Computers and the Symbolization of Knowledge[1]

Zenon W. Pylyshyn

University of Western Ontario

London, Ontario, Canada

Introduction

Anyone who reads newspapers or magazines these days has heard the term “Artificial Intelligence”. It inspires awe in the business community and concern or puzzlement in most other places. Discussions of the limits of machine intelligence, though they still occur, have largely been replaced by concern over the social and economic significance of this new technology.

I’m one of those who is awed and impressed by the potential of this field and have devoted some part of my energy to persuading people that it is a positive force. I have done so largely on the grounds of its economic benefits and it potential for making the fruits of computer technology more generally available to the public — for example, to help the overworked physician; to search for oil and minerals and help manage our valuable resources; to explore, mine, and experiment in dangerous environments; to allow the non-computing public access to vast libraries of important information and even advice, and in the process give real meaning to the

1. Text of a talk delivered to the International Semiotics Institute, Toronto, June 9, 1987.

freedom of information act; and last but not least, to entertain people — for if we do not blow ourselves off this planet, entertainment in its most general sense, including *education*, will surely be the major industry of the future.

These are fruits of the new technology that many of you may be familiar with, at least in a general way. You may also be aware of the many pitfals that stand between us and their realization, including the military misuse of such work. What I would like to do today is concentrate on another aspect of the new technology — and to some extent of technology in general — which has received much less press. I want to concentrate on the *intellectual* significance of the recent developments in computer science, and particularly in Artificial Intelligence, and to try to make the point that the ideas of computing, information, and artificial intelligence bring with it changes in one's world view, and in particular to one's view of human nature. Since this may not strike some of you as obvious, I would like to take a moment to reflect on the role of technology in shaping a world view, since this role has not always been adequately recognized.

The process of understanding and of articulating one's tenuous grasp on an evolving picture of nature is not unlike the process of creating a work of art, such as a sculpture. In both cases there are three major elements: the imagination and curiosity of the creator, the nature of the tools available, and the resistance offered by the materials. A sharp knife on soft wood yields a very different result from a massive chisel on granite rock, quite independent of the sculptor's initial intention. Similarly, a brilliant man tackling the world's mysteries with only his eyes, ears, and the concepts with which his native tongue provides him, will have different experiences, and will carve them up differently than will a person with the tools of a technological culture — even though the world may be essentially the same in both cases.

Technology, by which I mean to include not only the design of artifacts, but everything that involves a codified system of methods or *technique*, provides both instrumental tools to help us make observations and calculations, and conceptual tools that help us to see things in new ways.

Such conceptual tools may be thought of as imagination prosthetics because they typically extend the range of the conceivable. Conceptual tools dominate periods of intellectual progress. Speaking of the importance of one such set of concepts — that provided by pure mathematics — physicist Freeman Dyson (1964) says,

> One factor that has remained constant through all the twists and turns of the history of physical science is the decisive importance of mathematical imagination. . . For a physicist, mathematics is not just a tool by means of which phenomena may be calculated; it is the main source of concepts and principles by means of which new theories can be created.

==============================

Insert SLIDE 1: DYSON QUOTE about here

==============================

Similarly, speaking of the development of new philosophical views, Susanne Langer (1962) put it this way:

> In every age, philosophical thinking exploits some dominant concepts and makes its greatest headway in solving problems conceived in terms of them. The seventeenth- and eighteenth-century philosophers construed knowledge . . . in terms of sense data and their association. Descartes' self examination gave classical psychology *the mind and its contents* as a starting point. . . . Hobbes provided the genetic method of building up complex ideas from simple ones…Pavlov built intellect out of conditioned reflexes and Loeb built life out of tropisms.

==============================

Insert SLIDE 2: LANGER QUOTE about here

==============================

I believe that history will record that around the mid twentieth century many classical problems of philosophy and psychology were transformed by a new notion of process: that of a *symbolic*, or *computational* process. Although the foundations for this idea was laid nearly a half century ago by mathematicians like Alan Turing, Alonzo Church, Kurt Godel, Steven Kleene, Emil Post, A.A. Markov and others, it was not until the late 1950s that the availability of digital computers made it possible to begin the transformation that we see today in both Artificial Intelligence and the closely related field of study called Cognitive Science (a term used widely in recent years to designate the scientific study of the processes that underly knowing, reasoning, imaging, planning, remembering, perceiving, and the like — the exact boundaries being under constant review).

