The Relevance of Manipulation to the Process of Perception

Edward Ihnatowicz

Published in The Institute of Mathematics and its Applications, May 1977, pp 133-135

The idea of making machines that think has an unfailing fascination, not only for science fiction readers, but for all who can see it is a possible way of gaining some understanding of the working of our own minds. Thinking, however, is not an easily defined phenomenon, although it is often considered to be the process of solving problems. We are accustomed to thinking of problem solving as essentially a mathematical process which may explain the popular view that the ultimate mechanical brain will be some mathematical machine, some universal problem-solving, theorem-proving and game-playing digital computer. It is pointed out that, after all, most of the great advances in our understanding of the world were made by people who learnt to identify and quantify various aspects of nature and by mathematically manipulating the resultant numbers have arrived at the universal laws which now form the basis of science. It is easy to overlook the fact that there are two distinct processes at work here. One is the sensory perception, or the means by which we observe life, and the other is the cerebral speculation on the results of our observations. It is our ability to perform the second of these, which includes all the intellectual operations of analysis, inference, association and generalisation which we particularly value in ourselves, and it is to these that the term "thinking" is generally restricted. If our goal is an understanding of the ways of our brains, however, it may prove that the more humble sensory perception is more important, and more revealing.

Many of our intellectual abilities are computational in nature and as such can be carried out, and usually much better, by machines whose structure or method of operation have no relevance to those of our brains. Consequently, even if we should succeed in writing, for instance, chess-playing programs capable of beating Grand Masters, these would not necessarily increase our understanding of thought. Perception, on the other hand, since it is by definition the function of the mechanism doing the perceiving, can be expected to be more directly instructive.

In perception itself, two distinct processes can be discerned. One is the gathering of the primary, sensory data or simple sensing of such things as light, moisture or pressure, and the other is the structuring of such data into information.

The relation between data and information is a functional one; information is only what data can be organised into or reduced to when related to the use to which it can be put by the system for whose benefit it is extracted.

Obviously, perception cannot be considered in the abstract, nor can the form that the resultant information may take inside the system. Both are dependent on the characteristics of that system and neither can be assessed without some knowledge of that system's purpose. In natural systems we divine such purpose intuitively as we intuitively assess any creature's intelligence by observing the efficacy with which it solves the problems of its survival. In artificial systems the concept of survival is somewhat arbitrary and the purpose seldom self-evident. To demonstrate the existence of perception under such circumstances, some objective test is necessary to show that the data are being organised in a way relevant to the task being performed. For such a test to be truly objective the task must be such that the response of the system to any stimulus can be evaluated on the basis of its benefit to the system without any reference to its internal workings. The most convincing response is one in which some physical motion causes a change in the relationship between the source of the stimulus and the system. Artificial systems which respond to stimuli by printing messages, however appropriate, are likely to be merely decoders or classifiers and therefore satisfy only one of the stipulated requirements of perception, that of sensing, leaving the structuring problem unresolved.

It may be argued, of course, that any arbitrarily chosen data structure is as good as any other as long as it is uniquely identifiable by the higher-level computing mechanisms but this is merely shifting the problem into another area. It implies the existence of some higher organisation capable of identifying some elements of a given set of data as corresponding to some aspects of the outside world, which in turn implies that the system is already aware of the existence of the world and can, therefore, attach some significance to such data, thus converting them into information. It would be much more interesting to discover a way by which any system could arrive at such an awareness spontaneously.

However difficult defining thinking, perception or intelligence may be, their intuitive recognition is easy and it is worth considering what qualities an artificial system would need to exhibit to appear intelligent to a casual observer. Consider a cigarette packet which, by means of some concealed apparatus, was able to move about the surface of a table in such a way that it successfully avoided being burnt by cigarette ends, having coffee spilt on it, falling off the table or being crushed by someone's elbow. Such behaviour would deserve to be described as sensible and, whatever mechanism controlled it, as intelligent.

I am not really suggesting that we should construct cybernetic cigarette boxes, merely that the situation described contains all the elements which would have to be taken into account in a design of a practical cognitive system. These can be described as:

1. The actual physical system (the box), whose construction is such that conditions hazardous to it can be identified;

2. A disordered and changing environment containing such hazards;

3. A motor ability capable of removing the system from the vicinity of danger;

4. Means of perceiving these aspects of the environment, capable of signalling the danger in advance,

5. A brain or a mechanism for deciding what to do on the basis of the perceived information.

The inclusion of complementary positive elements such as advantageous situations and the desire to seek them, would not affect the general scheme in any essential way. Such a scheme accentuates the importance of perception and the need to understand its mechanics.

