

## Chapter 2

# Where to Start?

### Brains and Computers

It is certainly not obvious where to start in a description or study of “understanding” or “intelligence”. The terms are roughly synonymous, the main difference being that the second carries rather more suggestion of action. One feature that the brain has in common with a modern computer is a degree of versatility such that no short description can be given of the capabilities of either and no clear indication of a starting point for investigation or study. For the computer, one approach would be to go back to binary information storage and processing, and the use of stored programs, and to follow developments from there, roughly in historical order. Actually, these elementary aspects of computing are now submerged under such a mass of later developments that they would probably not now be the chosen starting points for a textbook or a taught course, but at least we have access to the ideas and intentions of the pioneers of the subject, and examination of these offers one possible starting point.

For the brain there is no corresponding access to ideas or intentions of a designer, but following Ashby (1960) it can be seen as a product of biological evolution and hence of selection according to survival value, and this gives some indication of where we might start.

### Discrete Logic

It is often assumed that discrete logic is the basis of effective thought, but it must have been a relatively late evolutionary development. The means of processing sensory data and using it to determine actions favouring survival is far from being peculiar to humans. The capability can even be ascribed to single-cell organisms, so is not dependent on a nervous system. Reactions of a particular single-cell organism are described by Alberts *et al.* (1989) and will be discussed in Chapter 4. The observation that microorganisms can show such behaviour gives indirect support to the suggestion, referred to in the past chapter, that microtubules may be involved in computations, possibly even when coexisting with nervous systems. This is of

course highly speculative, and it is impossible to know whether human intelligence is usefully considered as an evolutionary development from that of microorganisms.

Nevertheless, although the details are obscure, there can be little doubt that human intelligence evolved from something more primitive. The evolutionary viewpoint is important because new developments do not obliterate traces of previous forms, and traces of the primitive modes are likely to be effective “behind the scenes” in the intelligence of humans and other animals. There are many indications that human thinking is far from being fully represented by deductive logic, or in natural language, even if the outcome is often promptly rationalised and expressed linguistically. What can be expressed linguistically is just the tip of the iceberg of the thought process. The view is supported by no less a thinker than Albert Einstein in a letter to Hadamard quoted by Penrose (1989):

The words of the language, as they are written or spoken, do not seem to play any role in my mechanism of thought. The psychical entities which seem to serve as elements of thought are certain signs and more or less clear images which can be “voluntarily” reproduced and combined . . . The above mentioned elements are, in my case, of visual and some muscular type. Conventional words or other signs have to be sought for laboriously only in a second stage, when the mentioned associative play is sufficiently established and can be reproduced at will.

Some elusive aspects of the thought process have been acknowledged by the emphasis on heuristics in mainstream Artificial Intelligence. The limited success of AI suggests that there are further aspects still more elusive, probably also describable as heuristics but of a kind not yet explored, and it will be argued here that attention to continuous processing may be a key.

An additional reason for rejecting discrete logic as the basis of effective thought is that it is only useful once some concepts have been formed. Attention has been given to this under the heading of “binding” in neurophysiology, but something more is needed. Here and in the next chapter it will be argued that the manipulation of continuous signals is an important part of neural processing and should be regarded as more primitive than concept-based processing. It is useful to consider how the use of concepts could have evolved from continuous processing, as an evolutionary development that can throw light on processes operating “behind the scenes”.

### *Meaning of “Logic”*

The word “logic” is usually assumed to imply the manipulation of discrete concepts and as such is the basis of much work in AI, and it is definitely understood in this sense in computer technology. However, the word also denotes the science of thought, or thinking about thinking (with, as one logician put it: “no essential difference between the thinking that is thought about and the thinking that thinks about it”).

Much confusion has resulted from the common assumption that these two meanings of “logic” are equivalent. If they were, there would be no need to look for anything “behind the scenes” in the thought process. It is perhaps hardly surprising that humans tend to emphasize the discrete or concept-based aspects of thought,

as it seems to be their superior capacity for it (or at least for processing that finds expression in such terms) that distinguishes us from other animals and has allowed the development of science, technology, philosophy, and so on.

Formal logic, like mathematics, is often preferred to natural language because of its apparent precision, but this is illusory if it depends on distorting the representation of a problem or situation to make it fit the formalism. The preference is reminiscent of certain advice given to young men in Victorian times, where they were urged to remember that white is white and not “a kind of light colour”. Unfortunately for such a rule, colours described loosely as “white” vary considerably among themselves, as becomes clear when buying paint to match some existing decoration, and the less precise “kind of light colour” may be a more accurate description.

There is evidence, from Jerry Lettvin quoted by Anderson (1996), that Walter Pitts was unduly impressed by the possibilities of formal logic. It also seems clear from numerous papers in the earlier collection of the works of McCulloch (1965) that he attached much significance to the fact that the model neurons of the 1943 paper could perform the operations of formal logic. This is a somewhat tricky issue, as the operations of formal logic can form the basis of a general-purpose digital computer and can therefore underlie digital arithmetic and hence processing of continuous variables to any required accuracy. However, no evidence has been found of the nervous system using a binary (or decimal, etc.) system, and it is not useful to think about its processing of continuous data in such terms.

The last observation should be qualified by noting that it is possible to form a number system with radix one, and therefore no “carries”, such that magnitudes are represented by the length of the number. This can correspond to the number of pulses in a burst, or the number of active elements in a bundle of parallel binary channels, and this can be seen as a meeting point of analog and digital representations. The analog interpretation is in a sense more “natural”, however, and likely to be relevant to the evolution of intelligence. There is a connection with the emphasis on the “fabric” of computing by Beer (1959), which will be referred to again in the next chapter.

When McCulloch and Pitts (1943) produced the famous theory, they were well aware that their model neurons were unrealistic in many respects. One is that real neurons usually fire repetitively in a way that suggests the significant feature is the frequency and hence a continuous signal. It was argued, very reasonably, that it would be instructive to see what could be achieved with these highly simplified models of neurons, but reservations about the correspondence to real neural nets have often been ignored. Even the originators have arguably fallen somewhat short of the ideal in this respect.

