FAILURE IS NOT THE SPUR

Margaret. A. Boden

This paper will be published in _Adaptive Control in Ill-Defined Systems_, eds. O. Selfridge, M. Arbib, & E. Rissland (New York: Plenum, 1982/3)

Cognitive Science Research Paper

Serial no: CSRP 015

The University of Sussex
Cognitive Studies Programme
School of Social Sciences
Falmer
Brighton BN1 9QN

OlT

NATO Workshop on "Adaptive Control in Ill-Defined Systems" (1981),

Abstract of: FAILURE LS NOT THE SPUR

Margaret A. Boden

University of Sussex

An ill-defined system is one with respect to which some *prima facie* appropriate theoretical description is not suitable, because it treats the system as being more well-defined in a specific respect than it really is.

One way in which the mind is an ill-defined system is that its adaptation is not initiated and controlled primarily by response to failure. The concept of failure is not appropriate to all intelligent behaviour, because it implies an attempt to achieve some specific goal: one can always sensibly ask "Failure to do what?". But organisms do not always aim at a specific goal. Rather, their behaviour is often spontaneously exploratory in nature.

Adaptation thus has to be conceptualized in terms of some sort of creative urge. This should be thought of as an essentially structured phenomenon,, whereby formal potentialities are explored and/or radically transformed, and the results evaluated for their usefulness or interest.

Piaget realized that the life sciences need some concept of autonomous adaptive creativity, characterized in structural terms. His concept of "equilibrium" was supposed to illuminate how new, more differentiated, structures arise out of simpler ones — whether in evolutionary, embryological, or psychological development. But his account was too vague to be theoretically useful.

Failure does not always lead to adaptation, even when it is recognized. It has been said that time is needed for the "consolidation" of new theories. Consolidation involves a variety of meta-activities focussed on the currently developing structure, which lead to increased elegance, economy, clarity, and control. The system must be able to recognize when consolidation has been achieved, so that more radical structural changes may appropriately be explored.

A key concept is "interestingness," for two different questions: (1) Has the generative potential of the current cognitive terrain been adequately mapped, or is further consolidation required?; (2) Is this novel result of my exploratory activities worth pursuing, or should it be ignored? A theory of adaptive control should specify the criteria for such evaluations, some of which will be domain-specific.

Unpredictable contingencies can be allowed a creative role by a structural-transformational theory of adaptive control, provided that they can be integrated into the exploratory activities of the mind concerned.

MARGARET A.BODEN: BACKGROUND AND RESEARCH INTERESTS

My qualifications are in the medical sciences, in philosophy, and in psychology, and I have also followed up an interest in artificial intelligence. I studied these fields because of their relevance to these issues: the nature of the mind and mental phenomena; the role of purpose and intentionality in theoretical psychology and the philosophy of mind; how it is possible for the mind to be embodied; the philosophical relation between psychology and physiology or biology; and the implications of various types of theoretical psychology and biology for moral thinking about individuals and society. These questions are the focus of my books:

- Purposive Explanation in Psychology (Harvard Univ. Press, 1972; Harvester Press, 17787;
- Artificial Intelligence and Natural Man (Basic Books; Harvester Press, 1972);
- Piaget (Fontana; Viking Penguin, 1979);
- Minds and Mechanisms: Philosophical Psychology and Computational Models (Harvester Press; Cornell Univ. Press. 1981 – in press).

MARGARET A. BODEN: Brief c.v.

Degrees: M.A. (Cantab) – Natural Sciences (for Medicine);
Moral Sciences (Philosophy).
Ph.D (Harvard) – Social Psychology (Cognitive Studies).

Awards: Major Scholar, Newnham College, Cambridge.
Harkness Fellow.
McMichael Premium of Inst. Electronic & Radio Engineers

Current Position: Professor of Philosophy and Psychology,
University of Sussex, England.

