ON DEFINING THE INTELLIGENCE OF BEHAVIOUR AND MACHINES

by

Dr. P. H. Millar

Keywords

Aims

Anthropomorphism

Development of A. I.

Intellectual task

Intelligence

Psychological analogy

Task-selection

Teleological machinery

Turing test

ABSTRACT

Despite its merits the Turing test does not analyse intelligence in a practical way. The "intelligence" of psychologists is little better but consideration of psychological and neurological ideal experiments suggests that intelligence is judged on the basis of the decision-making in the chain ' input-decide-output'. We cannot separate any specific behaviour syndrome as 'intelligence'-display.

Studying 'intellectual tasks' instead of 'intelligence' is little help since the criteria for identifying and assessing these are either anthropocentrically subjective or based on 'intelligence' itself.

Intelligence is best seen as measuring the performance of the highest task selection and control mechanism in a machine. We cannot however escape the need to specify aims.

While the Turing test for supposedly intelligent machines has served a definite purpose - for by setting an easily comprehended goal it has encouraged research in artificial intelligence - the test does have many drawbacks. The two most serious of these are that it offers no information as to what we are to regard as reasonable human-like reactions, and that it does not offer any clue to a step-by-step approach to the creation of an artificially intelligent machine.

Turing's proposal, it will be recalled, was that instead of vague talk about computers "thinking" or being "intelligent" we substitute discussion of the level of performance achieved by computers in one specific activity. The activity concerned is simply that of responding humanly to our interrogation. If a person is presented with two terminals, one linked to a computer and the other to a human respondent then the computer will have "passed" the test if that person is unable to distinguish which terminal has which linkage.

That the Turing test is not sufficiently closely defined is clear if we consider various devious or even "dishonest" ways of satisfying it. We might, for example, construct a machine which gave the human being who was interacting with it the impression that conversation would only be continued on subjects acceptable to the machine and which then steered the conversation in such a way that the initiative always remained with the quite trivially pre-programmed machine. We can all think of individuals who behave in just this way and we do not necessarily regard this behaviour as showing lack of intelligence; prima donna-like behaviour is likely to increase rather than diminish our conviction that we are dealing with a real human being. But do we want prima donna computers?

Many other tricks can be devised which will deceive a human observer into thinking that he is dealing with another human being, but insofar as these rely on quirks of human psychology we tend not to want to build them into A. I. machines. Clearly we have a more subtle definition of intelligence in mind and are not interested in such "trivial" ways of passing the test.

The second drawback to Turing's way of setting the desiderata is that it does not provide us with any guidance for discerning the constituent skills which go to make up intelligent behaviour, and the way in which they have to be combined. When faced with the task of constructing a machine we have to

begin somewhere and are bound to ask of any proposed definition of intelligence that it be analysable in such a way that we are presented with a succession of steps which would lead us towards our goal. In the case of the Turing test we are simply referred back to human intelligence and presented with the task of analysing that before we can even begin to plan the task of constructing the machine.

These criticisms are not intended to detract from the positive aspects of Turing's suggestion: among these are (i) that it avoids any reference to particular mechanical devices or systems theoretic entities such as feed-back loops, and (ii) that it lays an emphasis on the relationship between intelligence and human-like behaviour. It is this latter relationship which I now wish to explore in more depth.

To start the consideration of machine intelligence with a consideration of human intelligence is not to prejudge the issue in favour of the conclusion that machines cannot be intelligent. It is in fact no more objectionable than commencing an examination of how machines can be made to walk by trying to define exactly what we mean by 'walking' in the case of human beings and animals.

If the attribution of "intelligence" to machines is to become anything more than a metaphor we must either invent a new and independent meaning for the word "intelligence" or develop (in Carnap's sense "explicate") the existing term so that it loses its anthromorphic connotation and acquires criteria of applicability which can be applied equally well to machines and men. It is this latter approach that I now wish to pursue. Whether this approach can be successful or not is in my view an open question, and is indeed a pragmatic question: although certain terms which are at present applicable only to human subjects could in principle be explicated or adapted so as to apply equally easily to non-human subjects we should find that the consequences of such adaptation were unacceptable in several cases - e.g. "virtuous", "sinful", "vertebrate". We shall have to decide whether a definition of "intelligence" which is so adapted as to apply equally well to machines, does so much violence to our existing system of beliefs and necessary truths about intelligence that we would be better to abandon the search. But let us begin, and let us start by examining human intelligence itself.

