EXPLAINING PERCEPTION:
AN ASSESSMENT OF CURRENT ECOLOGICAL AND COGNITIVIST APPROACHES

by

Christopher Albert Fields
B.S., Louisiana State University, 1977
M.S., University of Colorado, 1980
M.A., University of Colorado, 1983

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Philosophy
1985

This thesis for the Doctor of Philosophy degree by
Christopher Albert Fields
has been approved for the
Department of
Philosophy
by

Robert Cummins
Robert C. Cummins

John A Fisher
John A. Fisher

5/6/85

Fields, Christopher Albert (Ph.D., Philosophy)
Explaining Perception: An Assessment of Current Ecological and Cognitivist Approaches
Thesis directed by Professor Robert C. Cummins

Ecological realism and cognitivism are the two major current contenders in the field of cognitive perceptual theory. This thesis examines these theories, and the debate between them. It shows that the debate, as it exists in the literature, is inconclusive, primarily because of problems in the current formulations of the two contending theories. The most obvious difficulties in the two theories are removed, leaving reconstructed versions of both. The debate is then re-examined in the context of the reconstructed theories. It is shown that ecological realism is a special case of cognitivism dealing with the detection of properties of objects in the environment by resonant transduction. It is also shown that neither theory, as it stands, can adequately describe changes in the perceptual abilities of even very simple animals.

ACKNOWLEDGMENTS

Rob Cummins has seen this project through from its beginning in 1981. I would like to thank him especially for his support while a visitor here, and while at Chicago Circle. It has been a great pleasure to work with him.

I also thank my committee: John Carnes, John Fisher, Gene Gollin, Steve Leeds, and Gary Stahl. They have cheerfully attended many meetings and criticized many drafts. I appreciate their support.

Fellowship support from the graduate school at CU is gratefully acknowledged.

Thanks go to Eric Dietrich and Katrina Walker for stimulating, and fun, discussions, and to the many other friends who have endured endless rambling about perception models.

I would like to thank my parents for their support and encouragement. Finally, I thank Cindy for her love and support in this and many other things.

CONTENTS

CHAPTER

CHAPTER I

INTRODUCTION

If cognitive psychologists should agree about anything, they should agree about perception. Perceptual processes provide the informational input on which all cognitive processes act. Unless this input can be characterized, theoretical models of cognitive processes cannot be tested, and cognition cannot be understood. For example, one can ask what information about the world the eye, and the rest of the visual system, provides to the rest of the brain. Unless this question can be answered, we have no way of saying what the other processes occurring in the brain that use visual information have to work with. Unless we can specify this, however, we cannot say what those processes do. Understanding perception is therefore essential to understanding cognition in general.

What does a theory have to provide in order to explain perception? The broadly phenomenological answer is that explaining perception is explaining what it is like to perceive, i.e. to experience the world by means of perception (e.g. Merleau-Ponty, 1962; Nagel, 1974). The broadly mechanistic answer is that explaining perception is explaining how the physical, biological, and psychological processes that constitute perception work (e.g. Lindsay and Norman, 1977; Marr, 1981). This dissertation is concerned solely with mechanistic accounts of perception.

Mechanistic investigations of perceptual processes currently procede at several different levels of description [1]. On the one hand, one can investigate the biophysical processes involved in the interaction of sensory receptors with perceptual media such as the electromagnetic field, or the biochemical processes involved in the transmission of signals across synapses. In contrast to these "low-level" approaches to perception, one can consider perception more abstractly as a process in which information is transferred from an environment to an organism, as shown schematically in fig. I-1. This approach to perception, i.e. that explanations are to be couched in terms of the flow of information, is the one taken by cognitive science. Cognitive science, therefore, attempts to explain how an organism can "find out" the things about its environment that its behavior suggests it routinely does find out by means of perception. The level of description adopted by cognitive science thus overlaps that of much of the philosophical tradition, in particular that of traditional epistemology (e.g. Rock, 1975, pp. 12-18; Pylyshyn, 1984, pp. 1-6; Rorty, 1979, Ch. 5).