There is much other evidence of an unwarranted prejudice in favour of discontinuous models of brain functioning, some of it reviewed in the next chapter.

## *Laws of Form*

The making of a distinction, for instance between inside and outside a boundary, has been suggested as the most fundamental component of intelligence and is the starting point for the treatment elaborated in George Spencer Brown’s *Laws of Form*

(Brown 1969). This is tantamount to advocacy of discrete logic so does not offer the most fundamental starting point. This is not to deny that it offers a valuable introduction to a large area of finite mathematics, but the value is to students who, along with the rest of humanity, have already progressed beyond the most elementary stage of development of intellect.

## Associative Recall

A number of approaches appear to assume that the basic problem is achievement of associative recall. People certainly have a remarkable capability to recall remembered items using a wide variety of associations. Postulated mechanisms allowing such recall include neural nets as postulated by Hopfield (1982), and others depending on holography (Gabor 1968; Willshaw 1981), as well as on a method mathematically equivalent to optical holography but not requiring mechanical rigidity (Willshaw *et al.* 1969).

These theories require a prior mechanism to select the items that are worth remembering, so even if valuable in themselves they do not provide a starting point. Another objection is that the kind of similarity that is effective in forming their associations is inherent in the method and not modified by learning. The theories therefore fail to account for the subtleties of human learning, which operates on more than one level by allowing different criteria of similarity to be learned. That different criteria are effective in different contexts was demonstrated by reference to a simple task of pattern recognition by Selfridge and Neisser (1963), who showed that spatial correlation as a measure of similarity could be misleading in the classification of hand-blocked characters, even where the images were adjusted to standard size and orientation. Introspection readily supports the idea that different similarity criteria are used in different tasks. For example, those that are effective in classifying leaves of trees by their shape are not the same as are effective in classifying hand-blocked characters.

At the same time, as Pribram (1971) points out, and as discussed by Andrew (1997, 2000), the similarity (or invariance) criteria effective in elementary recognition of objects, irrespective, within limits, of orientation and size, can be argued to correspond closely to the achievements of holography. Though not central to the current search for a starting point for consideration of brain activity in general, it is interesting to note that there are alternative approaches to the study of visual perception, one with reference to holography and the other to analysis of the image in terms of constituent features. Pribram refers to experimental demonstrations of spatial frequency selectivity that support the holographic version, but it is also clear that he does not suggest one single holographic reconstruction as an adequate model. It is possible that the multilevel application he visualizes is reconcilable with a feature-based scheme.

The remarkable human facility for associative recall is also apparent in essentially verbal contexts, for example, where the comment “she didn’t bat an eyelash”

evokes a similar phrase referring to “eyebrow” but with the understanding that the substitution implies supreme imperturbability. These aspects have a bearing on the principle of “heuristic connection” treated in the next chapter, but they do not provide the needed fundamental starting-point.

## *The Binding Problem*

In studies of neurophysiology, the way that sensory stimuli come to be associated, or bound together, as representations of recognised objects, has been termed the “binding problem” (Andrew 1995a). Experimental studies have failed to reveal special cells responding to discriminable categories, sometimes indicated by saying there is no “grandmother cell” (Gross 2002).

Experimental results bearing on the problem have been obtained by Singer (1990), who made simultaneous recordings from pairs of neural units whose outputs might be bound at a later stage of processing. The experiment was repeated with numerous variations, ranging from relatively simple cases where recording was from primitive cortical units as described by Hubel and Wiesel (1962), to others in which the paired units were in different brain hemispheres or responded to different sensory modalities.

It was found that, throughout this wide range of situations, the characteristic feature of paired discharges that were suitable for subsequent binding was a high degree of temporal coincidence of impulses in the two streams, indicated by a sharp peak at zero time in their cross-correlogram. The significance of this for underlying mechanisms is not easily judged.

Attempts to treat the binding problem theoretically have tended to focus on selection of patterns of activity that recur frequently. An early model of binding on this basis was made by Chapman (1959). A similar emphasis on recurrence appears in a comment by Einstein, quoted by Negoita (1993):

I believe that the first step in the setting of a “real external world” is the formation of the concept of bodily objects and of bodily objects of various kinds. Out of the multitude of our sense experiences we take, mentally and arbitrarily, certain repeatedly occurring complexes of sense impressions (partly in conjunction with sense impressions which are interpreted as signs for sense experiences of others), and we correlate them to a concept – the concept of the bodily object. Considered logically this concept is not identical with the totality of sense impressions referred to, but is a free creation of the human (or animal) mind.

In considering origins of intelligence, presumably governed by needs of survival, it is necessary to look for something other than frequency of occurrence of patterns or complexes that leads to their “binding”. At a later stage of evolution it is possible that patterns are chosen arbitrarily, as suggested by Einstein, and later retained only if they prove useful, but such a process (which could be described as operation of the heuristic of curiosity) cannot account for the emergence of binding and conceptual thinking in the first place. A view will be presented here that is plausible from an evolutionary viewpoint, consistent with Ashby’s insistence that the brain should be seen as a specialised organ of survival.

## *Models*

Craik (1942) viewed the nervous system as a calculating machine capable of modelling or paralleling external events. The idea of a model is not so simple as at first appears, because clearly a mental model is not a scaled-down version of something external, in the same way as, say, a model train is a scaled-down version of a physical entity. A model that can exist in memory, whether human or computer, has a less tangible relation to its prototype. The model and prototype correspond in only some of their characteristics, and the fact that the model is useful means that these characteristics are appropriate for some purpose.