Selfridge, M. Arbib, & E. Rissland (New York: Plenum, 1982)

FAILURE IS NOT THE SPUR

Margaret A. Boden

The concept of "ill-defined system" collapses into triviality if it is used to refer to any system that has not yet been well defined. One might take it rather to mean one that can never be understood in a well-defined way. This interpretation, however, invites troublesome disputes over what is to count as "well-defined", and also prejudges the question of whether human knowledge will ever be adequate to the system concerned. For instance, Schrodinger's wave-equations are mathematically well-defined, but they concern quantum phenomena which many would regard as a paradigm case of ill-definedness; and though the Copenhagen School believed this ill-definedness to be grounded at the ontological level, Einstein cited his conviction that "God does not play at dice" in interpreting quantum indeterminacy as a merely epistemological matter.

So let us rather say that an ill-defined system is one with respect to which certain _prima facie_ relevant types of theoretical description are inappropriate, because they treat the system as being more well-defined in a specific respect than it actually is. This prompts us to specify the ways in which certain currently available theoretical approaches are inadequate to characterize the system, while Leaving open the question whether any satisfactory theoretical description of it can in principle be found.

Minds (especially human minds) are ill-defined in a number of ways. That is, there are several sorts of theoretical description that one might expect to apply to mental phenomena but which are in fact inappropriate, because each wrongly assumes that minds are well-defined in some specific way in which they are not. As I have argued at length elsewhere CBoden, 19723 there are at least five aspects of mental life with respect to which this is true. (It is sometimes claimed that the insights of Merleau-Ponty, Wittgenstein, and Godel show that minds are ill-defined in other ways too; I have discussed these arguments in CBoden, 1977, ch. xivD.)

First, minds are intentional systems, whose actions are mediated by inner representations; so they _cannot_ be described by the natural sciences, whose terms do not make any distinction between the "subject" and the "object" of thought. Second, because of what Perry called "the independent variability of purpose and belief" and others have called "the hermeneutic circle" [Perry, 1921; Taylor, 19713 there can be _no_ behavioral definition of their goals, intentions, or beliefs. Third, they are symbolic systems within which complex structural transformations take place, transformations that _cannot_ be defined in the (quantitative rather than qualitative) information-processing terms of classical cybernetics — i.e. the cybernetics of feedback and adaptive networks, prior to the incorporation of notions of symbolic knowledge structures, _etc_., gleaned from artificial intelligence (AI). Fourth, they are systems with a rich generative potential, which _cannot_ be conceptualized in terms of general laws linking dependent and independent variables. And last, they use many purely conventional symbols, so that statistical averaging over the stimulus input may *fail* to capture crucial differences: compare the effects of a telegram saying "Our son is dead" and one saying "Your son is dead", where only a single letter distinguishes the two messages.

All these characteristics, each of which identifies a sense in which minds are ill-defined systems, apply also to complex AI programs. Since the latter are, in one important sense of the term, eminently *well-defined* — for they are rigorously specified by the instructions comprising the program — it should be clear that, as I noted above, "ill-definedness" is not an absolute term: it is always relative to some type of definition of explanation which one might *prima facie* expect to be relevant to the system in question.

A sixth way in which minds are ill-defined systems, which is especially important in relation to theories of adaptive control, is that they do not function merely so as to minimize and correct failures. They are not controlled simply by the degree of match-mismatch between the current state and a specific goal-state. That is, failure is not the prime trigger of cognitive change, and adaptive control cannot be explained in terms only of differential response to distinct classes of failure.

There are two reasons for doubting the centrality of failure in development. First, even where behavior is of a sort to which the concept of failure can sensibly be applied, failure is neither necessary nor sufficient for adaptive change. Adaptation without the goad of failure is shown, for example, by children drawing maps, who spontaneously improve their spatial representations although they are successfully solving the problem set [Karmiloff-Smith, 1979]. To be sure, their increased computational efficiency enables them to solve problems later which they could not have solved before, but it was not forced on them by any earlier mistake. Failure without adaptation occurs when it is recognized and yet ignored, treated as an inexplicable nuisance rather than as a spur to development. For instance, studies of children's understanding of balance-problems suggest that events initially ignored as anomalies are only later taken seriously as counterexamples enabling refinement of the child's current theory [Inhelder & Karmiloff-Smith, 1975].

Second, a point which will turn out to be intimately related to the previous one: theories focussed on failure implicitly assume that the mind is well-defined in a way in which it is not. Failure can be the main factor in adaptation only if organisms always aim at some specific goal. For a failure is always a failure to do something specific, in that the concept of "failure" invites the question "failure to do what?" Insofar as behavior is not goal-directed, then, the mind is an ill-defined system requiring explanation in other terms.