The epistemological characterization of perception is not uncontroversial. Behaviorists, in particular, reject the epistemological characterization, claiming that behavior can be explained by appeal to laws directly linking stimuli specified in the restricted vocabulary of physics to responses specified in the same vocabulary [2]. Post-behaviorist cognitive science has, by and large, rejected this claim. There is, however, deep disagreement among cognitive scientists as to what is wrong with it. There are two fundamentally

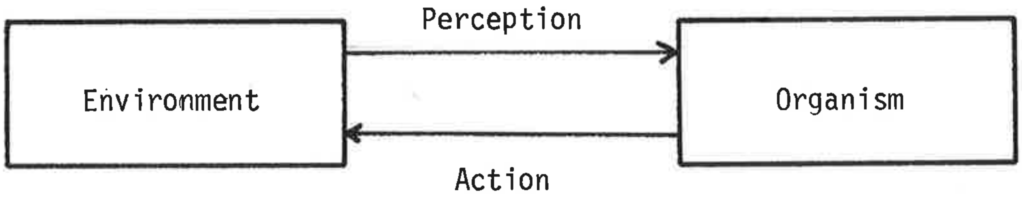

Fig. I-1: Perception can be viewed as a process in which information flows from an environment to an organism (Dretske, 1981). Similarly, action can be viewed as a process in which information flows from an organism to an environment (Fields, 1983a).

opposed views. The dominant view - the information-processing view - is that stimuli and responses should be specified in the restricted vocabulary of physics, but that the s-r relation is generally not lawlike at the informational level of description. Pylyshyn (1984), for example, claims that descriptions of "[perceptual] transducer inputs must be stated in the language of physics" (p. 165), but at the same time, argues that "if we attempt to describe human behavior in terms of the physical properties of the environment ... human behavior is essentially random" (p. 12). On Pylyshyn's "cognitivist" view, the s-r relation must, instead, be explained by appeal to inferential processing of the information contained in the stimulus by the subject (Fodor, 1975; 1980; Fodor and Pylyshyn, 1981; Marr, 1981; Rock, 1983; Pylyshyn, 1984). In particular, cognitivists claim that organisms must infer the significance of what they see, i.e. that they must infer the contents of perceptual beliefs, from percepts together with stored beliefs or knowledge (Fodor and Pylyshyn, 1981).

In contrast to the received view, the ecological realist view is that stimuli and responses cannot be specified in the restricted vocabulary of physics, but must be specified in the vocabulary of ecology, i.e. in a vocabulary containing such terms as 'predator', 'prey', 'shelter', 'conspecific', etc. When this is done, ecological realists claim, the s-r relations are lawlike at the informational level of description; for example, "a core claim of the ecological approach is that an organimsm ... and its environment ... are bound together as a synergistic system by laws

(Turvey et al, 1981, p. 254, their italics). Ecological realists therefore deny that organisms inferentially process the information contained in stimuli (Gibson, 1950; 1966; 1979; Michaels and Carello, 1981; Turvey et al, 1981; Reed, 1983). In particular, they claim that the perception of the affordances of the environment is direct: "Gibson rejected the idea that organisms have to infer the properties of their environments ... from the properties of other, putatively more primary, objects" (Turvey et al, 1981, p. 240) [3].

The central debate in current cognitive perceptual theory is that between proponents of ecological realism and of information processing. Since about 1980, however, a second debate has arisen within the information processing camp. This second debate concerns whether the information processing in question is inferential in the narrow sense, i.e. in the sense of being defined over stored propositional representations. Traditional cognitivists, or "propositionalists" hold that cognitive information processing is inferential in the narrow sense, and that, therefore, all information is represented in an explicit propositional form, e.g. as sentences in a "language of thought" (Fodor, 1975; 1980; 1983; Pylyshyn, 1980; 1984) [4]. Members of the "new connectionist" camp, however, deny this, claiming that information can be represented without being stored in propositional databases, and that, therefore, information processing need not be inferential in the narrow sense (McClelland and Rumelhardt, 1981; Feldman and Ballard, 1982; Anderson, 1983; Anderson and Pirolli, 1984). For example, McClelland and Rumelhardt (1981) claim that their model of letter perception "accounts for apparently rule-governed performance without any actual rules" (p. 375).

The present discussion focusses on the debate between traditional cognitivists (hereafter, "cognitivists") and ecological realists. Connectionists have not, thus far, entered into this debate. The secondary cognitivist-connectionist debate will not be considered here.

At first glance, ecological realists and cognitivists appear to disagree about almost everything. Cognitivists claim that the acquisition of perceptual knowledge, or of perceptual belief, is inferentially mediated; ecological realists claim that it is direct, i.e. unmediated (e.g. Gibson, 1979; Ullman, 1980). Ecological realists claim that there are "ecological laws" correlating directly perceivable ecological properties of the environment with actions on the part of the perceiver; cognitivists deny the existence of such laws (e.g. Turvey et al, 1981; Fodor and Pylyshyn, 1981). Cognitivists view the brain as a device for representing and processing information, i.e. as a computer. Pylyshyn (1984) is quite explicit on this point: "I want to maintain ... that computation is a literal model of mental activity" (p. 43). Ecological realists view the brain as a proper part of a "perceptual system" that "resonates" to information present in the environment. This resonance is direct; it does not involve computations or inferences: "resonance is to higher-order variables of stimulation, to information ... a perceptual system directly registers these variables; it does not calculate them" (Michaels and Carello, 1981, p. 66).