Intelligent entities are usually viewed as working towards goals, or acting purposefully. Any such entity that is a product of evolution can be interpreted as pursuing the goal of survival and usually as pursuing other goals subsidiary to this. The way of interacting with the environment may be described as a set of rules, or a policy, and may or may not invite description as involving a model. The acceptance that a model is only justified in terms of a purpose suggests there is no firm distinction between a purposive system that forms a model and one that does not. In fact, any set of rules that favours achievement of a goal implies modelling of the environment because it amounts to the assertion that the environment is such that action according to these rules is likely to have a good outcome. There is correspondence here with the view of scientific theory expressed in the past chapter and with the general observation that access to absolute truth is not necessary in order to operate effectively and to believe that explanations are satisfactory.

Highly intelligent systems (of which we usually take ourselves examples) often form something that compellingly invites description as an internal model of a part of the environment. A person can sit still and imagine himself engaged in some activity, such as climbing a mountain or changing a wheel on his car or asking for a loan from his bank manager. If he is familiar with the environments in which these tasks are carried out, his mental images may contain quite a lot of detail. Even here, though, the model is related to purpose, and details that are not relevant to any purpose are usually absent from the image. It seems safe to assume that, in considering the early stages of evolution of intelligence, the focus should be on useful interaction with the environment rather than on modelling.

Of artificial systems for adaptive control, some are classified as forming a model and some not. The distinction is based on how a designer went about constructing the system and on how it is most readily understood. The overall performance of a model-based system may be indistinguishable from that of a system not explicitly embodying a model. This will be discussed in more detail in Chapter 4.

The observation that all adaptive control systems can be described as forming models helps to provide the “missing link” referred to in the title of this book. A servo system has more in common with the kind of cognitive system that is studied under the heading of AI than is generally thought. To achieve rapid and accurate response, a servo usually embodies more than simple negative feedback. A fairly usual expedient is to combine components denoted by the acronym PID, for

“proportional, integral, and derivative”, where straightforward negative feedback would have only the “proportional” component.

If all the properties of the servo and its environment (the load it will drive) are known, established theory allows the derivation of optimal gains for the three components (in the situation where everything is linear). Alternatively, the gains may be set by successive adjustment while the servo is in operation (possibly with additional components if there is serious nonlinearity). Humans and animals are remarkably good at making such adjustments, especially when they themselves constitute the servo. They do so in a wide range of manipulative tasks and learned skills, from walking to riding unicycles and flying jumbo jets. All of these involve modelling of the environment to a greater degree than is generally realized.

### *Conditioned Reflex*

The phenomenon of conditioning, famously investigated by Pavlov, has been the basis of numerous psychological studies, as reviewed for example by Dickinson (1987). Pavlov studied in particular the occurrence of salivation in dogs, which is an unconditioned response to the appearance of food and is also produced as a conditioned response by another signal such as the sound of a bell when this has been repeatedly paired with presentation of food. Here the unconditioned response of salivation clearly has survival value because it helps the dog’s digestion, and its elicitation by a different signal may be advantageous if the signal is indeed linked to food. It might for example allow salivation to begin sooner than if the animal waited for the materialisation of food as such.

This classic or Pavlovian form of conditioning agrees very well with the mechanism for synaptic modification proposed by Hebb (1949), expressed by him as follows:

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased.

Another type of conditioning experiment, particularly associated with B.F. Skinner and his invention of the “Skinner box”, is called operant conditioning. Here an animal is rewarded, typically by receiving a food pellet, when it performs a particular action such as pressing a lever. The first press is accidental as the animal explores its environment, but if it is hungry it repeats the action and receives more food. This offers a starting point for the study of intelligent behaviour favouring survival. An action of the animal is found to correlate with something favourable to survival and the response is to enhance the action. It seems reasonable to regard responsiveness to such correlation as a primary feature, perhaps the essential primary feature, of an intelligent entity.

The connection between “conditioned reflex” and statistical “conditional probability” was stressed by Albert Uttley (1956, 1959), who built a special-purpose computer to model the conditioned reflex by computing conditional probabilities



(Andrew 1959a). The conditional probability of a particular event, given that another has occurred, is an indication of the kind of linking referred to as “correlation”, and in this context it is a more appropriate numerical indicator than is the correlation coefficient.

Uttley’s conditional probability computer had a set of yes/no input channels, five in number as implemented, and arbitrarily labelled  $j$ ,  $k$ ,  $l$ ,  $m$ ,  $n$ . The device was “trained” by presentation of a succession of inputs consisting of nonempty subsets of the five channels (31 possibilities). It had 31 units that responded to respective subsets and counted the numbers of occurrences. The units counted all subsets, so that, for example, simultaneous inputs through the  $j$ ,  $k$ , and  $l$  channels would activate the  $jkl$  unit and also those responding to occurrences of  $j$ ,  $k$ , and  $l$  and of the pairs  $jk$ ,  $jl$ , and  $kl$ . The counts in all seven units would be incremented.

Once sufficient counts had been accumulated, the device was able to make inferences of activity, based on computation of conditional probability as a ratio of counts. The conditional probability of activity in the  $l$  channel, given that there is activity in the  $j$  and  $k$  channels, is given by:

$$p_{jk}(l) = (\text{count in } jkl\text{unit})/(\text{count in } jk\text{unit})$$

and it was arranged that, when such a ratio had exceeded a preset threshold value, an inference would be made of activity in an input channel, here the one labelled  $l$ .

In the device as implemented the counting was done in analog fashion, by incrementing the charge on a capacitor in each unit, and it was arranged that the charging law was roughly logarithmic, so that the division needed to estimate conditional probability could be replaced by subtraction. At the time the device was constructed, digital computers were rare and expensive, but apart from that it was believed that an analog solution was likely to have closer correspondence to events in neural circuitry. This argument was strengthened by the ease with which a “forgetting” feature could be introduced, simply by connecting resistors across the storage capacitors to allow their “counts” to drift towards zero. The inferences made by the computer could then be seen to be more strongly influenced by recent inputs than by earlier experience, conforming to human and animal behaviour.