It might help to consider the simplest, irreducible situation where perception might be reasonably expected to manifest itself. Such a situation might be represented by some finite volume of space filled with a uniform medium and containing two entities whose relationship is such that a physical contact between them is beneficial or harmful to one of them. In such circumstances contacts may occur as a result of random movement of either the medium or of one or both of the entities. Perception can be stipulated to exist if the occurrence of contacts is higher or lower than that resulting from random motions and such conditions being advantageous to one of the entities. Such a situation is clearly impossible without a provision of voluntary motion coupled with some mechanism for sensing the presence of the other entity over a distance. Furthermore, for such sensing to be possible the element to be detected must signal its presence by somehow affecting the medium in the area between itself and the other element as a function of that distance. In other words, whatever the method of such signalling, and that could be temperature, light, sound or chemical diffusion, its effect must be to create a gradient in the surrounding medium. The importance of such a gradient is in providing a criterion by which the system can monitor its rate of progress between the start of the motion and the physical contact. Such conditions occur frequently in nature and there are many organisms that employ various taxes or methods of migrating towards a favourable stimulus or away from an unfavourable one by climbing up such gradients. One of the more elegant ones is a bacterium called Escherichia coli. It propels itself with the aid of a bundle of flagella located at one end of its body. When these rotate in one direction the bacterium moves forward in a straight line, but when they move in the opposite direction the flagella splay out, causing the bacterium to tumble and lose direction. In the absence of any gradient in the stimulus to which the bacterium is sensitive, the backwards and forwards movements occur at random, resulting in Brownian motion, but when a gradient is detected, the ratio of occurrence of the two types of movement changes and the total motion of the bacterium becomes directional towards or away from the source of the stimulus. Such behaviour, being very mechanical, could be easily implemented in an artificial system and indeed many man-made missile-guidance systems or autopilots are a good deal more sophisticated. It is, nevertheless, interesting for, simple though it is, it contains an indication of an elementary perception mechanism. Under conditions where the gradient of the stimulus is stationary such as, say, a chemical diffusion around some source of food, the animal working its way towards it will expect the slope of such a gradient to be proportional to its muscular effort and to disappear when its motion stops. If then the level of the stimulus changes without the corresponding change in the movement, a situation which can occur only if the source of the stimulus is moving, some part of the control mechanism of the animal will be triggered off and the discrepancy between the stimulus and the muscular activity discovered. Such a discrepancy, which would be a physically measurable quantity, could perhaps be considered a very elementary, mechanical basis of perception.

It appears that not only is the physical motion of animals, when it is not random, controlled by some form of perception, but perception is equally dependent on some form of motion. There is some evidence that this applies to higher animals also. In a classical experiment Held and Hein have shown that vision in cats fails to develop if they are prevented from moving during a certain crucial phase of their early life. Also Piaget suggests that during the early development of a child manipulation forms a very important part of the development of visual perception.

It is true that we can both perceive and think while remaining perfectly motionless and that consequently a possibility exists that an artificial system might also do the same, but the fact remains that as yet there are no such systems, nor is it likely that they will appear in the near future. But then there is a great disparity between the number of computing elements in animal brains and in computers and there is no telling what artificial system might be able to achieve when their capacity begins to match that of the natural systems. Nor is it simply a question of numbers. Most of the 10 to the power of 10 neurons in a man's brain, for instance, are linked together into a active data-processing network both serially and in parallel, while in the bulk of our computers only a minute part of the computing elements are actively engaged in processing while the remainder are reserved for storage, with the data being retrieved and processed in the serial mode. A few experimental parallel processors have been developed whose method of operation has some correspondence to that of natural networks but their development is still in its infancy and their programming is proving very difficult.

The size of our computers is not in fact a very serious restriction. If the present rate of progress in the area of large scale integration is maintained, the possibility of construction of networks of elements comparable in size to natural ones is very real. The difficulty lies in specifying their organisation and their programming. No one would seriously contemplate specifying interconnections for a network of a hundred million elements. Equally no one is contemplating trying to unravel all of our own neural networks. And yet our brains have developed to their present state of complexity and efficiency from, in the last analysis, a disordered group of elementary particles by a method which must be eternal, universal and presumably still available.

