ARTIFICIAL INTELLIGENCE AND BIOLOGICAL REDUCTIONISM

Margaret. A. Boden

This paper will be published in Beyond Neo-Darwinism, ed Mae-Wan Ho. · Academic Press.

gnitive Science Research Report

rial no: CSRP 021

e University of Sussex
gnitive Studies Programme
hool of Social Sciences
lmer
ighton BN1 9QN

ABSTRACT

Artificial Intelligence and Biological Reductionism

Margaret A. Boden

Anti-reductionism is not antiscientific if it offers positiv suggestions about empirically-based concepts to explain phenomena nc explicable in currently accepted terms. Artificial intelligence (A.I. is anti-reductionist in a number of respects. Computational concepts ca provide explanatory power over and above that of the more basic theorie in the life sciences, while being entirely compatible with them. Huma and animal intelligence, and psychological phenomena in general, ca usefully be thought of in these terms. Behaviourist psychology ar neurophysiology cannot express the phenomena concerned, because thei vocabulary has no room for the concepts of representation c intentionality. But A.I. is concerned with symbol-manipulating systems and these concepts are theoretically central to it. A.I. can be usefu in physiology, by clarifying what are the computational tasks which th nervous system is performing, and it may even illuminate some aspects c evolutionary and morphogenetic biology. A computational approach to lif and mind is entirely compatible with notions of human freedom. By it emphasis on the subject's representation of the world, it counters th mechanization of the world-picture brought about by the natura sciences.

ARTIFICIAL INTELLIGENCE AND BIOLOGICAL REDUCTIONISM.

Margaret A. Boden

(To be published in Beyond Neo-Darwinism, ed. Mae-Wan Ho. -- Academic Press.)

Reductionists usually see anti-reductionists as enemies of science, as a warring intellectual army whose aim is to oust science and substitute mystery. For the reductionist assumes that all problems in a given domain are soluble in terms of available types of concepts, preferred as empirically well-grounded and ontologically basic -- an assumption the anti-reductionist denies. In truth, however, anti-reductionists often wish to oust mystery and substitute science. For the mysterious, though unknowable by familiar methods of enquiry, is potentially intelligible in other ways (hence the promise implicit in titles like "The Mystery of the Red Room"). Only the mystical is unknowable *tout court*. Science and mysticism thus make strange bedfellows, but science and mystery can enter into fruitful liaison.

Far from being attempts to oppose or undermine science, anti-reductionist claims may be (though admittedly they are not always) intellectual challenges aimed at strengthening it. Anti-reductionist critiques are partly negative exercises, intended to show the incapacity of current approaches to express or explain the problematic phenomena. But ideally, the anti-reductionist will also offer positive suggestions about what alternative types of concept might be better suited to this task. If these concepts are clear, empirically fruitful at least on their own level, and grounded in or demonstrably compatible with the more basic scientific concepts favoured by the reductionist, they may hope to be welcomed into the scientific community instead of being shunned as intellectual outlaws.

With respect to living things, the reductionist temperament shows itself in such assumptions (for instance) as that neo-Darwinism can answer all evolutionary questions, molecular biology all questions about individual morphogenesis, stimulus-response psychology all questions about behaviour, and neurophysiology all questions about the mind. These assumptions are sometimes made explicit. More often, they remain unspoken -- but they are no less powerful for that.

Such forms of biological reductionism have been countered by various anti-reductionist critiques, some of which are represented in this volume. The negative poles of these critiques have much in common, but their positive aspects are more diverse. They do not deny the truth of neo-Darwinism, but see it as conceptually inadequate for addressing certain important biological problems. Positively, they offer other concepts they regard as better-suited to these problems. These differ from case to case, but in general they are not *alternatives* to the reductionist's conceptual base so much as theoretical *complements* to it.

I shall concentrate on the positive exercise, suggesting ways in which *computational* concepts (drawn from computer science and artificial intelligence, or "AI") might offer explanatory power that is not provided by the more basic theories in the life-sciences -- theories with which these concepts are nonetheless entirely compatible. They can be useful to biologists, to physiologists, and to ethologists, but their most obvious application is to psychology: the science not of body, but of mind.