In fact, creatures do not always aim at pre-defined goals, but often appear to delight in activity for its own sake. Or rather, they commonly aim at goals (such as economy, clarity, elegance, and interestingness) which are high level meta-goals that control the adaptive exploration of their own cognitive processes. Because people commonly think of "goal-directedess" in terms of relatively specific, well-defined goals, to conceive of adaptation in terms of response to failure is to risk losing sight of the exploratory aspects of thought to failure, which is neither necessary nor sufficient for adaptation. They are controlled by more general considerations than match-mismatch with a specific goal-state, and they are grounded in some relatively autonomous creative urge (n.b. this concept is not "ill-defined", as I am using the term, but vague).

How can we conceptualize such an urge (the evolutionary advantage of which is obvious)? Like the "track and trail" strategies of lowly organisms [Selfridge, in press], it must both lead the creature to engage spontaneously in novel behaviour, and enable it to take advantage of the results when this exploration throws up something useful. This of course implies that the creature has available some form of evaluation function in terms of which it can recognize something "useful", or at least "interesting". "Track and trail" may describe some aspects of the exploratory activity of the higher animals (for instance, the locomotion of predators who have not yet scented game). But it cannot be used as a paradigm for all adaptive control, because (as posited above) the mind is a richly-structured system that cannot be described by classical cybernetics. If creative exploration is not to degenerate (sic) into mere chaotic novelty, the generation of new forms, as well as their consequent evaluation, must take place within certain structural constraints. In general, to understand a creative phenomenon 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.

For instance, mathematical, scientific, and artistic creativity involve the deliberate exploration and disciplined transformation, or relaxation, of structural constaints. Non Euclidean geometry originated in the (deliberate) dropping of Euclid's last axiom. Kekule's discovery of the benzene ring involved the (unconscious) transformation of one topological structure into another: it is significant that Kekule dreamt not of a little girl's hoop, but of a snake biting its tail -- in other words, of a closed curve that one would have expected to be an open curve. A prime root of Einstein's creative achievement was his query whether the concept of simultaneity can be analysed and the resultant "parts" variously combined to form distinct conceptual structures. And the development of tonal into atonal music, in broadly the way in which this happened, can with hindsight be seen as intelligible, and even inevitable [Rosen, 1978]. In all of these cases, of course, the initial transformation (which may or may not have been consciously effected) was followed by some sort of evaluative assessment, whether by mathematical proof, experimental method, or artistic discipline. Everyday adaptation presumably involves similar transformational processes, with varying opportunities for conscious initiation and/or control, and varying criteria of "interestingness".

In the biological domain, a theory of morphological creativity 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 could only have been imagined, not created [Boden, 1981]. In biology as in psychology, we need some account of processes that can explore the space defined by background creative constraints, and of processes that can transform these constraints themselves.

the development of cognitive skills, positing autonomous motives such as "competence" or "curiosity", but they have said little or nothing about mental structures [White, 1959]. Piaget, by contrast, realized that the life sciences need some concept of autonomous adaptive creativity, characterized in structural terms. In his account of "equilibration" he 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, as I have argued elsewhere [Boden, 1982] his concept of equilibration was so vaguely expressed that it provided no clear questions, still less clear answers. Attempting to explain psychological adaptation, he made vague reference to "positive and negative disturbances," "regulations," and "regulations of regulations". These cybernetically inspired concepts cannot express qualitatively distinct mental structures and processes. Also, Piaget overestimated the role of failure in development. He spoke of "reproductive, or functional assimilation" as the natural, spontaneous, propensity to exercise one's new skills. But he assumed that this exercise leads to adaptation by way of gradual corrective processes of "accommodation," processes initiated by increasingly demanding assessments of failure.

One might expect computational concepts to be helpful in clarifying Piaget's remarks about cognitive exploration, because they articulate complexities and transformations within symbolic structures (and so provide intentional, or hermeneutic, explanations rather than physical ones). Indeed, much of what Piaget had to say about equilibration can be better stated in these terms [Boden, 1982]. However, current "learning" programs are mostly concerned with the fine tuning of already structurally adapted computational systems, rather than with their structural adaptation itself. Some of them, too, overemphasize the role of failure [e.g. Winston, 1975] (although others [e.g. Samuel, 1970] use a hill-climbing strategy controlled by some notion of "getting better"). Almost nothing has been done to model systems where the structuring principles themselves change over time in an integrated fashion. This is a key notion in understanding development and creativity, and this was what primarily concerned Piaget.