These disagreements about the facts of perception give rise to methodological differences as well. Ecological realists claim that perception can only be understood by taking ecological niches to be the proper psychological units of analysis (Michaels and Carello, 1981); cognitivists, in contrast, seek to understand the capacities of organisms in abstraction from any specific environmental setting (Fodor and Pylyshyn, 1981). Consequently, ecological realists reject the standard interpretations of most traditional laboratory procedures in principle. In particular, they reject the standard interpretation of experimentally-induced illusions, i.e. that they represent failures of inference (Turvey et al, 1981). Cognitivists, on the other hand, employ illusion studies as a basic experimental paradigm for the study of perceptual inference (Gregory, 1970; Fodor and Pylyshyn, 1981; Marr, 1981; Rock, 1983).

The ecological realist-cognitivist debate is, to all appearances, taken by most of its participants to be a debate to the death for the theories concerned. Fodor and Pylyshyn (1981), for example, characterize ecological realism as "fundamentally wrong" (p. 195). Pylyshyn (1984) similarly argues that "the attempt to explain perception by linking percepts directly to the perceived properties of the world (as was Gibson's goal) clearly fails" (p. 141). On the other side, Turvey et al (1981) claim that "the ecological approach is a scientifically acceptable approach to cognition and what passes as the Establishment's approach, the one championed by Fodor and Pylyshyn, is not" (p. 238). Michaels and Carello (1981) are almost as blunt: "whether ecological *theory* can stand alongside and peacefully coexist with traditional theory is quite another matter. In our view, such a marriage - or even a detente - would not work" (p. 164, their italics).

These claims are quite strong; if they are to be substantiated, strong arguments are required. When the debate is viewed "from the outside," however, the suspicion that the arguments of either side do not contact the opposing position is almost inescapable. One suspects that there are more axes being ground than issues being addressed by the debate. There are several reasons for this. First, neither side appears to make a serious attempt to appreciate, let alone employ, the terminology of the other. The statement of Pylyshyn quoted above, for example, makes no note of the fact that in ecological realism, nothing corresponding to a percept exists. Second, the arguments on both sides tend often to degenerate into ridicule. Michaels and Carello, for example, parody the cognitivist view as a "ghost-in-the-machine" theory in their discussion of motor control, even though they offer no counter to standard arguments against this interpretation (e.g. Dennett, 1978a; Cummins, 1983) [5]. Third, both sides, for the most part, confine themselves to a fairly limited set of paradigmatic examples, and neither side makes a serious attempt to address directly the issues raised by the specific examples employed, and taken to be paradigmatic, by the opposition. Fodor and Pylyshyn, for example, emphasize that cognitivism is concerned with understanding belief fixation, and do not consider the cases of perceptual coordination of behavior in lower animals that are the stock in trade of ecological realism. Turvey et al, in their reply to Fodor and Pylyshyn, emphasize that ecological realism is concerned with the perceptual coordination of behavior, and brush off the concern with belief fixation with the remark that such cases "do not represent a coherent set of examples" (p. 240). This is especially disturbing in view of the fact that one of the central issues of the debate, at least from the cognitivist perspective, is the question of whether perceptual episodes leading to the fixation of a particular belief form a natural kind for psychology.

The primary purpose of this dissertation is to get to the bottom of the debate between cognitivism and ecological realism. I will attempt to show that the arguments that have been deployed, by and large, do not achieve their goals, but that they do show that a re-evaluation and reformulation of both theories is required. I will argue that the debate between the reformulated theories concerns two main issues: the role of appeals to stored information in psychological explanations, and the role of appeals to inference in psychological explanations. Using data from both lower organisms and humans, I will then show that ecological realism is not an alternative to cognitivism, but rather a special case of cognitivism.

The discussion that follows is divided into three parts. The first part considers the ecological realist-cognitivist debate as it has appeared in the literature (Ch. II), and then clarifies the explananda of perceptual theories (Ch. III). It argues that the reasons for suspicion outlined above indeed warrant a re-examination of the theories involved in their own terms.

The second part comprises a critical analysis and reconstruction of ecological realism (Ch. IV), and of cognitivism (Ch. V). In Ch. IV, I argue that ecological realism must abandon the subsumptive explanatory strategy in favor of an analytic strategy if it is to compete with cognitivism. I then show that a consideration of the central concept of resonance demonstrates that ecological realism not only invokes processes that make essential use of stored information, but also invokes processes that can be described as inferential.