The conditional probability computer could be made to simulate some forms of animal learning. It had no intrinsic appreciation of reward or punishment, so an indication of a favourable response had to be given. The device was sometimes demonstrated as modelling the training of a dog, where the five input channels were labelled so that two of them ( $b$  and  $s$ , say) represented the commands “beg” and “sit” and another two ( $B$  and  $S$ ) represented corresponding actions by the dog, and the fifth ( $r$ ) represented a reward in the form of a bone.

The dog was “trained” by presenting the inputs  $bBr$ ,  $bS$ ,  $sSr$ ,  $sB$  repeatedly in random order. It would then make the appropriate response (as an inference) to a command ( $B$  to command  $b$ , and  $S$  to command  $s$ ) provided it was “made hungry” by being “shown the bone”. That is to say, there would be an inference  $B$  in response to the input  $br$ , or  $S$  in response to  $sr$ , though no inference to inputs of  $b$  or  $s$  alone.

The need to “show the bone” does not invalidate the scheme as a model of biological learning because a tendency to seek particular inputs can be seen as a



consequence of the needs of survival. Nevertheless the demonstration with discrete input channels, in common with other models or accounts of classic and operant conditioning, does not provide the required starting point for the examination of intelligence. This is because the inputs depend on complex processes underlying pattern classification. The dog modelled in the above example has to be able to distinguish spoken words as well as the image of the bone, and in reports of classic conditioning it is assumed the subject can differentiate between the sound of a bell and for example that of a whistle. A familiar human example is the feeling of appetite that is aroused by quite specific sound patterns that customarily precede eating, such as the rattle of teacups.

However, although the phenomenon of conditioning is usually studied and described in ways that require the participation of complex processes of sensory analysis, appreciation of correlation is the basis, and could become effective at a more primitive level. This would be such that pattern classification is not required and the organism responds to correlation between continuous variables. This seems to be an appropriate starting point for the study of intelligence, with pattern classification and concept-based processing emerging as later evolutionary developments.

This viewpoint was reached after an attempt to apply the principles of the conditional probability computer in a practical control situation.

## **Application to Process Control**

The problem of reconciling continuous and conceptual information processing arises in attempts to make useful “learning machines” or self-improving devices to perform practical tasks. There is an implicit imprecision in any term (such as “self-improving” or “self-organising”) that implies that an entity modifies itself. Whether the entity is seen as self-improving, or usefully self-organising, must depend on how it is described, as there has to be a mode of description according to which the purported improvement was included from the start. There is a problem in identifying the “self” that organises. The modification, of course, is not expected to occur in isolation, but in interaction with an environment.

A useful viewpoint is due to Glushkov (1966) who describes a self-organising system as being decomposable, at least notionally, into a learning automaton and an operative automaton. The latter determines the behaviour of the system as tested and described (for example, the response of a person to the posing of a question). The learning automaton is the part of the system that modifies the operative automaton as experience is gained. The division between the two must be somewhat arbitrary and may depend on the length of time over which the behaviour attributed to the operative automaton is observed. An interesting point is that Glushkov considers the possibility of more than two levels of automata, so that there could be a learning-to-learn automaton that would operate to improve the operation of the learning automaton, and a learning-to-learn-to-learn one, and so on. (The term “learning-to-learn” appears in the psychology literature, and biologists are very ready to account for aspects of animal behaviour, or acceptance of pollination by plants, as favouring

appropriate levels of genetic diversity, and hence as evidence for “adaptation-to-adapt”).

Despite these complexities, there is fairly general agreement about what is meant by “learning” as a characteristic of people and animals, and as something that it is useful to imitate in artifacts. At a relatively simple level this can amount to “tuning” or automatic adjustment of numerical parameters, for instance of the gains of the three components in a PID servo mechanism. The adjustment would be made with respect to a suitable criterion of fidelity with which the servo output conforms to its input. Such adjustment can be considered a simple form of learning, although Pask (1962; see Andrew 2001a) insisted that the term should be reserved for processes that altered the representation or function more drastically.

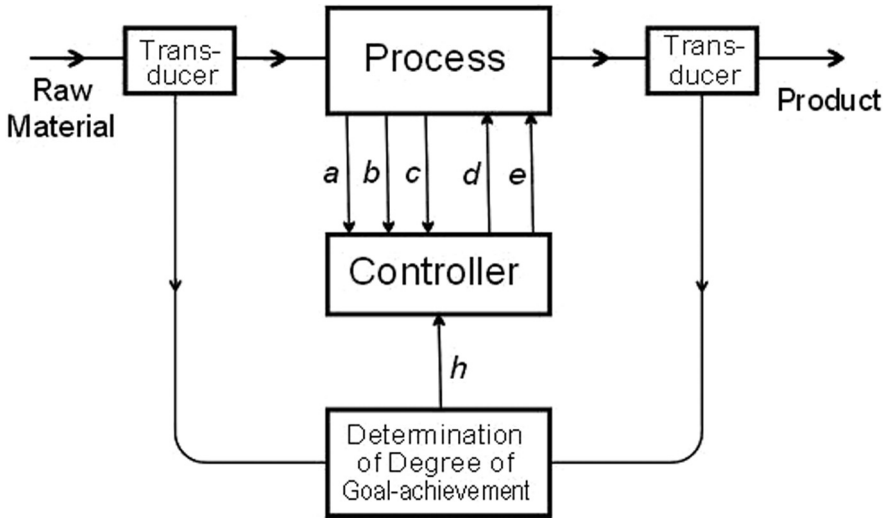
Pask’s insistence was that learning should necessarily involve what is discussed below as “classification”. The relatively simple process of “tuning” or parameter-adjustment was termed by him “adaptation”. However, it will be argued here (a) that “tuning” can be a much more complex process than at first may appear, and (b) that there is a route by which classification can arise as a development of “tuning”, and that this offers a highly plausible theory of the origins of biological intelligence. “Classification” can be roughly equated to concept-based processing, though in a primitive form where a concept need not be named in any language external to the learning organism.