For example, Winston's program can recognize only a few properties of the input, these being all and only the relevant ones, each of which was initially defined by the programmer [Winston, 1975]. Given these predefined properties, the program can learn from counterexamples as well as from examples, but only if they are presented in an epistemologically suitable (highly constrained) order. A more "Piagetian" self-modifying program is Sussman's HACKER, which can diagnose five classes of mistake and adapt differentially to them, generalizing its adaptive insights so that they can be applied to many problems of the same structural form [Sussman, 1975]. It can even modify successful procedures, for example by removing redundant steps, and by replacing repeated code by subroutines. The structure of HACKER's problem-solving does become more complex and differentiated, and better adapted to the specific constraints of the problem-situation.

exploring only a small number of transformational paths.

Humans can try out many new ways of thinking to see what they will find. They need not have any problem in mind, to which they hope the exploration may be relevant; even if they do, they may be very unclear about how this could be so. Admittedly, where a newly acquired structure is concerned, people often explore its generative potential without making any attempt to transform the structure itself. For instance, children practise grammatical permutations of words when left alone in their cribs [Weir, 1962]. And, as we have seen, children sometimes ignore a failure as an irrelevant anomaly instead of treating it as a spur to adaptation. The experimenters suggested that they need time for "consolidation" of their new theories -- but what is consolidation, and why is it necessary?

I am reminded here of Descartes' fourth rule, "Recapitulate!". At first sight this seems banal in the extreme, qualifying perhaps as mental hygiene but hardly as 25% of "The Right Method of Conducting the Reason". However, it is a fact that recapitulation of an argument -- for example, going over and over a geometrical proof -- may lead one to "see directly" relations which earlier one could only infer by remembering a series of steps. In Descartes' terms, deduction gives way to intuition. And rereading a proof, paper, or program enables one not only to eliminate "howlers," but to find economies, generalizations, and improvements in clarity. Recapitulation, that is, seems to be a method of achieving consolidation.

In general, during the period of consolidation one does things that enable one to represent and improve the structure of the cognitive structure itself. That is, one engages in meta-activities, and/or one follows high-level meta-goals, of various kinds. [cf. Rissland, 1978]. So one eliminates redundant steps; one constructs higher-level representations that economically summarize a number of already available sub-structures; one explores the sorts of transformation allowed by the unfamiliar structure; one classifies the states that can be generated within it; and one compares this structure with others in various ways.

Until one thus understands the general potential of a structure, one may not be motivated to change it, nor able fruitfully to pay attention to counterexamples to it. Kuhn's [1962] account of the activity of normal science, which continues despite theoretical anomalies, seems to fit this description. So does the exploration of any new artistic style before it is superceded by another. Once consolidation has been achieved, other meta-processes may come into play (sic) by which relatively radical changes can be effected. For example, dropping a very basic constraint (such as Euclid's last axiom, or the "stringlike" nature of molecules) may lead to a coherent structure with generative properties very different from its predecessor. Again, combining sub-structures or procedures in different combinations, and with different orders of priority, may be a general method by which the mind can spontaneously generate new structures out of old. The technique of "brainstorming" is based on this principle.

If consolidation is a mapping of the generative geography of one's current structures, a theory of adaptation should explain how one realizes that the terrain has now been reasonably well-mapped, so that judgments about the extent of unexplored territory must influence the evaluation of what is "interesting." This self-tuning (whereby what is of interest today may be boring tomorrow) is analogous to the successive strengthening of evaluation criteria in "track and trail" adaptation. But in the human case the criterion is non-metric and multi-dimensional: one may have explored some aspects of a structure to one's satisfaction -- indeed, satiation -- but not others.