We may distinguish two types of truth about human intelligence. First there are the insights which we as speakers of English enjoy about what necessarily "goes with" intelligence; whether we call these insights necessary truths, rules of logical grammar, attributes of the deep structure of sentences employing the word 'intelligence' or its cognates, or ultimately entrenched truths, these are the insights which lead us to say that someone's utterance betrays unfamiliarity with the notion, rather than showing a deviant theoretical approach.

Secondly, there are the empirical findings which we postulate to be related to what has been diagnosed as intelligence. It is the overriding tendency of the scientific approach to use such empirical findings and particularly the theoretical frameworks used to unify them, as a basis for adaptation of the preliminary definitions on which the first type of insight is based. This seems the best policy here, and I propose to review some (psychological) findings related to what has been informally described as intelligent behaviour.

Psychologists are loath to use the term 'intelligence' except in the very specific context of I. Q. testing, and even in this field it is customary to disclaim any intention of testing general intelligence. The limitations of psychological tests are well known and it is probably enough to recall that I. Q. does not seem a very good parameter for inventiveness, theoretical insight or creativeness. The question is however a vexed one if only because I. Q. does to some extent correlate with ability to express oneself - which may lead to more recognition of the abilities of the high I. Q. subjects - and because the social and personal factors determining actual performance may lead even high I. Q. subjects to patently 'unintelligent' behaviour.

The state of the art in I. Q. testing is therefore that although some progress has been made we still recognise a difference between what is tested and that which we would wholeheartedly call 'intelligence' while, farther, the tests reveal capacities which may or may not be exercised in

1 Ryle
2 Chomsky
3 Quine, Goodman

practice whereas the intelligent machine will, we hope, be so not just potentially but also in its action.

It may well be the difficulty of isolating the "intelligence-factor" in specific acts which has led to the reticence of psychologists on this subject. There is a general feeling that the parameter along which man differs most significantly from other primates may fairly be described as that of intelligence, but this parameter does not readily admit of identification in particular activities - as for instance the factor of walking upright does.

A striking example of the difficulty is automatism. In such states human beings will exhibit prolonged sequences of well-integrated and apparently goal-directed activity without subsequent awareness of their actions and with no special evidence of premeditation. An obvious problem is that of responsibility in cases where a crime is committed in such a state. But what of intelligence? There seems no reason to believe that an individual in this state could not perform any of the actions which we would regard as relevant to establishing intelligence (whether by I. Q. testing or according to the everyday sense). Yet the individual is to all intents and purposes not conscious, and it is difficult to justify the ascription of intelligence to the individual in these circumstances. One way round this case would be to conclude that consciousness was necessary if we were to ascribe intelligence to the individual <u>as such,</u> but that we might still talk of this as a case of "display of intelligence" without attributing the intelligence to the person concerned. This is to say that we should be able to judge the behaviour objectively and ignore the psychological state of the agent. One might even suggest that we have no direct access to the <u>psyche</u> of the agent in any case and <u>must</u> in practice judge intelligence in such a way. For A. I. purposes it is clearly of great advantage if we can eliminate the more metaphysical aspects of the everyday concept of intelligence, so let us pursue this approach.

We are now interested only in the <u>display</u> of intelligence and take the human being as the vehicle for such display simply because we are most used to judging human actions. We need not remain with the case of automatism but can rapidly progress to cases which can be experiment-ally manipulated with greater ease.

The experiments of W. G. Penfield and others who have followed him in the stimulation of various areas of the human cerebral cortex are clearly of relevance. As well as being able to identify the specific areas of the brain which are associated with particular types of sensory input in given individuals, and also identifying motor centres which at least mediate voluntary activity, these experiments have been able to isolate association areas. When the visual association areas are stimulated electrically the subject has the impression of seeing scenes of quite considerable complexity and integration: this is in contrast with the effects of stimulating the primary visual areas for in those cases the typical experience is that of illumination of one point in the visual field.

In an analogous fashion stimulation of the primary motor areas produces muscular contraction in the appropriate limbs or bodily parts; in the supplementary motor areas more complex behaviour can be evoked. * Although up to the present evoked behaviour has been of a stereotyped nature, there seems no reason to believe that given multiple stimulation and full freedom of movement specific 'remembered' actions could not be reproduced.