The reason for this is that AI uses the methodology of computer programming to study the content and function of various sorts of knowledge [Boden, 1977]. It focusses on the structure and organization of internal models, or symbolic representations, and on how these can be transformed in complex and flexible ways so that the representational system can cope with changing and largely unknown situations. It is this flexibility, and the large degree of autonomy that it involves (so that the system can achieve its ends, within limits, irrespective of external conditions), that constitute *intelligence*.

Accordingly, AI programs are quite unlike the more familiar types of computer program (dealing with wages, tax-returns and the like), which are rigid, unimaginative, and inflexible -- in a word, stupid. AI programs specify computations enabling computers to do such things as: recognizing objects seen in widely varying positions or lighting conditions; planning complex tasks involving unpredictable conditions; conversing (by teletype) in natural language; understanding spoken speech; making sensible guesses where specific knowledge is not available ... and the like. It should be evident from these examples that "computation" here does not mean "counting", but *any* symbolic process of inference, comparison, or association. The symbolism may be numerical (for counting is one example of computation), or it may be of some other form (such as verbal, visual, or logical).

AI is sometimes thought of as a branch of cybernetics, and if one accepts Wiener's [1948] broad definition of cybernetics as "the science of control and communication in the animal and the machine", then so of course it is. Cybernetics itself was an anti-reductionist exercise. Negatively, it claimed that the concepts of physics and chemsistry alone could not express the nature and functioning of homeostatic control in biological organisms or man-made mechanisms. Positively, it showed how these matters could be expressed, and rigorously discussed, in terms of new -- information-processing -- concepts which were logically independent of their embodiment in a given body or artifice.

That a steam-engine needs a part functioning as a governor is a truth of cybernetics. That the governor is embodied in two metal balls linked to a rotating shaft is, rather, a truth of nineteenth-century engineering -- for *any* mechanism with equivalent functional properties would do. The physics of centrifugal forces explains how it is possible for the steam-governor to do its job, much as biochemistry explains how it is possible for blood-sugar or body-temperature to be regulated. But physics and chemistry concern the possibilities for embodiment of cybernetic principles, rather than their essential nature. That must be expressed at a different level, in information-processing terms defining signals and feedback.

Wiener himself suggested that not only biological functions, but psychological regulations too, could be cybernetically explained. In an influential discussion of behaviour, purpose, and teleology written with biologist colleagues, Wiener used the key terms so that teleological or purposive behaviour was synonymous with "behaviour controlled by negative feedback,[11] by which he meant behaviour that is •'controlled by the margin of error at which the CbehavingD object stands at a given time with reference to a relatively specific goal ... The signals from the goal *are* used to restrict outputs which would otherwise go beyond the goal" CRosenblueth, Wiener, & Bigelow, 1943, p. 193. A paradigm case of "purposive[11] behaviour on this account would be, he said, that of a machine designed to track a moving luminous goal, or target.

But truly purposive behaviour involves more than this. It is guided by internal representations of the "goal", of plans of action for reaching the goal, and of various criteria (including moral values and personal preferences) for evaluating these plans, wherein the overall goal may not correspond at all closely to <u>any</u> external state of affairs or target CBoden, 19723.

For example, people can engage in the highly complex behaviour of going on a unicorn-hunt, and several centuries ago they commonly did so. But, clearly, we could not explain their activities in terms of feedback signals from the fleeing unicorns ~ for there *ne\/er* were any unicorns. Rather, our explanations must be in terms of the hunters' ideas and inferences about unicorns, including their beliefs that unicorns are likely to be found in forests, their heads resting in the laps of virgins, and that their horns have magical properties. We need to show how these ideas and inferences might be represented, organized, and accessed, in such a way as to generate the behaviour — and experience — of the people concerned.