Learning machines operating to adjust continuous variables have been widely assumed to be uninteresting as models of learning, and there have been several efforts to adapt schemes having discrete inputs and outputs. One such scheme is the conditional probability computer due to Uttley, already described as a means of modelling conditioning. The desire to introduce these discrete methods was prompted by the fact that human control is at least partly concept-based. It does not, however, follow that discrete operation should play a part from the start. According to the viewpoint that will be developed here, attention should first be given to continuous control, from which classification and concept-based processing can emerge.

An adaptive controller applied to a continuous industrial process is shown in Fig. 2.1. The controller receives signals  $a$ ,  $b$ , and  $c$  from transducers in the process plant and generates control signals  $d$  and  $e$ . Associated with the controller is a means of evaluating the degree of goal achievement, normally an estimate of profit for an industrial process though other considerations might also enter. This is labelled  $h$  for “hedony”. The controller is required to find a way of computing  $d$  and  $e$  from the inputs so as to maximise  $h$ .

## **Boxes**

One way of making the continuous signals acceptable to a controller accepting discrete inputs is to divide the range of possible values from each transducer into a number of intervals and to let all values within the interval be equivalent. This means that the multidimensional space (Three-dimensional if there are three transducers) in



**Fig. 2.1** Process with adaptive controller

which a point represents the state of the process (insofar as it is monitored by these transducers) is divided into rectangular “boxes” within each of which all points are equivalent as far as control action is concerned. There will usually be a large loss of information in going from the continuous to the discrete representation, but this may be acceptable.

Michie and Chambers (1967) developed a system they called BOXES to explore adaptive control. One of the applications they described was learning to play noughts-and-crosses (tic-tac-toe) by machine, though because this operates in a discrete environment, the special feature of division of a continuous representation space into “boxes” was not involved. They also considered a control task of a continuous nature, namely the balancing of an inverted pendulum, or pivoted pole, attached to a driven trolley. (This has also been used as a sample task for experiments in machine learning by Donaldson [1960, 1964], Widrow and Smith [1964], and Barto *et al.* [1983], but using quite different approaches.)

In the experiments of Michie and Chambers the pole, or inverted pendulum, was simulated in a computer program, but it is convenient to describe it as a physical mechanism. Like a clock pendulum, it was constrained to move in a plane. The pivot at its lower end attached it to a motor-driven trolley running in the same plane on a length of track. The task to be learned was that of controlling the motor of the trolley so as to keep the pendulum from falling over while also not allowing the trolley to run off either end of the track.

Signals made available to the controller, as though coming from attached transducers, indicated the position of the trolley along the track, and the angle of the pendulum, and the time derivatives of each of these. The range of the angle was

divided into six intervals, and the ranges of the other three variables into three intervals each, so that the total number of “boxes” was  $6 \times 3 \times 3 \times 3$ , or 162.

The motor driving the trolley operated in what is termed “bang-bang” fashion, being always fully energised in one or the other direction. Time was suitably quantised, and at each step a decision was made as to whether the motor would drive to the left or to the right in the next interval.

A pair of storage locations was associated with each of the 162 boxes. In one of them was stored a value for the mean time to failure after a decision to drive to the left when the point indicating the current state fell inside the box. The other stored a value for the mean time to failure after a decision to drive to the right. The values were “smoothed”, or weighted in favour of recent experience. In the early stage of learning the decisions were made randomly, but as experience was gained they were made to favour the choice that was associated with a longer mean time to failure. Failure occurred when either the pendulum fell over or the trolley ran off the track.

The task as described, as well as that of learning to play noughts-and-crosses, differs from the situation represented in Fig. 2.1 in having to wait for the end of an episode of operation before receiving a signal corresponding to  $h$ , or hedony. Once the system has learned to keep the pole from falling, and the trolley from leaving the track, there is no further change in control policy. It would of course be a simple matter to include in the hedony feedback an extra component allowing reinforcement of policies that minimise the deviations of the pole from its upright position.

As a general model of learning in a continuous environment, the BOXES scheme has obvious drawbacks. It depends on partitioning of the operational space in an arbitrary way, whose effectiveness only becomes apparent in retrospect. A possibility that could be considered would be to let the partitioning be modified adaptively as learning proceeds, such that adjacent boxes might be merged if the feedback signals to each were found to be similar, and boxes whose feedback signals were particularly variable might be split.

In any version, however, if the partitioning is too coarse, there must be a large loss of information in going from continuous values to the discrete form, and this almost certainly means that the quality of control is less than what other techniques would allow. On the other hand, if the divisions are too fine, the number of “boxes” is increased and can readily become unmanageable. In the limit the boxes could be made to correspond to the limits of discrimination of the transducers, so that there would be a separate box for every discriminable state of the system. The penalty for this is not only the increased cost of the necessary storage, and of complexity of the means of updating it, but more significantly in the time needed for learning. Learning a control policy would require that the point representing the state of the system entered each of the boxes a substantial number of times, so as to allow the collection of statistical data relevant to it.

The number of discriminable states of the system, equal to the product of the numbers for all the monitored variables, is such that the time needed for learning in this way is prohibitive in any practical situation. The conclusion has to be that continuous variables have to be treated as such, so that interpolation and

extrapolation are possible. These clearly play a part in human and animal learning, as can be seen for example by considering a person learning to ride a bicycle. It is unimaginable that the person would know what to do when the bicycle had tilted, say, three degrees from the vertical, and also five degrees, and yet have no idea what to do for a tilt of four degrees.

An obvious disadvantage of the BOXES method, if used to generate a continuous output, is the inevitable discontinuity in crossing a partition between boxes. A smoother output could be obtained by associating the collected statistics with an array of key points distributed throughout the state space, for instance with a key point at the centre of each “box”. Then the control output for any state of the system could be generated by interpolation from the key points nearest to its representation in the space. Such a method takes account of continuity and is classifiable under the heading of “learning filters” discussed in the next section.