An adaptive system must also be able to judge the "interest" of the results of its (more or less radical) explorations, and of the exploratory paths it decides to follow. Without this control, which determines what the exploring mind will find worth pursuing, exploration would degrade into mere chaotic thrashing about. Literary criticism, criticism of music or the visual arts, critical history of ideas, and scientific discussion all aim to express our intuitions in this regard. A psychological theory of creativity should try to make these insights even more explicit. This problem is not seriously addressed by theories of "adaptation-levels", "discrepancy principles", and the like [e.g. Helson, 1959], which equate degree of interest with degree of novelty (novelty being measured with respect to the subject's current schemas, or competence). Adaptation is not a matter of nobbling the new, but of pursuing the promising.

A study of creativity which has addressed this question is D. B. Lenat's automatic mathematician [Lenat, 1977]. This program uses several hundred heuristics (not just a few transformational rules) to explore the space defined by a hundred primitive concepts. "Exploration" here means asking about certain facets of a given (primitive or constructed) concept. For instance: is it named; is it a generalization or a special case of some other concept; what examples fit the definition of the concept; which operations can operate on it, and which can result in it; are there any similar concepts; and what are some potential theorems involving the concept? One of the facets is "interestingness": Lenat attempts to control the exploration by guiding it into areas likely to be more adaptive than others. For instance, he provides it with the general heuristic that if the union of two sets has a simply expressible property not possessed by either of them, then it is probably worth exploring. Lenat claims to have identified several very general heuristics, but also stresses the need for large numbers of domain-specific, knowledge-based, heuristics (some of which are specializations of the more general ones).

Granted that the heuristics were thought up by Lenat rather than by the program, it is significant (and surprising to many people) that this sort of fruitful exploratory thinking can be formally represented at all. However, the degree of creativity evinced by the program is difficult to assess. Critics [Hanna & Ritchie, 1981] have remarked that Lenat does not list all the concepts regarded by the program as interesting: perhaps a high proportion were mathematically trivial. It is not clear from the publised accounts whether some crucial "discoveries" were made possible only by the use of unacceptably ad hoc heuristics, nor is it easy to draw the line between an acceptably specialized expert heuristic and a disingenuous programming trick. Certainly, many of the heuristics are highly domain-specific, relevant only to set-theory. But it is a prime theoretical claim of Lenat's (and of many other workers in AI) that intelligence depends heavily on expert knowledge, as opposed to very general skills. To the extent that this is

Perhaps simitar considerations concerning creative exploration might illuminate various biological phenomena which, on a neo-Darwinist account of evolution, are very puzzling. These include the facts that the fraction of DNA that does not code for the synthesis of specific proteins increases phylogenetically; that species have evolved remarkably quickly, and that the more complex species have if anything evolved at a greater rate than their predecessors; and that the speed at which a species evolves morphologically seems quite unrelated to the rate at which its individual proteins evolve (so frogs have protein-synthesizing mechanisms of comparable complexity to those of man). Such facts are not explicable in terms of "Random-Generate-and-Test", the mutational strategy favoured by neo-Darwinism. This is because (as was discovered by the early workers in automatic programming), the combinatorics of such a process are horrendous tcf. Arib, 1969D. Switching to a higher-level biological language (£f. "consolidation"[11]), might be effected by random processes of gene duplication and recombination; but this merely reduces the exponent without preventing an exponential explosion.

Instead, some strategy of "Plausible-<3enerate-and-Test" is needed, whereby mutations of a type likely to be adaptive become increasingly probable. The initial heuristics must evolve by random mutation (since there is no suggestion of teleology here), but these survive by natural selection and can eventually enable a form of biological bootstrapping by modifying each other. This is possible because they are embodied as DNA and their "target" for interpretation is itself DNA. That is, they are heuristics recommending certain "copying errors" and preventing others. The sort of transformational processes they influence *are* gene substitution, insertion, deletion, translocation, inversion, recombination, segregation, and transposition.

In a recent, speculative, paper CLenat,198OD, Lenat likens these transformational processes to Production Rules, saying 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.