Although the computer is a logical continuation of the path of technological development that has been going on ever since homo sapien began to make tools, yet in some ways it is also a radical departure from this path. In this talk I would like to explore the parallel between earlier technological trends and the development of computers. At the same time I will also examine what is *special* about computing. This discussion will lead me to consider what computing has in common with human thought that recommends it as a vehicle both for the understanding and for the excercise of intelligence. In doing this I shall consider three characteristics of this new technology: its capacity to increase the quantity and quality of complexity, its capacity to exhibit very nearly unbounded plasticity of behavior, and its capacity to determine actions on the basis of knowledge of the external world and certain goals related to this knowledge.

Computers and the Quality of Complexity

The industrial revolution made technology supreme by giving us mass production based on the two fundamental ideas of *division of labour* and *standardization (and hence* *interchangeability) of components*. These two are closely related. The products of technology are decomposable into component parts — parts that are independently specifiable in terms of the *function* that they perform, and which can be used in *any* sample of the manufactured product. Because of this it is possible to have such components designed and manufactured by specialists, and then independently assembled at some later date. This idea is important in part because it leads to high productivity. Yet perhaps even more important than the increase in productivity is the fact that the dual principles of division of labour and interchangeability of components make possible a certain kind of *accumulation of complexity* which had never been seen before the twentieth century. Before the widespread adoption of these principles the overall complexity of a project was limited to what could be conceived in the mind of the master designer (even the great pyramids and the wonderful cathedrals of Europe are no exception). Nowadays no single person has more than a highly sketchy understanding of the major products of technology.

The two elements, which together may be refered to as the *modularity* of technology, have had an effect on every phase of modern life. Indeed it might be argued that Guttenberg's invention of moveable type is an instance of the discovery of just this sort of modularity. In recent times a more self-consciously systematic and structured version of modularity has been developed to help organize large complex projects. Much of the dramatic successes of modern technology can be attributed to the exploitation of this idea. For example, there is very little in the way of new scientific principles or discoveries that went into the contemporary space program. The dramatic achievements in such things as landing a person on the moon is due entirely to the development of means for organizing complexity — for carrying the principles of modularity and division of labor to its extreme.

One of the fundamental characteristics of computer system design is that it makes the idea of the modular and hierarchical organization of complexity into a fundamental prescriptive principle, and capitalizes on it to create enormously complex systems — yet systems that are still understandable, and even repairable if they should fail to work properly. One way in which computing makes possible an extreme degree of modularity is worth examining because it sheds some light on the nature of the phenomenon of organization of complexity.

The one example I want to sketch relates to the idea of levels or layers of organization. Computer systems are typically designed by first implementing a set of general purpose facilities in hardware, then another independently designed set of facilities which might take the form of an interpreter for a programming language. Such facilities are specifically designed to supress certain logistical details from the person who is concerned with the design of a system for some particular application. So-called high-level programming languages, such as Basic, Fortran, LISP, and so on, make it possible for a system designer to concentrate on the concerns of the particular task at hand. Such an individual almost never knows in detail how the computer electronically carries out the sequence of commands that he specifies in his design. Similarly a designer can also use the computer language layer to create a new set of tools, such as a data base management system, which another designer can subsequently use to design still another complex system with no knowledge of how the facilities he is using are realized at the lower level. This deliberate *suppression of detail* is crucial to the orderly growth of complexity, because it allows different designers to each concentrate on a natural class of tasks that share common design principles.

This layering principle is fundamental to computer design and is carried to an extreme in such applications as the setting of standards for certain large-scale cooperative ventures, such as intercomputer communications[2]. The point of this layering is to provide a way to mediate between the diversity of specific machines and forms of information transmission on the one hand, and the need for identical standards on the other. This layering is just the idea of interchangeability of components and division of labour raised to a high art, in response to the

2. For example, the International Standards Organization has devised certain standards and codes to be used for communications among machines (the so-called Open Systems Interconnect standards). These standards consist of a set of 7 strictly independent layers — each of which can be realized in whatever way is appropriate for any particular computer or communication link.

enormous complexity of modern technological artifacts. It is the fact that technical objects can be organized in terms of this kind of conceptual layering that makes it possible for people to have sufficient understanding of such devices to operate them, repair them, and indeed even to design them.