An obvious difficulty here is that action normally involves contact with external objects. If such external objects as are necessary do not exist then the action cannot be brought to fruition, and even if they do exist it is all but essential that the sensory processes be integrated into the evoked reaction. This is by no means implausible; in the case of normal activity it seems that the higher cortical centres are not burdened with the task of adjusting bodily motion to slight variations in the position or movement of objects. Instead a much more direct coupling is established between the visual input and the motor output and this coupling is biased to achieve the desired effect.

This brief sketch is somewhat speculative, but not outrageously so. What it is intended to show is that we can well imagine a situation in which a humanoid being was equipped with a multi-point stimulation matrix implanted in the brain, and that we could - whether on the basis of theoretical understanding or by dint of experiment - evoke any one of the large range of skills which that being had previously acquired, and could arrange the integration of the sensory input necessary for the effective exercise of those skills.

* e.g. Bowsher, D. Introduction to the Anatomy & Physiology of the Nervous System. 2nd Edition, pp. 133-4.

Now consider the case of such a humanoid being who is under our control. We are able to elicit any activity which he himself might have used to demonstrate his intelligence. If anything can pass the Turing test then this can, but at the same time ail the reactions may be dictated by a human controller. One naturally rejects this approach and does so for two good reasons;

1. the use of learned abilities whose neurally coded triggers are merely pulled by the stimulation, leaves too many questions unanswered. (Note that we cannot exclude the possibility that goal-directed learning is one of the abilities which we can elicit; hence even on tests of learning the Turing condition might be satisfied)

2. although we may well accept that the outward behaviour did display intelligence we cannot accept that the system to which that intelligence should be attributed is bounded by the body of the humanoid agent.

The second objection is the crucial one. It is important to contrast induced "intelligent" behaviour with spontaneous intelligent behaviour, and by means of this imaginary experiment we are able to do so. In the case of the induced behaviour we are uneasy not about the physical presence of the electrodes, but about the fact of external control. The same problems would arise even if the control were exercised by hypnotism or by telepathy.

Further it is not the fact that the behaviour is controlled, that troubles us. The spontaneous intelligent behaviour is identical in form and may be postulated to be neurally identical at all points of the efferent chains below the points of the neo-cortex which are stimulated. The principal difference arises because the determination of the response is in one case mediated by the sensory input and cortical functions of the agent himself and in the other by the sensory input and cortical functions of (and transmission lines from) a separate individual. If we treat the controlling and the controlled individuals as a single system then it becomes plausible (though not entirely unproble-matical) to talk again of genuine exercise of intelligence.

* We may compare this puzzle to that posed by Puccetti, Analysis 29.3, January 1969.

It is particularly instructive to note that resulting two-person system acts on the basis of the perception and decisions of one individual but by means of the actual movements of the other. Since the controlled individual is not taken to display intelligence we might reasonably conclude that the seat of the intelligence in the system lies in the controlling individual and his functions.

We can however subdivide these into the perceptual, the decision-making and the executive functions. Incapacitation of any one of these will prevent the appropriate behaviour being manifested, but we shall almost certainly not wish to say that incapacitation of the perceptual or the executive controlling function is as such a removal of the potentiality for intelligent decision-making.

This point is obvious when we consider individuals suffering from various sensory or motor disabilities. In practice we make allowances for the disabilities before trying to judge intelligence. Even in cases where the disability is so extremely severe that we do not know how to make the necessary allowances, we still feel they should be made. It is exactly this fact which makes the Turing test plausible: after all the reactions of the machine which is being tested will automatically be judged in the light of the supposition that it is a human being whose perceptual input has been restricted to the one channel provided for the test communications. Similarly we shall not regard it as a sign of lack of intelligence in a Turing test situation that the respondent cannot display behaviour except by means of replies transmitted over a communication channel of more or less arbitrary characteristics.

These considerations support those already adduced that although the only way of detecting intelligence is the examination of behaviour, there is no particular behaviour pattern which is sufficient to constitute the display of intelligence. We can also add that although an intelligent system is so by virtue of its reactions to input stimuli there are no particular input channels or inputs which are necessarily implicated.

This leaves us with the decision-making as being crucial to intelligence, and such a conclusion would hardly be a controversial one. Briefly stated, one might argue that only things done as a consequence of decision ought to be taken into account in the assessment of intelligence and that it was the quality of the decision and not the quality of its execution or the accuracy of the suppositions leading to it which ought to determine our

conclusions.