Fuzzy set theory, as introduced by Zadeh (1965), recognises the need to combine continuous and concept-based forms of processing. It has allowed new and valuable approaches to adaptive control, as treated by Pedrycz (1989) and in a vast literature in journals and conference.

The fuzzy approach combines the discrete and continuous in a way that is opposite to the evolutionary viewpoint advocated here. The fuzzy approach accepts concept-based processing as a basis and extends it by allowing concepts to be “fuzzy”. An example is the concept of a “tall man”, which is imprecise because there is no exact threshold of height beyond which a man is recognised as being tall. Zadeh showed that such “linguistic variables” can be treated effectively in terms of membership of fuzzy sets. In this particular case, the degree of set membership is a function of the actual height, rising from zero at some height (say 1.7 metres), which would definitely not be considered tall, to unity for heights of, say, 1.9 metres and above.

Because fuzzy techniques assume concept-based processing, they do not seem to provide the required starting point for the examination of intelligence. Their practical value is indisputable but is realised in combination with human judgement. Their application to adaptive control is particularly valuable in nonlinear situations, but, as with the use of BOXES, they generally depend on manual setting of parameters. Examples are the upper and lower limits mentioned above for “tallness”. Fuzzy relationships are usually represented by concatenations of straight-line segments, the  $x$ -coordinates of whose junction points are somewhat arbitrary.

## *Learning Filters*

Adaptive controllers treating continuous variables as such have received a great deal of attention. An early example was by Draper and Li (1951) who applied such a controller to adjust the ignition timing and the proportions of the fuel/air mixture of an internal combustion engine so as to maximise the power developed. One general approach to adaptive control is referred to as Kalman filtering (Sorenson 1985; Visser and Molenaar 1988). Methods for optimisation of computer models

are reviewed by Schwefel (1981) and these are generally applicable also to real-time online optimisation. The topic is discussed further in Chapter 4.

An important point that will be developed in Chapter 4 is that the changes possible within an adaptive controller, as represented in Fig. 2.1, can include selection of terms as well as adjustment of numerical parameters. One possibility is to let the control signals  $d$  and  $e$  be computed as polynomial functions of the sensed variables  $a$ ,  $b$ , and  $c$ , such as:

$$d = K + La + Mb + Nc + \text{higher-order terms}$$

where the parameters  $K$ ,  $L$ ,  $M$ ,  $N$ , and others appearing in the higher-order terms are adjusted during operation so as to maximise  $h$ . The adjustments can be determined by fairly simple statistical estimates discussed in Chapter 4, though the situation is complicated by influences on  $h$  that are not connected with the control action, referred to as noise, and also the uncertain time delay between a control action and its effect on  $h$ .

The means of adjusting the polynomial can readily be extended to allow changes in the set of terms included in the polynomial. These can be seen as changes in the form of the polynomial, and of the connection pattern of a set of neuron-like elements continuously evaluating it as the control signal  $d$ . A controller working in this way can merit description as self-organising, which is understood to imply the possibility of fundamental changes in structure. The meaning to be attached to “fundamental” is not precise, as for example in a network the making of a new connection is functionally equivalent to turning a gain setting from zero to a nonzero value. Description as self-organising is probably warranted when the number of connections effective at any one time is small compared to the number of possibilities.

Certain measures of correlation can indicate the desirability of introducing new terms to the polynomial. (An alternative would be to introduce new terms randomly “on approval” and then to eliminate a term if its coefficient remains close to zero for a sufficient time.) The introduction of terms of degree higher than one allows for nonlinearity. The method can also allow the introduction of new variables, additional to the  $a$ ,  $b$ ,  $c$ , and so forth, initially present. This shows correspondence to the establishment of the conditioned reflex, where a new signal such as the sound of a bell or whistle comes to be associated with the unconditioned reflex linking the appearance of food with salivation.

The manipulation of a polynomial in this way has much in common with recommendations for the forming of regression models, for instance by Draper and Smith (1966). There is however the important difference that the task of regression analysis has the characteristics of “learning with a teacher”. That is to say, a “correct answer” is made available to the learning system, so that the magnitude and sign of error can be determined. Any system required to form a model or to make a prediction has this character. A system required to optimise its control with only an indication of hedony, or degree of goal achievement, has to operate differently, as will be discussed in Chapter 4. The difference is sometimes indicated by reference to optimisation with odd and even objective functions.

The use of a polynomial is not the only way of representing a continuous relationship between variables. The possible use of a grid of “key points” in the multi-dimensional space defined by the monitored variables has been mentioned. A kind of self-organisation is possible here, too, allowing modification of the key point locations. This could operate by eliminating a point from the grid if its stored values were close to what could be derived by extrapolation from neighbouring points, and using other criteria to indicate where the density of points should be increased. There are possibilities to be explored here, but use of a polynomial, as is also customary in regression analysis, seems more convenient. A particular disadvantage of the key-point method is that the introduction of new variables requires a change in dimensionality of the representation space. Artificial neural nets offer yet other methods for representing continuous relationships in ways that allow adaptation.

Possibilities for self-organising controllers operating on continuous variables were discussed by Andrew (1959b) as the outcome of an attempt to apply Uttley’s discrete-event conditional probability computer to process control. At the time the discrete-event and continuous approaches were seen as irreconcilable alternatives, but the view of conditioning as a response to observed correlation allows a unified view encompassing both.

A similar scheme was implemented in hardware by Gabor *et al.* (1961) and termed a “learning filter”, a term that has come to be applied generally to schemes of the kind. Strangely, although they referred to the device as “learning”, Gabor and his colleagues denied that it had any relation to human intelligence. This was in reply to a question from Gordon Pask following the spoken presentation of the paper (in discussion appended to the printed text). Instead, these authors stressed the classificatory aspect of human mental activity and claim that an attempt to model human intelligence should embody this from the start. They are of course right in claiming that classification is important, but not to the exclusion of other functions.