This hierarchical style of design is both extremely natural for computing — for reasons having to do with the ease with which computers allow functions to be *composed* from subfunctions — and at the same time is the reason why computer systems can grow to such enourmous complexity. This, in turn is one of the factors that has contributed to the recent successes in artificial intelligence. Researchers now have the tools (i.e. some of the layers) with which to build systems whose complexity is sufficient to begin to produce behavior that we consider intelligent.

Those who remember the good old days of the early 1960's (how recent is the early history of artificial intelligence!) will recognise that programs such as the General Problem Solver were a very small fraction of the size of such current AI expert systems as Xcon or Mycin or Internist. This is critical because one of the lessons that has been learned in the past several dozen years is that intelligence arises from the interaction of a very large number of basic parts. In contrast with physics, cognition appears to have few general and powerful laws. Whatever else is required for the exercise of intelligence, it appears that one thing that is required is a large number of transactions on a large knowledge base. To a first approximation, to be intelligent is to know a lot and to be able to access the knowledge when it is appropriate. The commercial expert-system **Xcon** has a representation of some 2500 elementary facts, and **Internist** is said to have nearly a hundred times as much. And both of them are admittedly in the class of idiot-savants that can only perform intelligently over a narrow range of problems. Clearly in AI we are dealing with large numbers. Without some general scheme for keeping control of such quantities, the designers of such systems would soon lose their way.

Computing as Knowledge-Based

It is commonplace nowadays to accept computers as a major new medium for the storage, transmission and transformation of information. It is becoming clear to the public now, as it was not when I first began to teach computer science a number of years ago, that computers are not just lightening-fast calculators, but *information handlers*. Nonetheless it is not generally appreciated how the information handling capacity of a computer is different (except for their speed and efficiency), from the information handling capacity of such devices as books, files, tape recorders, telephones, libraries, television sets, and so on, all of which in some sense store, transmit, process, and transform information. But that it *is* different should be abundantly clear. Computers, unlike books or even television sets, *do* things because of the particular information they contain. Moreover, the nature of the behavior that they exhibit appears to be directly attributable to the *content* of the information they have — or what the information is *about*. This is not true of any of the other information handling systems we can think of — with the obvious exception of people or other higher organisms.

This observation is fundamental and bears some elaboration. Those who work with computers know that they can be coherently and consistently described as behaving in certain ways *because they know or have representations of certain things*. For example, a programmer might say that the computer predicted a rise in corporate spending at a certain time because it "has an econometric model" and was told the current values of certain parameters, or that it printed a check for so many dollars because it knows what the individual's salary is and what deductions have to be made according to the contract and the Income Tax law. The programmer might also correctly describe a computer as making a certain chess move by saying that it did so "in order to avoid having its knight captured by the opponent's queen." A more complex computer system, say one which makes medical diagnoses (e.g. the MYCIN system which diagnosis infectious diseases — to use a well-know older example), might be described as *inferring* that a patient has a certain infection and *recommending* a particular treatment based on its *knowledge* of symptoms and of the actions of certain drugs. Indeed most of these so-called 'expert systems' (e.g. the successor to MYCIN, called Teriesias) will even *explain* to the physician how it came to its decision by citing certain things it *believes* and certain inferential processes that it carried out.

Are these sorts of desriptions just convenient anthropomorphisms — like saying, for example, that the thermostat believes the room is too cold so it turns on the furnace? It's beginning to look very much like the answer has to be ***no***. For one thing the case of such artifacts as thermostats is easily excluded on the basis of a version of Occam's Razor (known to biologists as de Morgan's canon); *viz*, don't attrubute philogenetically higher capacitites when lower ones will suffice. In the case of devices such as thermostats, nothing whatever is lost in describing its functioning in physical terms. All its regularities are adequately captured in an electro-mechanical description, with no need to refer to goals, beliefs, knowledge, and so on. In such cases there is no genuine "level" of organization or of functioning beyond the electromechanical. But this is simply *not* the case for higher organisms, and it turns out not to be the case for the kinds of artifacts being designed in Artificial Intelligence. Such machines exhibit systematic patterns of behavior that could not adequately be captured by describing their physical structure. They must be described as being governed by internal symbolic representations: In other words, their behavior is dependent on the *knowledge* that is encoded in them. Since whenever I say this, some people look at me with suspicion, I had better take a few moments to at least hint at why this is so.