Let us for a moment turn to the practical matters of concern to those constructing A. 1. devices. Dr. Meltzer has recently proposed * that the most constructive approach is to abandon strict adherence to the Turing test and to concentrate on perfecting our algorithms and heuristics for "intellectual tasks". He clearly feels that we can easily overdo our efforts at slavish imitation of human levels of performance in skills which only human beings have much interest in acquiring (we may yet read of robots built with special attachments for riding bicycles). Dr. Meltzer would be only too happy if our computer-based "intelligence" were to outstrip the abilities of their designers - e. g. by proving theorems for the statement of which more than one million words are necessary.

One can hardly quarrel about the value of this approach. But the foregoing remarks do suggest that if our pursuit is genuinely artificial "intelligence" then this approach may be beside the point. It is therefore of interest to examine the link which Dr. Meltzer postulates between the skills and abilities which interest him and the human faculty which wc refer to as intelligence.

Dr. Meltzer suggests that we shall remain true to the spirit of A. 1. if we concentrate on the solution of "intellectual tasks", and believes that one can intuit relatively easily which tasks these are. Although not in agreement with this, 1 do recognise it as an important suggestion. First, because it brings out a second aspect of the connotations of the term "intelligence" - one which is not adequately dealt with in the foregoing discussion. Secondly, because it holds out the hope that we may be able to proceed in the systematic step-by-step fashion whose importance 1 emphasised at the beginning.

1 would not question the possibility of making a rough division of human tasks into physical and intellectual. This is in my view only a rough distinction and the successful completion of almost any task will involve both physical and intellectual components. However J take the point that mowing grass is physical and proving theorems is intellectual while writing a sonata is something else again.

Informally we feel a connection between "intellectual tasks" and "Intelligence". We might suggest that in order to measure intelligence we measure performance in intellectual tasks, and I think that most would agree that this is roughly in accord with our present use of words. We may say, therefore, that a task is intellectual insofar as the use of intelligence is necessary for its successful completion.

Something like this, at any rate, seems to be the basis for our feeling that intelligence is connected with the ability to perform intellectual tasks. But if we are to argue for substituting a test framed in terms of such tasks for one based on intelligence per se we must have a means of judging both the identity of these tasks and levels of performance in them, independently oi judgements about intelligence. This is the first point at which my scepticism comes to a head.

I am prepared to accept that this identification usually does not directly implicate assessments of intelligence but in this form it does seem to me to be based upon an introspection as to the sort of "trying" relevant to success, Just as we may identify physical tasks as those to which muscular tone and exertion are relevant, so we are tamiliar with the fact that our performance at theorem-proving or chess is related to our "trving" in a particular way. As this sort of trying also helps us react intelligently in perplexing situations, the link with intelligence is natural.

The alternative approach to "intellectual tasks" which consists of an attempt to specify the important task parameters, ends up in the same confusion as does the attempt to define intelligence solely in terms of the behaviour we take to exhibit it. Intellectual tasks are only shown to have been accomplished by the occurrence oi output the production oi which is external to the task itself - or by a process of intuition such as that based on the sort of "trying" involved.

The main objection to relying on intuition in this way is that we cannot adequately put ourselves in the place of the machine we are testing. Theorem-proving is an example where this does not seem necessary, but the difficulty of judging whether word associations and synonym tests are being done by means of intellectual activity rather than by merely consulting a table of synonyms, are obvious.

Dr. Meltzer might well answer that it doesn't matter how it is done - all that is important is that the result is the same as that oi the relevant

* A1SB Bulletin No. 12, 1071 (British Computer Society A.I. Group)

intellectual task as carried out by a perfect performer. Yet this we can only apply in cases we know what the standards of perfection are. Finding synonyms is taken as an intellectual task because of the sort of trying it involves us in; if we have a list of synonyms we can make a machine which performs this task and ipso facto performs an intelligent task. Playing chess is a similar case, except that there are no right answers - only playing well or badly, here again there are more or less objective externally visible parameters of success. But take other activities which we humans find hard such as thinking about n-dimensional spaces. Since standard L. P. programmes can easily solve multi-dimensional sets of inequalities we have good evidence that computers can handle this sort of situation. Yet most of us would doubt whether loading an L. P. programme endows a computer with an ability for intellectual tasks which we humans find difficult; the capacities of computers are so familiar that we have stopped thinking of this sort of thing as a difficult task for them.