Ch. V examines the claims of cognitivism that cognition is computation, and that computer programs provide explanatory models of cognitive processes (Pylyshyn, 1980; 1984). I clarify the cognitivist notion of the functional (computational) architecture underlying a cognitive process, and examine Pylyshyn's (1980; 1984) claim that empirical data can provide a principled distinction between a cognitive process and its architectural underpinnings. I then show that this claim is true only relative to an antecedent assumption of explanatory priority for a certain level of description (in Pylyshyn's case, that corresponding to propositional attitudes). The chapter concludes with a discussion of the role of appeals to inference in cognitivist models.

The third part, Ch. VI, reconsiders the cognitivist-ecological realist debate in terms of two questions: the role of stored information, and the role of inference. I argue that cognitivism and ecological realism are equivalent theories of perception in a

wide variety of lower animals, but that ecological realism can only describe the cognitively impenetrable components of human perception. Ecological realism is, therefore, a special case of cognitivism. Several problems that neither theory can currently answer are raised.

NOTES - CHAPTER I.

1. The relation between investigations at different levels of description has been the subject of considerable debate. A "level of description" can be regarded as a theoretical vocabulary used to describe phenomena at a particular physical scale, and at a particular degree of abstraction. This dissertation assumes a non-reductive account of the relation between investigations at different descriptive levels; that is, it assumes that the explanatory power of accounts at different levels is generally different. Non-reductive accounts of the relations between sciences are explicated and defended in detail by Fodor (1974; 1975; 1978), Boyd (1980), and Cummins (1983).

2. For example, Skinner (1953) praises Pavlov for designing experiments in which stimulus and response "could be easily described in physical terms" (p. 53).

3. What this claim entails is not completely clear. Ecological realists appear to reject the very idea of inference or information processing. It will be shown in Ch. IV, however, that they must accept the information processing involved in extracting signals from noisy media. The arguments that ecological realists direct against cognitivism (see Ch. II) challenge only the claim that organisms must infer what objects afford them, i.e. what useful or detrimental properties objects have; this at least suggests that this is the only form of inference that they object to.

4. The claim that information is represented in an explicit propositional form does not limit the format of the representation. For example, the information that snow is white could be represented declaratively as a statement, e.g. 'snow is white', or procedurally as a production, e.g. 'if 'snow', return 'white'' (McDermott, 1976). All that the claim requires is that information be represented by a proposition or set of propositions in some data structure, and that the rules of inference applying to the proposition apply to it as a proposition.

5. Cognitivism is, in fact, widely advertised as the only theory able to avoid both reductionism and dualism (Fodor, 1975; Cummins, 1982a). Ecological realists have consistently avoided countering these arguments.

CHAPTER II

THE DEBATE

The debate between cognitivists and ecological realists flared up in earnest with the publication of Gibson's The Ecological Approach to Visual Perception in 1979, in which he attacked cognitivism with such claims as, "Not even the current theory that the inputs of the sensory channels are subject to 'cognitive processing' will do ... the approach should be abandoned" (p. 238) and "The ecological theory of direct perception ... implies a new theory of cognition in general" (p. 263). This challenge could not be ignored; cognitivists immediately replied with heavy-handed critiques of ecological realism. The ecological realists, not to be outdone, responded in kind.

This chapter presents the positions of cognitivism and ecological realism as they are to be found in the literature, and the main moves in the debate between them. I argue that neither side in the debate has, thus far, achieved its objective of clearly showing the opposing theory to be false or misguided. While both sides present arguments sufficient to call the other side into question, none of the arguments are conclusive. It is my contention that the debate must be reformulated in terms of reconstructed versions of both ecological realism and cognitivism before serious progress toward its objective can be made. Such a reconstruction will be attempted in Ch. IV-VI. First, however, the debate must be examined on its own terms, for it is the evident failings of the debate as it stands that justify the reconstruction.

Cognitivism

The received view of perception, at least among cognitive scientists, is that perception involves information processing. This claim has two parts. First, it is claimed that somewhere inside the skin of perceiving organisms there are mechanisms that process the information about the world that is obtained from perception. Second, it is claimed that the only way to explain perception, and the perceptual coordination of behavior, is to theorize explicitly about these mechanisms. Ecological realists constitute the only major group in cognitive science that denies these claims.

The existence of a received view does not, however, imply anything resembling unanimity about details among its adherents. There are two main approaches to the problem of characterizing the information processing that is claimed to take place, the "bottom-up" neural modeling approach and the "top-down" computational approach.

The neural modeling approach of, e.g. Kandel (1976; 1979) and Levine (1983) involves the construction of detailed models of neural-system function that are based primarily on neuroanatomical, neurophysiological, and biochemical data. In this approach, individual neurons and their connections are characterized as realistically as possible. Computer simulations are used to model the firing patterns that arise in neural networks when particular component neurons are stimulated. The goal of this approach is to develop anatomically realistic computer models of neural networks that duplicate the perceptual and cognitive capacities of organisms. Major successes of this approach to date include models of memory storage and accessing (Levine, 1983), habituation (Kandel, 1979), and conditioning (Hawkins and Kandel, 1984). Patricia Churchland is a major philosophical exponent of this approach (Churchland, 1980a; 1980b; 1983).