Representation-Governed Processes

The behavior of intelligent organisms is typically explainable only if we assume that their actions are governed by *decisions* based on *knowledge* that is used in the furtherance of *goals*. I will not take time to defend this view here, though I might point out that it is implicit in virtually all of psychology — even that part of it committed philosophically to the opposing view. For example, even doctrinare behaviorists assume that it is the meaning, or semantic content of their instructions to subjects, together with subject's preferences and beliefs about the experimental setup, that will determine what the subject will do in the experiment.

What I would like to do instead is give you some examples of the sorts of skills that people in Artificial Intelligence are attempting to duplicate in a computer, and to ask what it would take in order for a machine to be able to exhibit these skills. In these examples I will ask about design requirements, rather than about experimentally inferred facts about how people work. Yet the result is the same: asking what the task *demands*, or what are necessary conditions for carrying out the task, is part of the methodology of Cognitive Science — which is why Herb Simon has appropriately classed this new discipline as one of what he called the '*sciences of the artificial*'. It also shows that the technology of synthesizing systems and the science of understanding natural systems may not be as far apart as some people have believed.

The examples I want to sketch are miniature problems that occur in systems falling into the category of artificial intelligence. The first two are from systems that analyse aerial photographs.

==============================

Insert SLIDE 3: Oakland Docks about here

==============================

Suppose that it is looking for ships in aerial photos such as the following. How can it do this? There are two different ways it might attempt to solve this task, and these two ways illustrate a prime difference between AI technology and other engineering technologies. The first way might be to search for image properties that are frequently (or even invariably) associated with ships — for example, patterns of a certain size and shape and appearing at certain coordinates or wavelengths. I will call such reliable indicators "signatures" of what we are looking for. When signatures are available — as they are for a few important features such as ground moisture — then this is certainly a useful approach.

However, signatures of man-made entities are rare because such entities are more often characterized by their function than by their shape or other physical characteristics. After all, a ship can be of many different sizes and can appear as quite different patterns on an image, depending on the season, the lighting condition, the surrounding geography, and so on. But human photointerpreters do very well at the task of locating them. How do they do it? The answer appears to be that they analyse the images by making use of what they know about ships, harbours, water bodies, rivers, the shipping industry, and anything else that might possibly turn out to be relevant.

In the sample image shown here, for example, some of the ships are barely visible. Yet they *can* be detected if you know what to look for and where; if, for example, you keep in mind that ships tend to be located on water near docks, that some rectangular protrusions of the shoreline may be docks, that contours separate water from shoreline have charactersitic shapes that provide evidence as to which is water and which shore (e.g. at the mouth of rivers), that one can usually rely on such very general facts as that a contour does not change from being a water-land boundary at one place to a land-water boundary at another place as we follow it around. There are dozens and dozens of such obvious common-sense facts that go into the photointerpreter's skill, as well as some not-so-obvious things he may look for. Thus although an expert photointerpreter may describe what he does as "just looking", we know that such "looking" covers a large amount of knowledge-based reasoning — reasoning that the interpreter may be unaware of doing. A system that did photointerpretation as well as an expert — and over a range of situations as broad as the expert — is unlikely to be able to get away with searching for a “signature”. It would invariable have to resort to inferences based on knowledge that is represented in a form that allows it to reason.

Professional photointerpreters use not only their special expertise in high altitude photography when they interpret such images; they use all that they know or remember about the area being inspected. I am about to show you a photograph taken from the Landsat 2 satellite which contains an anomaly that a photointerpreter was able to detect. The anomaly in question is much like the anomalies that are common in Escher-type drawings such as the one in the next slide.

================================

Insert SLIDE 4: IMPOSSIBLE FIGURE about here

================================

The anomaly in the following photograph is not only more subtle, but can only be resolved by appealing to a much wide base of knowledge. The inconsistency in the photograph of the Expo-86 in Vancouver (which I obtained from Bob Woodham) is related to the direction of shadows. Note that all shadows point to the sun being …

================================

Insert SLIDE 5, 6: EXPO 86 SLIDES about here

================================

Here is another example which shows that carrying out a certain task requires that we appeal to knowledge. Suppose we wish to design a system that understands English sentences. Clearly, in order to understand an English sentence the system will have to uncover what is called its "thematic structure": it will have to determine "who did what to whom." This requires that a number of things about the sentence be determined, one of which is simply to decide what all the pronouns or other anaphoras refer to. But how are we to determine what the italicised pronoun refers to in sentences such as the following?

