components of the latter clearly visible.

Here we see, in a single example, a complex web of all of the elements of intelligence and learning — interaction with the environment and social interaction, use of language — that artificial intelligence has been exploring for a half century, and a harbinger of its promise for continuing rapid development. Almost all of our hopes and concerns for the future can be examined in miniature in this setting, including our own role in relation to computers.

The robot soccer tournaments, like human soccer tournaments are a spectator sport. Although they have not yet erupted in riots of the kind that have infected the spectators at human soccer matches, they raise many questions about our relations with computers. For example, how far will and should we humans adopt the role of spectators in a world where much, perhaps all, of the world's necessary work can be done by computers?

Even before computers and the Web appeared on the scene, television and sports arenas, especially in the affluent societies of the first world, had caused a rapid development of spectatorism, which has been viewed, and I think correctly, with not a little concern. These developments remind us again that the myth of Pandora is as viable as the myth of Prometheus, and

that if the new knowledge is to bring us more good than harm, we must deal with the issue of how humankind will define its own purposes in the future world: purposes that, we would hope, would extend beyond the passive existence of the spectator.

When we look backward at the numerous historical examples of social elites that have lived more or less parasitically on their societies, what we see does not speak well for the ability of people to find useful and mutually supportive goals and challenges when these are not pressed on them by needs for survival. In the next round, we will have to do much better than such elites have done in the past. If we are to meet the three design criteria I have proposed for our future, or even come close to meeting them, we must find means to achieve yet another goal: to find activities for ourselves in this new world that will have the same wonder, and excitement, and intensity — going far beyond the crowd roars of spectatorism — as the most challenging activities that today's world provides us with: in the sports, crafts, professions, arts and sciences, and especially in the day-by-day subtle processes of living closely and warmly with other people.

Perhaps our very salvation will come from the severity of the problems we will have to solve: finding an ecologically sustainable state for the Earth and all its living inhabitants, injecting far stronger criteria of fairness into the allocation of available resources and their products, and disarming the vicious competitions that now take place between every imaginable sort of "we" and "they." If we accept this challenge of social design, there may be little spare time for excessive spectatorism. The soccer robots may have to design and build other robots to watch their games.