The second approach is computational. Computationalists take as their starting point functional characterizations of informational processes such as inference, memory retrieval, feature recognition, figure disambiguation, planning, language understanding and production, etc. (rewiewed by Lindsay and Norman, 1977). The goal of computationalism is to analyze these cognitive capacities into computational capacities, i.e. into capacities to manipulate formal objects according to a certain rule or algorithm (Cummins, 1983). Once a cognitive capacity has been analyzed as a computational capacity, the algorithm describing the capacity can be instantiated in a program run on a computer, allowing the computer to duplicate not only the performance, but also the putative capacities, of the natural system (Pylyshyn, 1984) [1, 2]. Computational psychology is, therefore, closely allied with artificial intelligence (AI) research (e.g. Dennett, 1978b; Pylyshyn, 1984).

The distinction between the neural modelling and computational approaches is, however, somewhat arbitrary. Many studies, notably those of Marr (1981; see also Marr and Poggio, 1979; Marr and Ullman, 1981) and of Arbib (1972; see also Arbib, 1981; Lara and Arbib, 1982), employ techniques from both approaches. As Marr (1981) points out, the goal of information-processing theories is eventually to understand cognitive processes at the hardware (i.e. implementation), algorithmic (i.e. realization), and computational levels simultaneously [3].

Of these approaches, computationalism claims by far the greatest current allegience among cognitive scientists. While many computationalists have considerable sympathy for the neural modeling approach, the difficulty of its task renders it unable, at least at present, to provide detailed descriptions of the information-processing capabilities of organisms much higher than Aplysia (a marine snail, discussed in detail by Kandel, 1979). In higher organisms, the brain is simply too complex for current computer-modeling abilities; the connection between the capacities that can be modeled neurally and overt behavior are, therefore, not understood at the neuronal level of description [4].

There are currently two main computationalist approaches to the modeling of cognitive capacities, cognitivism and connectionism, which are distinguished primarily by the architectures of the AI programs employed. Cognitivist architectures (Newell and Simon, 1972; Fodor, 1980; Pylyshyn, 1980; 1984) are constructed using propositional data structures that explicitly represent goals, knowledge, beliefs, plans, etc., while connectionist architectures (Mc-

Clelland and Rumelhardt, 1981; Feldman and Ballard, 1982; Anderson, 1983) can be viewed as nets of interconnected neuron-like processing elements. Connectionism differs from the neural-modeling approach in that it is, like cognitivism, top-down; no attempt is made to model individual neurons. Rather, connectionists attempt to model antecedently-specified information-processing functions on connectionist architectures. While the latter systems store data, they do not use explicit (declarative or procedural) propositional databases. Connectionism remains a minority view among computationalists (Dennett, 1984).

Cognitivism is the closest of the information-processing approaches to folk psychology, the everyday "common sense" psychology of beliefs, desires, plans, goals, etc. [5]. Fodor and Pylyshyn (1981), for example, explicitly characterize the cognitivist theory of perception as an attempt to explain the perceptual fixation of belief, a folk-psychological notion. As such, the cognitivist approach to perception is a theory of epistemic perception, i.e. a theory of perceiving _that_ things in the world have particular properties (Dretske, 1969, Ch. II; see also Barwise, 1981). Moreover, as is folk psychology, cognitivism is primarily a theory about humans, and the cognitivist approach to perception is primarily a theory about human perception.

Cognitivism adds to folk psychology the claim that cognition is computation; i.e. that the mind is a computer [6]. An immediate implication of the claim that cognition is computation is the _formality condition_, introduced by Fodor (1980). The formality condition is the claim that the operations of the mind on its input are to be defined in terms of the form in which the input is represented, not in terms of the semantic content of the input. All that mental processes "have access to" are the shapes, so to speak, of the represented inputs. As will be seen below, the formality condition is the basis of much of what is controversial in the cognitivist research program.

The formality condition is quite plausible in the case of perceptual systems. Consider vision as an example. The only information about the world to which the retina has access is the information encoded in the shapes of the patches of illumination falling on the receptor cells. It has no independent access to what these patches of illumination mean.

A consequent of the formality condition is the requirement that epistemic perception be inferential. Unless information about the non-formal properties of objects is specified (exactly) by the formal properties of perceptual input, such non-formal properties will have to be inferred from the available information, i.e. from the formal properties of the input together with additional, stored information (Fodor and Pylyshyn, 1981). In such a case, the information available from perception is "impoverished," i.e. insufficient to specify the meaning of what is perceived [7].