==============================

Insert Sentences about here

==============================

- The city councilors refused the demonstrators a permit because *they* feared violence.
- The city councilors refused the demonstrators a permit because *they* were communists.

There is no *general* fixed structure or mechanism for doing the right thing in all cases, in this case for assigning references. There is no way of wiring a machine once and for all so it gets all such examples right. Why not? Because it is patent that what the pronoun refers to depends on knowledge of such things as what city councillors are like, what the attitude of people in authority is to communists in certain countries, perhaps on recollections of recent history or the day's news, and so on without limit. Only factors like this would explain why the pronouns are typically assigned different referents in the two sentences and why the reference assignment can be easily changed by altering the context, and hence rendering some particular piece of knowledge relevant to the assignment. For example, I recently inadvertently provided an example of the effect of context myself when I used these sentences in a talk I gave in Florence, Italy. I had forgotten that the city council of Florence was in fact drawn primarily from the communist party. Because of this my audience had assigned the same referent to the pronoun in both sentences and the point of the example had been largely lost!

It is important to realize that all this talk about how certain regular features of behavior can only be understood in terms of the use of knowledge and the making of inferences and decisions based on goals is **not just a matter of expedience or convenience**. The fact is that there are certain systems in nature whose behavior is explainable only if we take into account the organization that they have at a certain level: namely, the level at which they can be described as having representations of knowledge and goals. No matter how accurately and completely you described their physical or neurological structure, you would still not understand why they displayed certain patterns of behavior — unless you also took into account this level of organization.

The existence of this level of organization is closely related to the behavioral plasticity of certain systems. One important reason why we have to postulate representations is in order to account for the rather radical plasticity in human behavior. Psychologists who conduct laboratory studies on human behavior must take extraordinary precautions in the way they instruct subjects and in what they might inadvertently lead subjects to believe about the experiment. The reason is that behavioral regularies are extremely sensitive to subjects' beliefs, which in turn are influenceable by the content of any information or clues that the experimenter might provide. Whatever regularities in behavior we might observe under some set of conditions can be altered in a systematic and logically coherent way (to a first approximation) by merely providing the subject with certain information — by telling or showing the subject something, or providing clues which together with other beliefs warrant some plausible inference (see Brewer, 1972, for excellent examples of this in human conditioning experiments). This responsiveness of behavior to the content of *information-bearing events*, by the way, is one of the main reasons for the downfall of behaviorism.

The radical alterability of behavior patterns has led some people to conclude either that a rigorous predictive science of psychology is not possible, in which case one should resign oneself to predicting statistical properties of behavior (i.e., what people will do *most* of the time), or they have concluded that we have been on the wrong track when we have been trying to model mental activity in certain mechanistic terms. The famous German Gestalt psychologist Wolfgang Kohler, was one of the people who took the second option.

Kohler viewed machines as too rigid to serve as models of biological or mental activity. The latter, he claimed, are governed by what he called "dynamic factors" — an example of which are self distributing field effects, such as the effects which cause magnetic fields to be redistributed when we introduce new pieces of metal. He contrasted such dynamic factors with what he called "topographical factors" which are structurally rigid. He says,

> "To the degree to which topographical conditions are rigidly given, and not to be changed by dynamic factors, their existence means the exclusion of certain forms of function, and the restriction of the processes to the possibilities compatible with those conditions… This extreme relation between dynamic factors and imposed topographical conditions is almost entirely realized in typical machines… we do not construct machines in which dynamic factors are the main determinants of the form of operation (Kohler, 1947, p65)".

==============================

Insert Kohler quote about here

==============================

Kohler clearly believed that "mechanical" systems are of necessity too contrained in their behaviors by their structure, or what he called "topographical factors." But in this Kohler was simply mistaken since, as the great mathematican Alan Turing was able to prove in the 1930s, certain kinds of systems wich do have a fixed topographical structure nontheless *are* able to exhibit arbitrarily plastic behavior; The way we would put it in today's terminology is that these systems are 'programmable' in a way that allows them to sumulate *any* formally describable system. Interestingly, order to prove the universality of the simple Turing machine, Turing had to introduce the idea of *symbols* or of reference, and hence of representation. In order to simulate another machine, the machine in question (called the Universal Turing Machine) has to have some representation of the machine being simulated. Thus what was missing from Kohler's analysis was the notion that a machine could have another level of organization — one in which it is described as having and using *symbols*, and hence *representations*.