Whether we are concerned with set-theory or genetic evolution, to explain creative development we need to posit specialist as well as general heuristics. Adaptation requires one (individual or species) to learn a large number of special tricks in terms of which to generate things likely to be interesting, and can be aided by special principles with which to evaluate the interest of anything that turns up. (Natural selection is a very general way of assessing interest *post hoc*). Insofar as this is so, the extent to which one can hope for a *general* theory of creativity is limited. But any such tricks and principles function against a background of creative constraints, in whose terms adaptive control is basically intelligible.

is an ill-defined system? Unpredictability is due in part to the fact mentioned previously, that the mind has a very rich generative potential. We should not expect to find psychological laws reliably predicting the generation of this symbolic structure rather than that one. However, unpredictability is grounded not only in idiosyncracy but also in contingency -- in the richness of the world, and the fact that it goes its own way independently of our designs. A structural-transformational theory of adaptive control need not deny a creative role to contingencies, provided that they can be integrated into the mental structure concerned by way of the criteria of "interestingness" already functioning within the mind.

For instance, a drummer suffering from Tourette's syndrome is able to make his involuntary tics the seed of exciting musical improvisations. Presumably, he notices (classifies as interesting) a novel rhythmic or tonal aspect of the noise produced by his hands during the tic, which he then explores in musically intelligible ways. If the whole process were random, the result would only on chance occasions be judged "exciting" (much as it would be a rare monkey that could type Hamlet). A structural theory can even allow that contingency is sometimes essential to creative intelligence. For example, because of the cognitive structures already active in his mind, Fleming was able to take advantage of the accidental appearance of the penicillium mould -- but he could not have originated it himself. A designer-God might even put a random-generator into creatures' brains, to produce (hardly to "generate") ideas by chance that might be useful to them during their insightful explorations.

A theory of creativity therefore has to be, in part, a theory of adaptive opportunism. Its core, however, must concern exploratory processes that can range over the space defined by background structural constraints, and that can even transform those constraints themselves. These processes are not controlled by pre-defined specific goals, though they may achieve high-level goals such as improved economy, elegance, and clarity. In short, failure is not the spur, and opportunism is not enough.

ooooo000ooooo

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

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

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

BODEN, M. A. (1981) Minds and Mechanisms: Philosophical Psychology and Computational Models (Ithaca, N.Y.: Cornell University Press), chap. 4, "The Case for a Cognitive Biology".

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

HANNA, F. K., & G. D. RITCHIE (1981) AM: A Case Study in A.I. Methodology. Dept. Computer Science Memo, Univ. Kent at Canterbury.

HELSON, H. (1959) "Adaptation-Level Theory", in S. Koch (ed.), Psychology, A Study of a Science, Vol I: Sensory, Perceptual, and Physiological Formulations (New York: McGraw-Hill).

INHELDER, B., & A. KARMILOFF-SMITH (1975) "If You Want to Get Ahead, Get a Theory", Cognition, 3, 195-212.

KARMILOFF-SMITH, A. (1979) "Micro- and Macro-Developmental Changes in Language Acquisition and Other Representational Systems", Cognitive Science, 3, 81-118.

KUHN, T. S. The Structure of Scientific Revolutions (Chicago: Univ. Chicago Press).

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.

PERRY, R. B. (1921) "The Independent Variability of Purpose and Belief", J. Philos., 18, 169-180.

RISSLAND, E. (1978) "Understanding Understanding Mathematics", Cognitive Science, 2, 361-383.

ROSEN, C. (1978) Schoenberg (London: Fontana).

SAMUEL, A.L. (1970) "Some Studies in Machine Learning Using the Game of Checkers. II:- Recent Progress", in F. J. Crosson (ed.), Human and Artificial Intelligence (New York: Appleton-Century-Crofts), pp. 81-116.

SELFRIDGE, O. G. (in press) Track and Trail in Adaptation. Cambridge, Mass.: MIT Press.

SUSSMAN, G. J. (1975) A Computer Model of Skill Acquisition (New York:

TAYLOR, C. (1971) "Interpretation and the Sciences of Man", _Rev. Metaphysics_, _25_, _3-51_.

WEIR, R. (1962) _Language in the Crib_ (The Hague: Mouton).

WHITE, R. W. (1959) "Motivation Reconsidered: The Concept of Competence", _Psychol. Rev._, 66, 297-333.

WINSTON, P. H. (1975) "Learning Structural Descriptions from Examples", in P. H. Winston (ed.), _The Psychology of Computer Vision_ (Cambridge, Mass.: MIT Press).