Two sorts of test case arise. The first are tasks which we find difficult but whose difficulty we attribute to our own mental confusions or habits since we avoid building such obstacles into our computers we hardly find it noteworthy or creditable that they can complete such tasks and we can readily say "For us it is an intellectual task but for the computer it presents no problem". * It is not clear why we should include all such tasks in the A. I. fold. The second are tasks for which optimal performance is not readily definable. Even theorem proving can form an example of this: if the task is to reach a decision about a given putative theorem then performance can be measured, however, if it is to give the best possible proof then only performance relative to human standards and insights can be measured. Yet the whole aim of Dr. Meltzer's note is to escape from the anthropocentric tendency in A.J. If we allow recognition or machine learning to become activities in which we admit the possibility that the machine will do better than the human being and perhaps even develop new types of response which we cannot or do not wish to achieve by whatever sort of trying, then although we have elevated the pursuits into studies of machine performance on autonomous tasks we have lost the flavour of intellectuality which remains with the tasks as long as they continue to be mere mimicry.

* e. g. doing arithmetic in binary

The purpose of this excursion was by no means to decry Dr. Meltzer's programmatic suggestions, but rather to point out that once we concentrate on the tasks which humans find require intellectual effort, we are left with a rather arbitrary collection and meet a considerable problem in applying the name A. I. once we start to deal with cases where human abilities fall short of 'perfection'.

The ultimate objection to any programme of concentrating exclusively on the development of competence in a range of tasks, is, however, that we obtain intelligent behaviour only by selecting and combining the various tasks to be executed. Here we may recall the original discussion of the role of the decision-making stage in intelligent activity. Just as behaviour which is determined by factors beyond the decision-making stage of the reaction sequence are not intelligent actions of the entity whose decision-making stage was bypassed, so also are cases where a given skill was exercised by an A. I. device hardly eligible for consideration if the decision which skill to exercise and on what selection of input to exercise it has been taken outside that device.

This simple statement may seem to demand too much of our present fledgling attempts at A.I. However we are excessively charitable in our dealings with less well-endowed beings and this may mislead us in cases where the A. I. devices are not equipped to receive enough inputs to require any decision. Our natural tendency is to say that an animal which has no means of knowing of, say, impending doom is not shown to be unintelligent by the fact that it did not try to avoid that doom. The computer which is so programmed that the only input it can deal with consists of theorems to proved, would be dealt with excessively leniently on this basis; yet it is hard to criticise the designer Ior not incorporating input modalities not relevant to theorem-proving. The best approach, I think, is to look at a theorem - prover as a task-machine which should be compared to a skilled hand-movement in the human- given that exercise of the skill is appropriate the rest can be left to the unit concerned.

Yet if we pursue this course we risk classifying almost all present day A. I. work as the construction of mere task-machines; we seem to be doomed to wait until these are available and then to build selector mechanisms to choose one or the other as appropriate, if we are to approach real A. I.

Not having been able to shake off the anthropomorphic ring of "intelligence" let us now return to the human case. Many writers have suggested that

the evolution of man can be seen in neural terms as a progressive liberation from reflex instinctive reactions and that this has been mediated by the progressive profusion and increasing complexity of the neural network particularly in the cortex.

A striking feature of the organisation of the primate nervous system is the ability of various centres to effect the same behaviour, and if we plot this in terms of the developmental origin of the various centres we find that one and the same muscular contraction can occur as a reflex, as an integrated part of an unconscious motor skill, as part ot a skill of which only the initiation or abstinence is conscious, or as a result of conscious decision within or outwith a context of consciously organised actions, these levels arc mediated by centres which correspond to successive stages in the develop-ment of the system* and the system is so organised that the more recent centres can intervene to alter the course of many of the activities arising at lower levels, and that many activities can pass from the conscious trying level to a level at which thev are more or less automatic only being influenced by the highest levels when a major disorientation appears.

If we now consider a simple action such as striking a match, we may observe that on the first occasion when it was performed it occurred at the conscious level with tactile and visual input made available at that level. Through gradual practice the action becomes highly stereotyped and may be subsumed in a larger, hall-automatic pattern such as that of lighting a cigarette. It seems at least probable that the tactile and visual input needed to control the match lighting no longer travels via the higher levels of the cortex once the routine has become habitual. It is present and may cause conscious intervention but this is a different matter.

If we treat this type of development as an analogy we need not be too despondent over the obvious task-obsession in current A. 1. Perhaps these applications arc still in the "conscious" stage and, given some habituation, can later be pushed to a lower level and only looked at again when gross deviations are signalled to the module which then conducts the higher controlling and task selecting functions.