It has been suggested that to emulate the evolution in terms of machines would inevitably take as long as it took nature to develop us, but this need not be the case. Nature, after all, had to operate under some very severe constraints which need not all appertain to machines. The most severe of them is that all changes on the genetic level, and these are the ones that control evolution, are random and that any experience gained by an organism during its life cannot be passed on to its descendents except through the process of natural selection. Thus, although nature has produced organisms that learn and reason, it itself does neither.

Learning can be thought of as a kind of evolution in which an organism does not have to die every time it makes a mistake, and if we can invent a type of evolution for our machines in which a system can pass on its experience to its progeny, that is, to tell us how to make it better, then the progress can be expected to be much faster.

It is worthwhile, therefore, to consider which of the identifiable elements of natural evolution need to be incorporated in an artificial system and how to construct an equivalent of trainable genetic coding. Only in such a way can we hope to arrive at the complexity of data processing which we now guess will be necessary before some of the capabilities of even the simplest natural organisms can be simulated.

Artificial genetics may appear an over-ambitious proposition, but it is essentially the task which the compilers of self-writing computer programs have set themselves, except that they try to do it entirely within the framework of their number-crunchers and have thus severely restricted their chances of success.

If, as has been intimated, mechanical interaction is a prerequisite of perception then a thorough understanding of mechanical information is inevitably a prerequisite of any success in this direction. Sadly, our appreciation of such information is very poor at present and if any attempt were made to compare or correlate artificially extractable mechanical data and, say, visual data about any aspect of a natural or artificial environment, the mechanical data would come a poor second.

Mechanical information is the most basic of all types of information, relating to such qualities as mass, force, inertia, consistency and friction. These are qualities shared by all objects living and inanimate, including ourselves, and in the detection and interpretation of which we are very efficient, but for the communication or expression of which we have no adequate method except, perhaps, poetry. We can, for instance, record or transmit by radio the sound that a bird makes and we can photograph or televise a bird's appearance but we cannot record or transmit what it feels like to hold a bird in one's hands, except by making verbal comparisons.

Yet it is information of this level of complexity which will have to be dealt with by artificial cognitive systems if they are ever to reach any reasonable autonomy in interpretation of the real world. It is impossible, for instance, to imagine how any system might "recognise" an image of an apple or of a table as being a representation of a solid object if it has no independent way of establishing that solid objects exist or, for that matter, that it is itself a solid object.

One of the characteristic properties of mechanical information is that it cannot be obtained passively. We cannot, in other words, determine whether a closed box is full or empty without lifting it, or a wing-nut tight or loose without twisting it, or, in general, any mechanical property of any object without disturbing it in some way. Quantum theory tells us that this is true of any form of measurement, but in mechanics it is particularly acute.

The implication of this for artificial intelligence is that the role of manipulating devices has to be extended from that of simple execution of commands to include the active interrogation of the environment. If we accept the somewhat restricted, but, for the present purpose, adequate definition of a brain as a mechanism for deciding what to do, and if what is to be done is some clearly definable mechanical action, then a basis can be established for evaluating both the usefulness of any sensory data and the efficiency of its structuring. Visual data, for instance, are embarrassingly comprehensive and a good deal of effort is required to extract the usually minute amount of directly usable information by a process which in most present-day systems, although often displaying great programming ingenuity, would not be practical in situations where the constituents of the image were not known a priori to consist of solid and usually geometrically regular objects. In such circumstances an introduction of mechanical sensing could help to eliminate all those constituents of an image not directly correlatable with the mechanical activity of the manipulator, such as those due to shadows, flat patterns or inaccessible elements.

In natural evolution, in fact, vision has occurred a good deal later than the ability to move and it would make much better sense for artificial systems to follow this example and to provide only such sensing abilities as can be utilised by the system's mechanics. Such a procedure would not only reduce the data reduction problems but might provide some insight into the higher-level abilities of generalisation and inference. If we construct our system so that mechanical and visual data about any object are extracted at the same time and have a comparable level of complexity, then there is a reasonable chance of a spontaneous discovery of their formal interdependence and of a functional basis for classification of external events leading, perhaps, even to a genuine concept of physical objects.

The over-all conclusion of such, admittedly, speculative considerations is that although our computers are not yet on the way to achieve an independent intellectual existence, some advance in our understanding of intellectual processes might be made with such machines, provided that means are discovered of somehow plugging them into the real world and imbuing them with a desire to survive as real objects in the world of real objects.