These are computational matters, in the sense defined above. They cannot be expressed in the terminology of traditional cybernetics, which focusses on feedback and adaptive networks, and which defines information-processing in quantitative rather than qualitative terms. It could not express the difference between a mermaid and a unicorn, for instance, nor the very different planning activities involved in hunting them (does one head for the seaside or for a forest?). To understand intelligence we need computational concepts, specifically designed to define complex transformations within qualitatively distinct symbolic representations such as these.

Such concepts are supplied by AI, for it is concerned with symbol-manipulating systems wherein the degree ~ and the nature — of match-mismatch between representation and reality (or even possibility) may vary enormously. So AI differs importantly from classical cybernetics, and is helpfully described as a "cybernetic[11] enterprise only in the broadest sense.

AI, then, sees minds (whether animal or human) as symbolic systems containing many internal representations of aspects of the world (and possible worlds), and a variety of rules for building, changing, comparing, and inferring from them. Psychological questions, on this view, concern the structure and content of mental representations, and the ways in which they can be generated, augmented, and transformed. Thinking, experience, and motivation — and psychological differences between individual people -- are all grounded in computational processes.

One main strength of the AI-approach to psychology is its emphasis on rigor. AI offers precisely definable concepts, because a program has to be expressed clearly (as a set of instructions defining specific symbol-manipulations) if the computer is to accept it. In aiming for rigor -- indeed, in setting a new standard for rigor in psychological explanation -- AI differs from other anti-reductionist approaches to the life sciences. The positive sides of these critiques commonly invoke concepts which, while they may have some phenomenological plausibility and intellectual resonance, are not expressed with clarity and are therefore difficult to apply in a scientific research-programme [Lakatos, 1970]. One example is Piaget's concept of "equilibration," mentioned below.

Another strength of the computational approach to psychology is its highlighting of process as well as structure. This derives from the fact that a program has to tell the computer not only what result to produce but also how to produce it. Non-computational psychologists often take psychological change for granted, assuming that it can be sufficiently specified by stating the initial and final mental states involved. However, the process of mental transformation is itself theoretically problematic. In a programming context, a failure to suggest any way in which the change might be effected will show up as a glaring gap in the program, a gap over which the uninstructed computer is unable to leap. Some computational account of how to make the leap must be supplied if the program is to function.

AI is relevant not only to psychology but to neurophysiology also, and this in at least two ways. First, it helps us achieve a clear account of what computational tasks the brain and peripheral nervous system must be performing. This can help the neurophysiologist in formulating questions and hypotheses about which bodily mechanisms are performing these tasks, and how. Second, AI may help us understand the potential and limitations of the computational or representational power of various sorts of hardware, including neurophysiological mechanisms.

These two "physiological" uses of AI differ from the familiar use of computer technology to model physiological theories, simulate nervous nets, etc. In those exercises, the prime intellectual input comes from the concepts already developed in neurophysiology. But in AI-based studies of the sort I have in mind, a crucial part -- sometimes the larger part -- of this input is supplied by computational research as such.

The first of these physiological advantages has been stressed by David Marr, whose early theoretical papers on the cerebellum and cerebrum attracted the interest of brain scientists, and whose more recent (computational) work on the visual system is deservedly one of the strongest current influences in AI [Marr, 1982]. As Marr put it when embarking on his computational research-programme:

> The situation in modern neurophysiology is that people are trying to understand how a particular mechanism performs a computation that they cannot even formulate, let alone provide a crisp summary of ways of doing. To rectify the

> situation, we need to invest considerable effort in studying the computational background to questions that can be approached in neurophysiological experiments.
>
> Therefore, although [my work] arises from a deep commitment to the goals of neurophysiology, the work is not about neurophysiology directly, nor is it about simulating neurophysiological mechanisms: it is about studying vision. It amounts to a series of computational experiments, inspired in part by some findings in visual neurophysiology. The need for them arises because, until one tries to process an image or to make an artificial arm thread a needle, one has little idea of the problems that really arise in trying to do these things [Marr, 1975, p.3].

Clearly, these problems arise for the human being trying to thread a needle or interpret a visual scene, but we are not aware of them because they are solved at computational levels beyond the reach of consciousness.