The most obvious examples of impoverished input involve language understanding. The token, 'the cat is on the mat', for example, conveys nothing about cats or mats to someone who does not know how to read English, and presumably would convey nothing at all to someone completely unfamiliar with the practice of writing. The cognitivist analysis of this situation is that the signal corresponding to the token is impoverished; it does not specify the meaning, i.e. the semantic content, of the token. The meaning can be inferred, however, by a perceiver who knows how to read English.

Cognitivists claim that signals are impoverished in general, i.e. that all epistemic perception requires inference (Lindsay and Norman, 1977; Fodor and Pylyshyn, 1981). The additional information used in such inferences must be stored by the perceiver in some sort of memory. The additional information must, moreover, either be learned (as, e.g. the meaning of 'cat' is) or be innate (as, e.g. the depth grammar of English presumably is; see Fodor, 1975).

In summary, cognitivism proposes that perceivers can be described by the model system shown in fig. II-1. Environmental signals are detected by a transducer, which is assumed to be informationally-encapsulated, i.e. which is assumed to receive no inputs from the rest of the system (Fodor, 1983) [8]. The output of the transducer is an encoded signal, a translation of the input into a form that can be used by the remainder of the system. This output is assumed to carry no more information than the environmental signal itself; it may, however, contain less.

The encoded output of the transducer serves as input to an inferential system that has access to information stored in memory. This system resolves the signal from background noise, and infers from it whatever information it can. Its output is used by whatever general cognitive processes the organism has available, for the