Although not usually recognised as such, the famous checker-playing program of Samuel (1963) incorporates what is really a continuous-variable “learning filter”. This is probably the most famous learning algorithm to emerge in Artificial Intelligence and although Samuel worked on checkers, or draughts, his methods have been carried over into programs to play other games including chess. That a learning filter finds application in these totally discrete task environments lends support to the suggestion that something of the sort participates in biological intelligence.

The making of a move by one player in chess typically requires a choice from about 30 legal possibilities. The number of possible courses the game might take in only a few moves is enormous, as each of the 30 moves could be met by any of about 30 responses from the opponent, and so on. A skilled chess player can look at the configuration of pieces on the board and can say something about the relative strengths of the two players’ positions. There is therefore interest in forming a “static evaluation function”, sometimes called the “scoring polynomial”, to indicate the relative strength of a player’s position.

Once a static evaluation function has been devised, the making of a move should be simple, as it is only necessary to choose which of the 30 alternatives leads to the most favourable outcome according to this criterion. However, the game of chess, or



checkers, is not so easily mastered, because no one has devised a static evaluation function that produces powerful play. The adjective “static” indicates that the function is evaluated without exploration of possible game continuations (i.e., without lookahead). The ideal method of play in which every possible course of the game is followed to completion is easily shown to be infeasible.

Static evaluation functions have been employed by chess players since well before the days of computer chess. Where a game has had to be abandoned without hope of resumption (which, in the romantic tradition of chess, would presumably be because an enemy was at the gates of the city, or one or both players was about to go to the gallows), such functions have been used to decide a winner.

It has been found that the quality of play is increased greatly if the static evaluation is applied after the tree of possible continuations has been explored to some depth. Once the evaluation has been made at each of the final nodes, or “leaves” of the tree, values can be fed back up through the tree, with the assumption that each player, at each stage, makes the choice most favourable to him. That is to say, if the exploration is made when player A is next to move, and the static evaluation function is computed so as to represent advantage to A, at each node of the tree that represents a nonfinal move by A the value is taken from the daughter node with highest value, whereas at nodes corresponding to a move by the opponent B the value is taken from the daughter node having lowest value. The feeding-back of values through the tree is referred to as minimaxing.

With a branching factor of about 30 at each node, it is obvious that the lookahead tree proliferates rapidly. Much ingenuity has been applied to reducing the complexity without compromising the effectiveness of the process, for instance by restricting attention to moves that are, by some criterion, plausible, or by recognising board configurations as turbulent or quiescent, with preference for the latter as final “leaves” at which static evaluation is applied. More valuable than either of these is the use of “alpha-beta technique” as a strategy for exploring the lookahead tree and performing the necessary minimax operation on it. The order of evaluation of nodes can be such that large parts of the tree can safely be ignored and the computation time correspondingly reduced. These methods can be referred to “as pruning the tree”. Plausible move selection and the like represent forward pruning, and alpha-beta technique is referred to as backward pruning.

The main point for the current purpose, however, is that the essentially discrete game situation is usefully represented by a numerical “scoring polynomial”, which, if of first degree, is simply a weighted sum. The variables in this are numerical estimates of features of the play position well known to experts, of which the most obvious is termed “piece advantage” or “material balance”. This is simply a comparison of the pieces still held by each player, the pieces being suitably valued. There is no point in attaching values to the kings for this purpose, as the fact that the game has not ended means that both kings are on board. The highest value is given to the queens, and then considerably lower ones to knights, bishops, and rooks, and the lowest of all to pawns.

Other variables in the scoring polynomial (at least, as set up by Samuel for checkers) include “mobility”, of which a rough indication can be formed by comparing the

number of legal moves open to each player, and “centre control”. Samuel’s primary interest was in machine learning, and he arranged for the coefficients in the polynomial (corresponding to  $K, L, M$  in the earlier case) to be adjusted, or “tuned” as experience was gained, and he also included means of modifying the selection of terms.

A difficulty is that there is no suitable feedback of the degree of goal achievement, or hedony, to guide the adjustments. In the BOXES approach to learning to play noughts-and-crosses (tic-tac-toe), all adjustments were deferred till the end of the game, and then counts associated with each of the possible play situations were updated according to whether the game was won, lost, or drawn. This is possible because for noughts-and-crosses the number of possible play situations is relatively small, but for checkers or chess the number is enormous.

For his checker-playing program, Samuel overcame the difficulty by exploiting the fact that the quality of play of a program improves with increase in the depth of lookahead. The greater the depth, the less the performance depends on the exact form of the static evaluation function (s.e.f.). This shows that the evaluation of a situation obtained by backing-up from a lookahead tree is more effective than that obtained by applying the s.e.f. directly. The learning algorithm devised by Samuel operates to reduce the difference between the backed-up evaluation and the evaluation formed directly. It therefore operates with an odd objective function because the difference between the two values is similar to an indication of error. One snag is that perfect agreement between the two values could be obtained by setting all of the coefficients to zero. To prevent the program finding this unwanted minimum, it was arranged that the piece-advantage term was always present with a fixed coefficient.

Samuel’s program was successful in improving its own playing ability in checkers, sometimes by playing repeatedly against another version of itself in the same computer. It is particularly interesting to consider the internal workings of game-playing programs in connection with an early debate between Ashby and Bellman after an address by Ashby (1964) in which he advocated a vigorous program of research on computer chess as a way of investigating intelligence in general. More will be said about this in the next chapter.

### *Classification versus Tuning*

The schemes that have been discussed under the headings of process control and learning filters all refer to learning devices dedicated to one particular task. On the other hand, a person or animal faces different tasks and environments at different times, and something that looks like a single task from one point of view may drastically alter in character. An example of the latter is the task of walking upright, which has one character when the feet make nonslipping contact with the ground, but another when slithering on black ice, and similar considerations apply to driving a vehicle. The essentially continuous “learning filter” or “tuning” process needs to be supplemented by classification of task environments, with distinct sets of “tuning” parameters, defining control policies, for the respective environments.