Doubtless, our evolutionary heritage has provided us with bodily mechanisms especially well-suited to such tasks, which mechanisms are the particular focus of neurophysiology. But to find out just what they are, what they do, and how they function, we need to understand the task itself. And since "task" is essentially a computational concept, albeit one of everyday English, we must understand the task in computational terms. In short, Marr's is an anti-reductionist exercise in the sense outlined earlier: physiology alone, physiology unaided by computational insights, cannot suffice to explain behaviour or experience.

Marr puts much greater emphasis on neurophysiology, as constraining the details of computational research, than do some other workers in AI. For instance, Marr's research on vision considers psychological optics in some detail and makes a concerted attempt to use (not merely to bear in mind) knowledge about the anatomy and physiology of the retina and early stages in the visual pathways. This is not simply because he happens to be interested in neurophysiology and biological mechanisms whereas others are not. For some AI-workers on vision consider optics only in very general terms and ignore neurophysiology, on the principled ground that physiological (hardware) implementation is theoretically independent of questions about computational mechanisms. This is a widely shared view in AI, and many physiologists are sceptical about AI accordingly.

In principle the view is true, as Marr himself admits (see the quotation above). It is analogous to the fact that the engineer's problem in actually making a governor is theoretically distinct from the cybernetician's problem in identifying the abstract principles of gubernatorial control. But in practice, it may be that we need to study the varying computational powers of distinct (electronic or physiological) hardware if we are to understand how tasks are performed by finite mechanisms in real time. Even if this is agreed, however, doing it is no simple matter. Within AI there are real, and largely unresolved, differences of opinion about the computational potential of relatively low-level, peripheral mechanisms (most of which seem to be biologically specialized to a high degree) as opposed to high-level cerebral mechanisms (which appear rather to be general-purpose in character).

This brings us to the second point noted above, that AI could help us appreciate what computational powers an organism might have <u>in virtue of its neurophysiology</u>. This would be useful to a wide range of biologists, not just to "physiologists" narrowly defined. I have discussed elsewhere some implications germane to a range of ethological problems, concerning both motor action and perception [Boden, in press (a)]. One relevant example is Hinton's recent research, which is a significant advance in the computational modelling of vision [Hinton, 1981].

Hinton's work is focussed on the computation of shape. Most AI vision-workers have assumed that the shapes of three-dimensional objects <u>have to be</u> computed by way of high-level, top-down processes. But Hinton has shown that this is not so. (It does not follow, of course, that top-down processes <u>are not</u> used by humans or other animals in the perception of shape.) He has designed low-level, dedicated hardware, mechanisms that are capable of cooperative computation, or parallel processing. These systems can perform shape discriminations -- such as recognition of an overall Gestalt -- commonly believed (even within AI) to require relatively high-level interpretative processes.

Hinton's mechanisms rely on excitatory and inhibitory connections between computational units on various levels that appear to have an analogue in the nervous connectivity of our own visual system. He not only provides an <u>example</u> of a mechanism that can compute shape in a surprising fashion, but also gives a <u>general proof</u> that many fewer computational units are necessary for the parallel computation of shape than one might initially have supposed. This proof lends some more physiological weight to the model, since the human retina apparently has enough cells to do the job.

Hinton's results suggest also that the way in which an object is represented may be radically different, depending on whether it is perceived as an object in its own right or as a part of some larger whole. This might account for the phenomenological differences between perceptual experiences of which we are reminded by those philosophers [e.g. Dreyfus, 1972] who argue that AI is essentially unfitted to model human minds -- or bodies. Hinton believes that his computational model of spatial relations can be used to understand motor control in a way analogous to the mechanisms of muscular control in the human body. His preliminary (unpublished) work suggests an efficient procedure for computing a jointed limb's movements and pathway through space -- a problem that can be solved by traditional computing techniques only in a highly inefficient manner.

Many commonly-expressed philosophical criticisms of AI and cognitive psychology, like that of Dreyfus just mentioned, may be invalidated by these recent computational developments. Such criticisms are explicitly anti-reductionist in nature, accusing AI itself of attempting to reduce all psychological phenomena to forms of serial processing. Evidently, to be an "anti-reductionist" can involve highly varying claims, and one may consistently support one form of anti-reductionism without endorsing another.