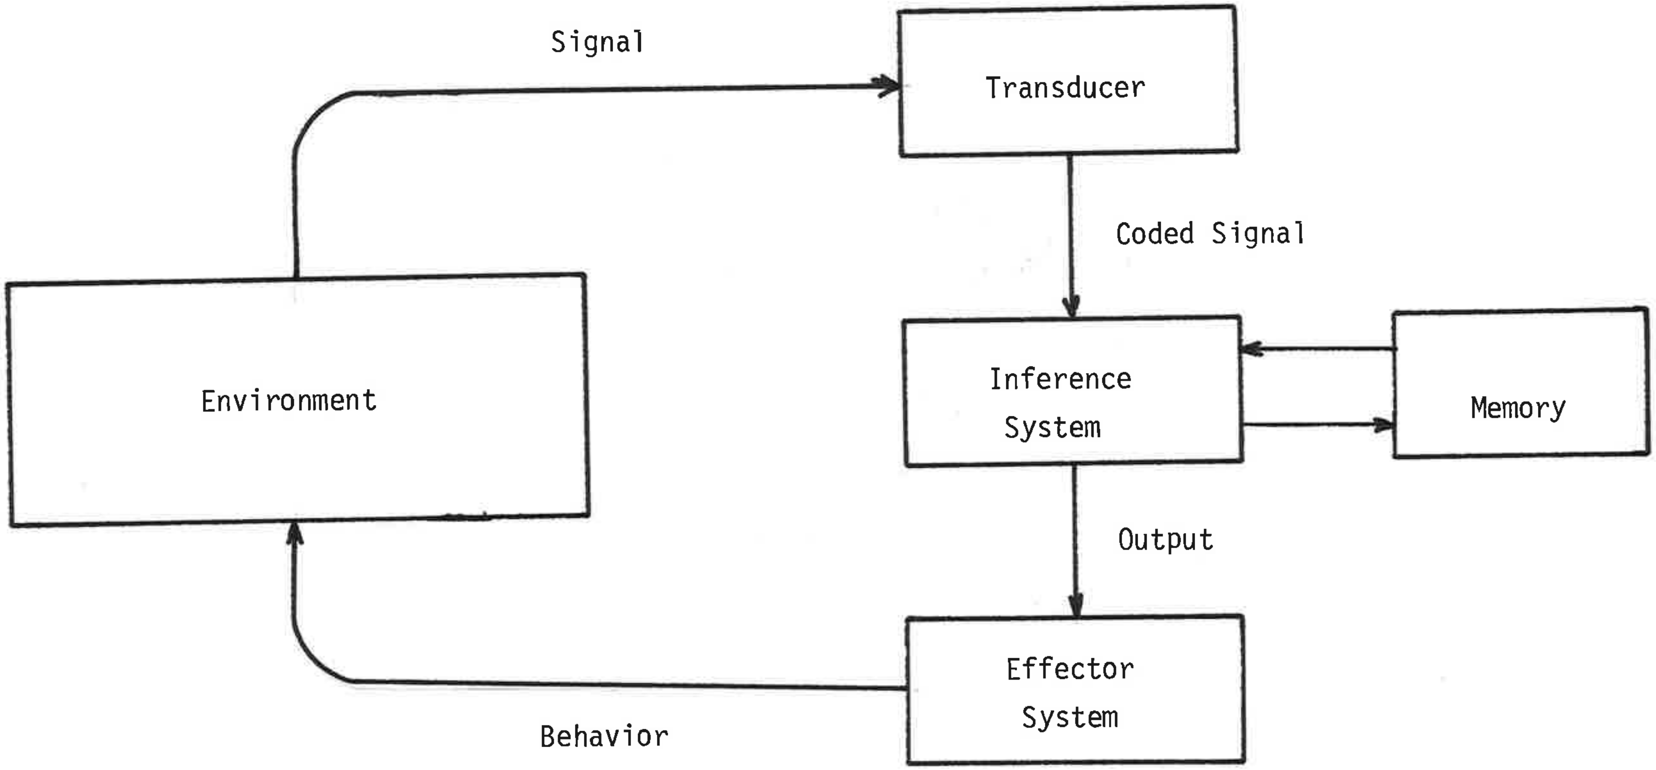

Fig. II-1: Schematic showing the major components of a typical cognitivist model. Only the inference system has access to information stored in memory. Input and output transducers are encapsulated in Fodor's (1983) sense.

coordination of behavior, problem solving, learning, etc. Because the way that the output of a perceptual system is used by general cognitive processes depends not only on the output itself, but also on goals, information stored in memory, etc., the relation between perception and behavior is not, in general, lawlike; hence Pylyshyn's (1984) claim that the perception-behavior relation is "essentially random" (p. 12; see also Ch. III).

As an example, consider a person who sees the North Star (Fodor and Pylyshyn, 1981). The light from the North Star itself does not specify, at least by any of its properties detectable by the unaided human visual system, its size, distance or spectral type, let alone the fact that the direction from the viewer to the apparent position of the star is north. The proximal stimulus, i.e. the light from the star, is therefore impoverished when it comes to these properties. A person can only see the North Star as the North Star, i.e. as a direction indicator, or as a distant, large, hot gaseous body if s/he has certain beliefs, e.g. the belief that the dot of light located at such-and-such a position in the sky is the _North Star_, and if s/he can use that belief as a premise in an inference. People can see the North Star as different things if they have different beliefs at their disposal, or if they use their beliefs differently.

Fodor (1983) argues that, while some inferences, such as the inference that a dot of light is the North Star, are performed by "central" cognitive operations that draw on a large general knowledge base, many important inferential processes are informationally encapsulated. Parsing, for example, is claimed to be unaffected by general knowledge; the only information available to the parser is information about grammatical structures (Fodor, 1983). The claim that a process is inferential does not, therefore, entail that it has access to general knowledge, or even that it has access to learned knowledge. It only entails that it has access to some knowledge base or other, which may, in fact, be unique to the process in question.

The cognitivist model shown in fig. II-1, and its more detailed extensions (e.g. Fodor, 1983; Pylyshyn, 1984), provide natural explanations for a wide range of experimental data on both human and animal perception (e.g. Lindsay and Norman, 1977; Rock, 1983; Fodor, 1983; Pylyshyn, 1984). The basic mechanisms involved in such processes as the resolution of signals from noise and the initial extraction of useful information (e.g. shape recognition) are well understood both mathematically (Grossberg, 1980; Levine, 1983) and in terms of neural architecture (Hubel and Weisel, 1977; Ottoson, 1983; Sterling, 1983; Gilbert, 1983). The computer model of Marr (1981) incorporates much of the available information about neural architecture into a system capable of duplicating significant mammalian perceptual discrimination abilities. Cognitivism therefore constitutes an established theory, and is a powerful opponent for any theoretical challenger.

Ecological Realism

Ecological Realism (ER) was developed by J. J. Gibson between 1950 and 1980 (Gibson, 1950; 1966; 1979). Current presentations of the theory can be found in Michaels and Carello (1981) and in Turvey et al (1981). ER is the best developed, and most serious attempt to incorporate into post-behaviorist cognitive perceptual psychology the idea that psychological research must be "ecologically oriented," i.e. sensitive to the details of behavior in natural settings, in order to be non-trivial (Gibbs, 1979 provides a general review of the "ecological orientation"; see also Bronfenbrenner, 1979).

ER is motivated by two primary concerns. First, it is an attempt to develop a theory of perception that is uniformly applicable across the phylogenetic spectrum. Turvey et al (1981), for example, employ examples of perception by plants, snails, wasps, and sharks as their primary illustrations of the explanatory capabilities of the theory. Michaels and Carello (1981) concentrate on human perception, but throughout refer to the perceiver as "the animal."

