MAGGOTY INTELLIGENCE

Edward Ihnatowicz

[Unpublished. Date unknown: pre 1988]

Two small and rather unappealing animals possess patterns of behaviour that have great relevance for the student of intelligent systems. These are the wood louse and the maggot of the common housefly, and it is the difference in their behaviour which is so illuminating. It has to do with the way in which they orient themselves to their environment. Wood lice like moist places and succeed in aggregating there by the simple device of slowing down their otherwise random movements as the humidity increases. The maggot, which, during a certain stage of its development needs to come out of the dark, finds light by a slightly more sophisticated system. It has a single, non-directional light-sensing organ at the forward end of its body and as it moves along it swings this end left and right, allowing the amount of light gathered during each swing to determine the extent of its forward motion. In this way it keeps altering its course until the amounts of light sensed are equal for both sides, by which time it must be heading straight for the light.

The difference between these two behaviours is slight, but important. The wood louse is clearly less efficient because, unless there is a continuous gradient between the dry and moist areas, it will reach its objective only by accident, while the maggot can direct itself towards the light much more purposefully.

This difference characterises an evolutionary step of great interest to any study of intelligence. It is a difference between the simplest animals who respond only, to such immediate influences as temperature, brightness, pressure or salinity and those higher forms which are able to survey their neighbourhood and actively search for the most advantageous conditions. This second group may be said to possess an awareness of themselves as something other than their environment. Surely a desirable quality in any intelligent system.

Our two animals are interesting because they do not exactly fit into either group but are somewhere in between, and with the maggot having the edge on the wood-louse in an important aspect, any differences in their physical organisations should prove illuminating.

The most obvious difference is that in the wood-louse the location of the moisture-detectors is of no relevance to their effectiveness while the light-sensing cells of the maggot must be localised at the end of its body so as to make them in some degree directional. This directionality is clearly, in some way, significant so we need to consider what can be its advantages to the animal. It certainly enables the animal to alter the amount of light received by simply turning away or towards the light, rather than having to move to a different location, but that this is an advantage is not obvious unless we assume a corresponding directionality in its nervous system.

If we assume that the animal is able to relate the variation in light intensity to the direction of its movements, we may assume that it is also able to work out the direction of the light. Such an ability would indeed be an advantage but an ability to tell direction of motion does not come automatically with an ability to move.

The wood louse, for instance, manages quite well without it. In other words, in addition to an ability to receive the stimulus, the animal must also be able to move and to know at every instant by how much and in which direction it has moved.

The maggot is such a simple animal and direction so basic a concept that it is tempting to speculate whether physical motion may not be a necessary condition for any form of perception. The problem is that we do not know precisely what perception is.

The difficulty which we experience when trying to define such terms as perception, learning or information stem from the fact that all these concepts are relative while definitions need to be absolute. They are relative in the sense that they all specify relationships which exist only in a specific physical situation, that of some self-contained system operating in some environment. Most commonly it is ourselves who are the system and our world is the environment so that it is we who know about the world, communicate with it, seek information about it, and are or are not intelligent in the way in which we deal with it. Such relationships cannot be quantified except in relation to some specific person, animal or artificial cognitive system and their environment. It would be quite impossible, for instance, to estimate the amount of information contained in some picture or sentence without knowing what it referred to and who was to receive it. In this way the classical definition of information as a reduction of uncertainty is not helpful, for it makes no reference to what it is that the uncertainty is about, nor who it is that needs to be certain.

Even if information, perception, and generally, intelligence cannot be defined or quantified, their general characteristics can still be investigated and this is what Artificial Intelligence attempts to do. In such investigations, however, it would be helpful not to rely on ourselves as the only reference. Ideally we should like to have a simple, artificial system whose properties could be clearly defined and well understood and which would behave in relation to its environment in a sensible way and obviously, it would be helpful to know what are the minimum requirements of such a system. Our considerations of the maggot seem to suggest that the minimum requirements for a cognitive system are: one directional sensory input, one proprioceptive feedback and one motor output. If this should prove to be a general principle applying equally to natural and artificial systems it would have important implications for such areas of AI as scene analysis and pattern recognition where most of the work is done without any reference to motor functions.