If AI is relevant to psychology, to physiology, and to ethology, what of biology more generally conceived? The idea that a computational approach might contribute to theoretical biology is not entirely novel, for various biologists have made comparable suggestions:

> Language ... I suggest may become a paradigm for the theory of General Biology CWaddington, 1972, p.2893.

> The view of the organism as an hypothesis-generating and testing system ... could transform biology by placing model construction and observation at the centre of the biological process, not at the evolutionary periphery, the phenomenon of Mind [Goodwin, 1972, p. 2673.

> The classical cases of pattern regulation whether in development or regeneration ... are largely dependent on the ability of the cells to change their positional information in an appropriate manner and to be able to interpret this change CWolpert, 1969, p.83.

> The problems of biology are all to do with <u>programs</u>. A program is a list of things to be done, with due regard to circumstances CLonguet-Higgins, 1969, p. 2293.

All these quotations except the last rely on familiar cognitive terms such as "language", "knowledge", "interpretation", or "information", rather than the more technical "program" or "computation". But, as I have argued elsewhere CBoden, 19813, their approach to theoretical biology is significantly similar to that of people influenced by AI. And all, of course, are anti-reductionist in the sense defined above.

Piaget is another example of someone who sought to complement the more basic biological concepts by others of a higher theoretical level. His lifelong opposition to reductionism in biology, as well as in psychology, was grounded in his interest in development. Developmental change brings about forms that are in some sense novel. In his account of "equilibration", Piaget tried to illuminate the way in which new, more differentiated, structures arise out of simpler ones, whether in psychological, embryological, or evolutionary development. He recognized the profound theoretical problem of how it is possible for harmonious structural novelties to develop, along with novel integrative mechanisms whereby the overall regulation of the system is maintained. And he tried to give an account of spontaneous, as opposed to reactive, structural development. However, his concept of equilibration was so vaguely expressed that it provided no clear questions, still less clear answers CBoden, 19823.

Computational studies have not given clear answers either. But they may help give us a better sense of the sorts of questions that need to be addressed in understanding adaptive development, whether in body or mind. From a computational point of view, the generation of new forms, as well as their consequent evaluation, must take place within certain structural constraints CBoden, in press (b)3. In general, to understand any sort of development or creativity would be to have a theory of the various transformations, at more or less basic levels, by which the relevant structural potential can be selectively explored.

A theory of morphological development, for instance, would explain how it is possible for a gill-slit to be transformed into a thyroid gland, or a normal blastula into a deformed embryo or non-viable monster; also, it would explain why certain fabulous beasts (such as mermaids, but <u>not</u> unicorns) could only have been imagined, not created. In biology as in psychology, we need some account of generative processes that can explore the space defined by background constraints, and of more truly creative processes that can transform these constraints themselves.

If an adaptively developing system is to be able to judge the "interest" of its explorations, it must have some form of evaluation criterion. Lenat's "automatic mathematician" is an AI program designed to shed light on this matter with respect to the example of mathematical creativity [Lenat, 1977]. The program starts with elementary set-theory, generates new concepts, and decides which to explore further. It uses several hundred heuristics (not just a few transformational rules) to explore the space defined by a hundred primitive concepts of set-theory. From this elementary base the program generates and follows up various concepts of number-theory (including one minor theorem that had not previously been defined). Among its heuristics are some which evaluate the (mathematical) interest of newly generated concepts.

In biological, as opposed to psychological, development it is natural selection that functions (<u>post hoc</u>) as the evaluative criterion. This fundamental tenet of neo-Darwinism is not in question. But there are various biological phenomena which, on a neo-Darwinist account of evolution, are very puzzling, For example, the fraction of DNA that does not code for the synthesis of specific proteins increases phylogenetically, and species have evolved remarkably quickly. Work in automatic programming has suggested that such facts are not explicable in terms of the neo-Darwinist mutational strategy ("Random-Generate-and-Test"), because its combinatorics are horrendous [Arbib, 1969]. Instead, some strategy of "Plausible-Generate-and-Test" is needed, whereby mutations of a type likely to be adaptive become increasingly probable. The initial heuristics evolve by random mutation and natural selection, but -- since they are embodied as DNA and their "target" for interpretation is itself DNA -- they can then develop by modifying each other. They function as heuristics recommending certain "copying errors" and preventing others, and the transformational processes they influence are gene substition, insertion, deletion, translocation, inversion, recombination, segregation, and transposition.