* it is not the intention to suggest that animals have existed in which, say, only the first, second and third types of control were present.

Attractive as this is, it highlights two points where we might improve A. I. research. First the suggestion just made that we might push the present taskB down to the level of habits incorporates no suggestion how we should do this. If we are to be true to the analogy we must do it in such a way that a hierarchy of habits within habits can evolve. It is not my intention to suggest the same types of activity as the use of sub-routines or modular programming - both of which have been with us some time now. Probably the most important development in terms of current programming styles would be the spontaneous labelling of statistically-detected recurring routines and their delegation complete with I/O adjustments to a slave processor without the intervention of a programmer.

The second point of importance is one which has been implicit in all I have said but which requires separate and emphatic repetition. This is that intelligence is a relative notion which still requires farther definition. The intelligence of human actions is relative to the circumstances, the abilities of the agent, and many other factors. Wc left the analysis of intelligent activity with the suggestion that the decision-making stage was erucial. I have also suggested that a machine can only be described as "intelligent" in relation to the operations ol the task selector and controller which it incorporates. In doing so 1 have argued for conscious attention to this particular function and for a development of A. 1. which amounts to the successive generalisation of the task selection and controlling function with a simultaneous sub-sumption of the previous highest level in the repertoire of the new highest level.

But if we are to proceed in this way there are still further questions to be posed at each stage in the development and we shall have to get used to asking these. Both decision-making and selection require criteria. An arbitrary decision can always be made - but that is not intelligence. An arbitrary selection of tasks can also be made and they can all be very sophisticated and perfectly executed but there will only be intelligence ii the selection was <u>appropriate</u>. What we ought to spend a great deal more time on is the definition of those criteria which are to decide amongst the alternative things which might be done. The charity with which wc exculpate the limited I/O machines allows this task to be stated already. Even within the chess-playing programmes decisions on broad strategy are quite in place, and are in fact the crucial aspects for assessing the intelligence of the computer's game.

Although this paper is primarily concerned with clearing the ground for a new look at this problem, 1 feel it would leave an unsatisfactorily negative impression if after having criticised other proposals 1 neglected to outline my own.

I would suggest that a definition of intelligence which leads naturally to the questions which ought to be asked about every real or abstract machine which is proposed in the A. I. field might run as follows:

> Intelligence is displayed when tasks for execution are selected from a range of alternatives in such a way that the aims of the agent are likely to be furthered with maximum economy of effort.

The key words and phrases in this statement will have to be interpreted anew from case to case. What precisely is to be meant by "task" or "economy" may not always be clear; the criteria for judging likelihood are also notoriously evasive. But most of all the relationship to "aims" is likely to prove a complication in computer applications, and before concluding 1 should like to outline the reasons for the inclusion of that term; 1 hope that we shall have an opportunity to discuss it in more detail presently.

If we recall the informal judgements we make of the intelligence of human beings it will be clear that these are made relative to certain presumed aims. Slashing one's wrists is a reasonable way of committing suicide but a less intelligent way of furthering a tempestuous love affair. Again, insulting one's benefactors is prima facie an unintelligent thing to do: however we might be given reasons which would show us that it was in fact quite intelligent in a given case. The aims of the agent are crucial, and when we have no other information we tend to suppose that the agent has very general "normal" human drives and aims.

Since one can neither assume this for a computer, nor interrogate it about other aims, it would be desirable to eliminate this particular relativity from our new enlarged definition. Unfortunately we cannot, I think: and it is through this notion that the contrast man-machine persists in my formulation.

Briefly, the taking of any decision on action must be based either on a selection with respect to the alternative actions as such or on a selection with respect to their consequences. If we are to avoid being drawn into a discussion of aims, we must deny that any evaluation of consequences can be relevant to the selections which we expect the "intelligent" computer to make for itseli. One can indeed state an even stronger conclusion, in that any evaluation of the alternatives which operates in such a way that predictable and differentially preferable outcomes must be compared is implicitly choosing between aims and must equally be avoided. If aims are not relevant we can only make decisions in such a way that the decision is its own payoff and is not rendered erroneous by any consequence whatever. My own feelings are that such decisions are either between vacuous alternatives or are to be compared to aesthetic judgements.

As I see rather greater scope than this for the applications of programmed machines, I suggest we grasp the nettle and start asking ourselves just how we do intend to generalise the aims of our existing programmes and machines.