So let us consider how profitable it is to speculate about perception learning or generally information processing in relation to a one-input system. Can the term "information" have any meaning in such a situation? In the first place, what is transmitted via any communication channel is not information but data. To consider data information implies that the data is on its way to some processing system which is in a position to interpret it by correlation with other data from different channels either arriving simultaneously or previously stored in some memory. In a single-input system such a possibility clearly does not exist.

Consider, for instance, the often quoted example of information being transmitted through a single channel, that of prisoners communicating by knocking on the wall. For this example to be valid we would have to consider the prisoners also isolated in every other way - no common language, wardens, food, or even view of the weather. In such circumstances, what could they be possibly communicating about?

Or again, consider the case of a congenitally blind man recovering his sight at a mature age: Our eyes are rightly considered the most efficient information gathering system but any visual impressions he may receive at that moment will be utterly meaningless and will remain so until he succeeds in establishing correspondences between the visual and other physical aspects of sense or objects known to him from his early life through other senses.

There is, of course, a way of looking at information in an abstract way in which we attach meanings to formal arrangements of elements within some set of data. This is what happens when we play games or do mathematics but mathematics is not information in the same way in which the colour green is not information. If we look through a telescope and all we can see is a uniform green colour it will not have any meaning for us unless we happen to know what the telescope is pointing at.

In the same way the number thirty-six is not information unless we know what it is a measure of. In general, the rules of mathematics like the rules of games or the laws of nature constitute information only to the extent to which they can be interpreted by a specific cognitive system.

Our stipulation that the motor output and feed-back are the necessary corollary of any intelligent system suggests that learning or perception cannot take place in the absence of active, physical interaction. This does not seem at all obvious but only because we tend to take our ability to perceive very much for granted and need people like Richard Gregory to tell us that an ability to see, for instance, is a skill we have to learn, albeit at a very early age. That we have to touch in order to see is something else again but that it is so has been proved in a number of cases, most dramatically in an experiment conducted by Held and Hein. Two kittens, brought up in darkness, were placed in two baskets in the apparatus shown in the drawing. The two baskets were mechanically linked; the kitten with the legs sticking out could reach the floor and thus rotate both baskets while the other one had to remain passive. The light was turned on giving them both a similar visual experience. After a spell in the baskets both kittens were tested to determine their ability to interpret visual patterns. Only the active kitten showed such an ability. The passive kitten remained effectively blind.

And yet cats are among the most highly developed animals. What clearly happens in a situation like this is that, although the passive kitten is aware of both the motion of his basket and the motion of the visual image which his eyes are receiving, it cannot discover that they are in any way related because it cannot control either of them.

That physical interaction is essential in any cognitive process seems likely when we consider that the only absolute reference any system has is its own physical frame. It is as if the animal had to use its own body as a measuring, or more precisely, conceptualising instrument; conceiving sizes and distances in relation to its own size, weights and forces in relation to its own muscular power and, most important, conceive time in terms of the rate at which its body could move. Such an arrangement might be thought of a king of perceptual filter and might explain how young children and animals manage to cope with the veritable flood of sensory data which must engulf them as soon as they are born. We know that very young children spend a lot of time trying to touch everything in sight as if to prove that as far as they are concerned the important thing about any object is not only what it looks like but what it does when touched, shaken, bitten or, for preference, dropped on the floor. They, perhaps, simply ignore all things with which they cannot physically interact. It is certainly difficult to imagine how any animal can form a concept of any phenomenon which it cannot in any way affect or be affected by.

It seems likely that the most promising way of investigating the problems of intelligence would be by constructing an artificial organism and observing its performance in a real environment and in real time. It sounds like a prescription for another robot but it is not. A robot is a machine which performs tasks, usually set by ourselves, while all we should ask of our machine is that it behave intelligently.