These genetic transformations have been discussed in the light of procedures known in AI as "Production Rules" [Lenat, 1980]. Production Rules are IF-THEN rules specifying a certain action, given a particular condition. Lenat suggests that the IF... part of the heuristic might be specified by proximity on the DNA molecule, whereas the THEN... part could direct gene rearrangement, duplication, placement of mutators and intervening sequences, and so on. For instance, one heuristic might be that gene recombinations should involve neighbour-genes rather than genes at opposite ends of the DNA string: in a creature where genes for morphologically related structures happened to lie next to each other, this heuristic would encourage mutations of both genes together, which would tend toward a structurally integrated evolution. Although these ideas are highly speculative, and their value is not proven, they suggest that concepts drawn from AI might be useful in thinking about biological evolution.

But perhaps AI, in relation to the life sciences, is not mere harmless speculation but rather a dangerous illusion? Certainly, many people assume that AI simply cannot be a humanely anti-reductionist project. Indeed, they scorn AI as an essentially inhuman reductionist influence, one that offers an "obscene" and "deeply humiliating" image of man CWeizenbaum, 1976D.

In general, anti-reductionists commonly complain that reductionist approaches to life and mind somehow rob us of our humanity — to which reductionists are unfortunately apt to retort that theoretical emphasis on "humanity" is no more than self-indulgent sentimentality. Skinner's attacks on notions like "freedom" and "dignity" are notorious, and the molecular biologist Monod has managed to epater les bourgeois by such relentlessly reductionist remarks as: "The cell is a machine. The animal is a machine. Man is a machine." Psychological approaches based on AI also assume that man is a machine, even daring to compare the mind with the metallic artefacts of the electronics workshop. So how can AI be anything but a form of reductionism? If man is a machine, what room is there for humanity?

If Monod's remark "Man is a machine" is taken to mean that the bodily processes underlying and generating human behaviour and experience (including moral conduct and insights) are describable by physics or molecular biology, it is — as far as AI is concerned — true. But if it is taken to mean that the concepts of natural sciences such as these suffice to express the nature and regulation of the human mind, then — according to AI — it is false.

Much as the steam-engine governor cannot be described, qua governor, in physical terms but only in cybernetic language, so the mind cannot be conceptualized physiologically, but only computationally. Describing a system (whether person or computer) ££ a symbol-manipulating system is conceptually distinct from describing the physical hardware that embodies the computational powers concerned. The former type of description requires computational concepts, whereas the latter employs the terms of physics, chemistry, and physiology. As the poet Blake foresaw, these natural sciences have had a dehumanizing influence, encouraging a "single vision" that has insidiously undermined people's sense of personal autonomy and responsibility. This is unsurprising, for no science that lacks the concept of representation can even acknowledge humanity, still less explain it.

Computational psychology does not support the mechanization of the world-picture brought about by the natural sciences, and by "scientific" forms of psychology such as behaviourism. For, unlike these, it can distinguish "subjective" truths (about ideas, aspirations, and beliefs) from "objective" truths (about brains and other physical things). What is more, it concentrates firmly on the former, admitting the influence on our lives of shared cultural beliefs, of individual ideas, interests, purposes, and choice, and of self-reference and self-knowledge CHofstadter, 1980D.