Finally, I want to conclude by saying something about the notions of ***cognizing***, which I have been using loosely all through this talk.

Is Cognizing a natural class of phenomena?

The qualititive growth in complexity and speed of information handling that computers have made possible has been more generally recognized than has their character as knowledge-processors, or as Dan Dennett has called them, "semantic engines", which enables them to act on the *content of knowledge*. At the turn of the century, philosopher-psychologist Franz Brentano argued that the mark of the mental (and consequently the mark of the uniquely human) was the possesion of what he called "intentionality" — by which he referred to the fact that mental states are "about" something — they have what we would call "representational content". The problem of understanding intentionality remains one of the tougher problems in philosophy of mind. Yet many of us believe that whatever the eventual satisfactory analysis of this notion, it will include features of what computers do, as well as what the human mind does. Indeed, computers are the only non-living systems we know that appear to support a level of organization which corresponds to having and using representations, and hence that corresponds closely to what Brentano called intentionality. They do this dispite the fact that at the moment nobody wants to say that such machines are *conscious* — hence it appears that intentionality may be independent of consciousness.

There is an important point here, if this analysis is correct. Revolutionary changes in our view of the world frequently have had the character of discovering new groupings of phenomena — groupings that philosophers call *natural kinds*. The discovery that violent and natural motion were not different kinds of phenomena, subject to different principles, as Aristotle had taught, was one of the first steps in the development of the new physics. It went hand in hand with the reclassification of the motion of heavenly bodies into the same category as the motion of middle-sized objects like stones and cannon balls. The modern concept of "physical object", as anything that has mass and location — whether or not it is visible, or even detectable in principle, and whether or not it is in motion or at rest — seems totally natural for us today. Yet, this was not always the case. Indeed, Galileo was ridiculed for his assumption that things he could see only through a special instrument (the rudementary telescope) were of the *same natural kind* as things that could be seen with the naked eye. And perhaps we can understand why. Classing such things in the same natural category was to make an enormous conceptual leap: the leap of seeing a new fundamental grouping of things.

What people who work in the fields of artificial intelligence and cognitive science believe is that certain aspects of human capacity must also be regrouped or reconceptualized. Man has been variously understood as a creature of special creation, as a social entity and, in the late 19th century as a biological object. What some of us now believe is that there is another natural category to which cognitive or rational action should be assigned. That category is one which also includes certain sorts of machines as members in good standing: machines whose behavior is governed by what they represent — by what they know. These are knowledge-driven systems, or what George Miller picturesquely refers to as ***Informavores***, or systems that are nourished and guided by information. If this regrouping or reconceptualization is correct, it means that certain forms of human behavior should be explained in precisely the same way that we explain certain forms of computer behavior. Thus, contrary to a widely held view, the computer is not a metaphor for mind, any more than mathematical structures are metaphors for the physical world, as Freeman Dyson insisted in my earlier quotations, or geometry was a metaphor for space to Galileo. Rather, computing is a literal description of aspects of cognitive processes, stated in terms of a more manageable member of the same natural kind (viz, the natural kind *cognizer*).

This, then, is the new heresy: man the informavore, not only a cousin of the ape, but of the computer.

If all this turns out to be true, and the new natural kind becomes assimilated to the general view, as did the Galilean categories over the Aristotelan ones, we shall be witnessing a revolution in our image of mankind, perhaps greater even than the Darwinean or the Freudian. We shall also be witnessing a revolutionary change in the nature of our environment as we extend ourselves electronically. For the extension will be unlike that brought about by electonic communication media, which, as Marshall McCluhan once pointed out in these very halls, simply extended our senses. This new extension will literally place replicas of some of our most cherished functions — like thinking, deciding, recommending, evaluating, and pursuing goals — out there in our environment, along with other people and animals. As autonomous active gatherers and exploiters of knowledge they will represent a new and still incomprehensible form of externalized intellectual activity. Although in a sense they will still be tools, they will also be active participants in our intellectual activities, and we shall have to learn to live with them on those terms.