For people and higher animals, the task situation may change because of internal decisions to pursue different goals, but at a primitive level the changes are likely to be imposed from outside. A means of classifying environments, so as to bring an appropriate “learning filter” into play for each, has somehow been evolved. The classification process is presumably invoked in response to a failure of a biological “learning filter” to converge, especially when the failure follows a particular pattern. A pattern that would strongly suggest a multiplicity of environments, and hence a need for classification, would be one where for a time the learning filter appeared to be converging on a policy and then abruptly began to diverge from it.

Details of how such evolution might have happened can of course be only speculative, but some such mechanism has to be visualised as underlying the transition from continuous-variable processing to classification and hence concept-based processing.

### ***Nontrivial Machines***

MacKay (1959) introduced a version of information theory in which he distinguished *logons*, or units of structural information, from *metrons*, or units of metrical information. As will be argued in the next chapter, the sensory inputs to a biological learning system are metrical, and the system derives structure from them. A basis for a possible mechanism was indicated in the foregoing section.

There is also a connection with the early work of Pask (1962; see Andrew 2001a) in which he maintained that any process to be described as “learning” must introduce a new descriptive language. Anything not having this feature he termed “adaptation”. The reference to “language” is in a special sense discussed by Arbib (1970) and does not imply the syntactic structure of a spoken language. The intention is to focus on fundamental changes of representation such as that from continuous representation to classification.

At the much more advanced level of human behaviour, it is possible to see evidence of innovation prompted by confrontation, in a way that is reminiscent of the postulated emergence of classification from opposing influences on a learning filter. Von Foerster (von Foerster and Poerksen 2002) has described humans as nontrivial machines. One illustration of this (not his) is that a person, shown the sentence “This sentence is false”, does not go into an endless loop of decisions (“true, therefore false, therefore true, therefore false ...”) as a trivial reasoning machine might. Instead of being trapped in such a loop a person classifies the situation, either inventing a term to describe it or recalling the existing term “paradox”.

### **Conclusion**

Out of a variety of possible starting points for the study and analysis of intelligence, it has been argued that the most valuable is the examination of schemes operating adaptively on continuous variables. The importance of continuous operation will be

defended from additional viewpoints in the next chapter. Concept-based processing is also extremely important, to the extent that facility in employing it should probably be seen as the main feature that gives humans a claim to mental superiority over other creatures. However, it should be seen as evolving from continuous operation, and suggestions are made as to how this could have come about.

The importance of continuous processing is defended on evolutionary grounds. This does not mean that it has to be apparent in all manifestations of intelligence. People are able to operate effectively in many areas without conscious reference to continuity, using faculties that have been shaped by billions of years of evolution. On the other hand, evolutionary developments usually leave traces of earlier forms, and it would be surprising if there were not remnants of primitive continuous processing operating subconsciously.

My own feeling is that the key to further progress in understanding intelligence, and to breaking the current impasse in AI, is to be found by examining the interface between continuity and conceptual processing. To some extent this is just a “hunch”, but arguments will be developed later, especially the discussion of the fractal nature of intelligence in Chapter 7, that lend support.

The term “continuity” has been used without definition, relying on intuitive interpretation. The distinction between discrete and continuous representations is treated by Klir and Elias (2003). As they point out, any observation of a value, whether by a physical instrument or a living sense-organ, is associated with some unavoidable finite error. It follows that the resolution level is limited, and as they say: “This implies that data are always discrete regardless of philosophical beliefs or the state of technology”.

In dealing with observed data, continuity “in the small” can have no meaning except as part of a constructed theory. By this is meant continuity as demanded by differential calculus and defined by saying that a function  $f(x)$  is continuous at  $\alpha$  if its limit, as  $x \rightarrow \alpha$ , is  $f(\alpha)$ . For the discussion here, continuity has to be interpreted “in the large”, implying linear ordering and recognition of distance. These are the necessary conditions to allow interpolation, extrapolation, and smoothing, of which at least the first is important in the many forms of learning that constitute skill acquisition.

Klir and Elias also observe that the advent of digital computers has encouraged the use of discrete rather than continuous methods of analysis. This of course is a comment on present-day practice rather than on the origins of intelligence, and also it remains true that most if not all digital computing makes some use of a number system, and hence of linear ordering and recognition of distance.

## Summary of Chapter 2

A number of possible starting points are considered for the description and study of “understanding” or “intelligence”. These include the rather usual assumption that discrete logic is fundamental, which finds expression in discussions by

McCulloch and Pitts and in the Laws of Form of Spencer Brown. Another candidate is a capacity for associative recall, which can be implemented in artificial systems using either Hopfield neural nets or analogues of holography. Each of these embodies an assumption that useful concepts have already been formed so must be underpinned by something more basic.

Neuroscientists have discussed the “binding problem” in perception, usually with an assumption that frequently occurring patterns of excitation come to be recognised. The emphasis on frequency is not consistent with a view of evolution as guided by survival value, both in determining binding and in the evolution of sense organs. Modelling of the environment has also been suggested as a fundamental property of an intelligent system, but it is argued there is no clear distinction between a model and a policy for interaction.

A useful starting point is provided by the study of conditioning, particularly operant conditioning in which the response to be conditioned is rewarded, typically by a food pellet. As usually described, this does not offer a starting point because it depends on a pre existing capacity for discrimination of complex stimuli. However, the basic principle of response to correlation between something external to the organism and something beneficial can refer to continuous variables and can be the basis of a continuous adaptive controller of the “learning filter” type. This can embody means of term selection that can be seen as self-organisation, but in order to operate in different environments it must be supplemented with a means of classification of environments. Such classification constitutes the beginning of concept-based processing and consideration is given to how it could arise. There is a connection with early theories of Pask and with von Foerster’s view of people as nontrivial machines.



<http://www.springer.com/978-0-387-75163-4>

A Missing Link in Cybernetics

Logic and Continuity

Andrew, A.M.

2009, XI, 139 p. 4 illus., Hardcover

ISBN: 978-0-387-75163-4