In sum, AI emphasizes the richness and subtlety of our mental powers, a richness that has often been intuitively glimpsed (at least by literary artists) but never theoretically recognized by scientists. Many humanists reject the reductionist influences of natural science with such passion that they come close to adopting mysticism. There may be good reasons for embracing mysticism, but the undeniable inability of natural science to conceptualize subjectivity is not one of them. Nor is the "wholistic" integration of biological phenomena such as embryogenesis or regeneration. In providing rigorous hypotheses about the mental processes that underlie humanity and make it possible, AI promises a scientific understanding of the most recent product of evolution, intelligence. Its promise has yet to be fulfilled, but the crucial point is that AI sees mind not as mystical, merely as mysterious.

oooooOOOOOOOooooo

REFERENCES

ARBIB, M. (1969) "Self-Producing Automata — Some Implications for Theoretical Biology,[11] in C. H. Waddington (ed.), Towards £ Theoretical Biology, Vol. 2 (Edinburgh: Edinburgh U. P., 1969), pp. 204-226.

BODEN, M. A. (1972) Purposive Explanation jm Psychology (Cambridge, Mass.: Harvard Univ. Press).

BODEN, M. A. (1977) Artificial Intelligence and Natural Han (New York: Basic Books).

BODEN, M. A. (1981). "The Case for a Cognitive Biology,[11] in M. A. Boden, Hinds and Mechanisms: Philosophical Psychology and Computational Hodels (Ithaca, N.Y.: Cornell Univ. Press), pp. 89-112.

BODEN, M. A. (1982) "Is Equilibration Important?," _Brit. J. Psychol._, 73, 165-173.

BODEN, M. A. (in press, (a)) "Artificial Intelligence and Animal Psychology," _New Ideas in Psychology_, 1 (to appear 1982/83).

BODEN, M. A. (in press, (b)) "Failure is Not the Spur," in O. Selfridge, M. Arbib, & E. Rissland (eds.), _Adaptive Control in Ill-Defined Systems_ (Plenum Press).

DREYFUS, H. L. (1972) _What Computers Can't Do: A Critique of Artificial Reason_ (New York, Harper & Row).

GOODWIN, B. C. (1972) "Biology and Meaning," in C. H. Waddington (ed.), _Toward a Theoretical Biology, Vol. 4_ (Edinburgh University Press), pp. 259-275.

HINTON, G. E. (1981) "Shape Representation in Parallel Systems," _Seventh Int. Joint Conf. Artificial Intelligence_, (Vancouver), 1088-1096.

HOFSTDATER, D. (1980) _Godel, Escher, Bach: An Eternal Golden Braid_ (London: Penguin).

LAKATOS, I. (1970) "Falsification and the Methodology of Scientific Research Programmes," in I. Lakatos & A. Musgrave (eds.), *Criticism and the Growth of Knowledge* (Cambridge Univ. Press), pp. 91-196.

LENAT, D. B. (1977) "Automated Theory Formation in Mathematics," *Proc. Fifth Int. Joint Conf. Art. Int.*, Cambridge, Mass.), 833-842.

LENAT, D. B. (1980) *The Heuristics of Nature: The Plausible Mutation of DNA*, Stanford Univ. Computer Science Dept., Report HPP-80-27.

LONGUET-HIGGINS, C. H. (1969) "What Biology is About," in C. H. Waddington (ed.), *Toward a Theoretical Biology, Vol. 2*, pp. 227-236.

MARR, D. (1975) *Analyzing Natural Images: A Computational Theory of Texture Vision*. (Cambridge, Mass.: MIT AI Dept. AI-Memo 334).

MARR, D. (1982) *Vision* (Cambridge, Mass.: MIT Press)

ROSENBLUETH, A., N. WIENER, & J. BIGELOW (1943). "Behavior, Purpose, and Teleology," *Philosophy of Science*, 10, 18-24.

WADDINGTON, C. H. (1972) "Epilogue," in C. H. Waddington (ed.), *Toward* a *Theoretical Biology, Vol. 4* (Edinburgh Univ. Press), pp. 283-289.

WEIZENBAUM, J. (1976) *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: Freeman).

WIENER, N. (1948) *Cybernetics, or Control and Communication in the Animal and the Machine* (New York: Wiley).

WOLPERT, L. (1969) "Positional Information and the Spatial Pattern of Cellular Differentiation," *J. Theoretical Biology*, 25, 1-47