Intelligence is a subjective notion. It cannot be defined in scientific terms and, consequently, its existence cannot be proved in such terms. We all know, however, what is meant by intelligence and we may fairly assume that if we were ever to see an intelligent machine we would recognise it as such.

The difficulty with building such a machine is that in order to build it we need some of the answers which we are building it to provide; we would like it to be able to process information but do not know how this should be done; we do not even know what information is.

We can, however, describe in general terms the characteristics we would expect it to have. It should be a completely autonomous system capable of demonstrating some sort of logic in its behaviour, in its relation to its environment; it would have to exhibit some form of awareness of itself as something different from its environment and to demonstrate an ability to discover and appreciate properties of any new objects it might come across and form discernible attitudes towards them. It would be a machine built without any external purpose whatever and the only criteria both for its construction and its performance would be internal to itself.

Such a machine would be difficult to design because, in a sense it would have to build itself, it would have to grow. We would have to start by making the best possible guess as to its minimal requirements and get it quickly to the point where it could exhibit some form of behaviour and proceed be devising means of improving its performance. Our criteria would need to be in a large measure intuitive and we would need to be prepared to adjust them frequently in the light of gained experience.

What is proposed is a very open-ended type of approach alien to most scientific work but it is one which artists are accustomed to and it is my personal feeling that artists who, no one will by now be surprised to learn, are frequently involved in science and technology should be involved in research of this kind. The area to be explored is so vast and the guideposts so few that it would be very unwise to consider any discipline or approach as irrelevant or unsuitable.

Such an open-ended project need not be vague and a specific starting objective can be defined: to investigate the hypothesis that concept forming in any machine or organism can only occur in relation to the organism's physical structure and only through a dynamic interaction with its environment.

If such an open-endedness is accepted than a new, broad-fronted approach becomes possible in which the way of working, the journey becomes more important that the arrival, in which no version need be the final one, in which every refinement can be incorporated and every alternative experimented with. A method of working can be devised in which any number of people from any number of disciplines interested in the general problems of intelligence could work together in close communication, if not necessarily proximity, on a machine which would, with any luck, show some signs of life from the beginning, which would be always improving, which would exist it many versions and which would never be finished.

It sounds visionary and impractical but it isn't. All that is needed to bring it about is the realisation of the futility of back room, single discipline projects.

The practical difficulties are enormous. Probably the most important one is the fact that our present-day, serial, algorithmic, digital computer is quite unsuitable for this type of work and that the vital true parallel machine is in its infancy and largely neglected. It is for instance a fact that although the present-day technology is quite capable of producing a mechanical arm flexible enough and fast enough to theoretically snatch a flying ball from the air, there is at present no digital computer fast enough to control such an action, nor is its present line of development designed to improve it in this respect. And yet the task is a trivial one by animal standards and the explanation is that unlike the computers, animals process all sensory data simultaneously, in parallel. Parallel computers working on similar lines are being developed but are for some inexplicable reason largely ignored and badly supported.

Similar difficulties exist in relation to the techniques of simulating sensory inputs. The visual inputs for instance are dominated by the TV camera, which, although capable of supplying us with an amount of data far in excess of our present or foreseeable processing ability, does it again in the wrong way. It is virtually a machine designed to accept parallel data and convert it into serial. It does it so fast that for most purposes this is of no consequence; it does, however, make it quite unsuitable for ours.

The area where most useful work can and needs to be done is that of mechanical sensing and manipulation. To achieve the sort .of dynamic interaction which an AI project of any complexity requires, a mechanical manipulator needs to be built of much greater sophistication than may appear necessary at first glance. What is wanted is a mechanical arm which could not only manipulate things with speed and precision comparable with that of the human arm but at the same time informs us of all their mechanical characteristics, such as inertia, direction of motion, velocity, acceleration, friction, weight, compliance and frequency of any oscillations. It should not only be capable of stopping something from falling over but also of find out, for instance, whether a closed tin contains soup or cat-food, by shaking it.