The second primary motivation is the desire to "put psychology back on the track of seeking lawful relations," a path proponents of ER claim information processing theories, including cognitivism, have abandoned (Turvey et al, 1981, p. 237). ER is prima facie a subsumptive theory, i.e. a theory that attempts to explain events by subsumption under causal laws (Nagel, 1961; Cummins, 1983). The goal of ER is to find "ecological laws" that relate the <u>activities of organisms</u>, as opposed to their percepts or beliefs, to relevant features of their environments (Gibson, 1979;

Turvey et al, 1981).

These two motivational factors lead ER to reject the mentalism inherent in folk-psychological notions such as belief and desire, and to reject forms of explanation that employ such notions. Ecological realists argue that it is completely implausible to assume that snails, fish, birds, and the like must infer the properties of their environments from impoverished stimuli together with stored knowledge or beliefs, or that they decide what to do by considering what they perceive in light of stored desires, goals, and plans (Gibson, 1979; Michaels and Carello, 1981; Turvey et al, 1981). Such concepts therefore cannot be invoked by a psychology that applies to all organisms [9]. Moreover, if the explanatory goal of perceptual psychology is to find laws relating perception to action, descriptions of the animal in question in terms of mentalistic predicates are as irrelevant to psychology as they are to, e.g. population biology. If the behavior of an organism is determined by its perceptions, and if it is to be explained by appeal to a perception-action law, appeals to internal states of the organism, such as its mental states, add nothing to (subsumptive) explanations of its behavior. Furthermore, ecological realists claim, such mentalism adds nothing to explanations of its perceptions.

The rejection of folk-psychological mentalism entails that ER cannot view itself as explaining the perceptual fixation of belief. Belief fixation is not an action, i.e. a behavior that occurs in the external world. ER instead focusses on the perceptual guidance of such everyday activities as walking, running, flying, swimming, avoiding obstacles and other dangers, locating and catching food, etc. (Turvey et al, 1981). In all of these activities, the acting animal's orientation and movements are coordinated and guided by perception. It is this relation between what an animal sees, hears, smells, etc. and what it does that the ecological laws postulated by ER are supposed to explain.

The rejection of belief fixation as the explanandum of perceptual psychology does not, however, entail the claim that perception is non-epistemic. On the contrary, ecological realists emphasize that seeing, for example, is always seeing-that. Gibson (1979), for example, claims that perception is "an experiencing of things, rather than a having of experiences. It involves awareness-of instead of just awareness" (p. 239), while Michaels and Carello (1981) characterize ER as describing the "biology of knowing" (p. 62).

"Knowing" in ER is not, however, regarded as being in a certain mental, or even brain, state. Ecological realists reject the equation of "knowing" with "having knowledge;" they rather view "knowing" as having an ability (Michaels and Carello, 1981, p. 62 ff). Perceptual knowing is knowing how to perceive, e.g. how to see things in a certain way. That is, it is being able to see things in a certain way. Knowing is, moreover, regarded as an attribute of a whole organism, not of a mind or brain [10].

Specifically, ecological realists claim that organisms see (hear, smell, etc.) not "bare" objects, but rather that objects have certain meaningful properties or "affordances" (Gibson, 1979; Michaels and Carello, 1981; Turvey et al, 1981; Reed, 1983). For example, organisms are claimed to see that other animals afford being eaten, or that caves or clumps of bushes afford shelter, or (to use an example of Michaels and Carello) that bottles afford throwing as well as drinking out of. The affordances of an object for an animal are what the object affords the animal, in terms of survival value, pleasure, pain, etc. Affordances are, in essence, the meanings of objects to the organism in question; they are what is ecologically significant about the object for the organism. Meanings are, again, not mentalistic constructs; they are natural, relational properties of objects, properties that specify the relation of the object to the organism in question. The claim that organisms perceive affordances is thus an extension and application of the notion of intentional perception, i.e. that what is perceived is a meaning, developed by Merleau-Ponty (1962).

ER represents the major current challenge to the cognitivist paradigm in perceptual psychology; indeed, even Fodor and Pylyshyn (1981), the staunchest defenders of cognitivism, allow that proponents of ER constitute a "substantial minority of the cognitive science community" (p. 139). Its proponents view ER not merely as an alternative theory to explain the part of perception that does not involve belief fixation, but rather as a revolutionary new approach to cognition destined to replace information-processing theories altogether.

The central foundational claim of ER is that "perception is a process in an animal-environment system, not in an animal"

(Michaels and Carello, 1981, p. 1). The traditional notion that organisms can be psychologically characterized independently of any particular niche is therefore abandoned. The animal and its associated environment, its niche, are claimed to form a system that can be characterized physically, e.g. in terms of matter and energy flow, biologically, e.g. in terms of biological evolution and predator-prey relations, and psychologically, in terms of perception and action. Neither perception nor action, ecological realists claim, can be defined for a system smaller than that of (organism + niche); in particular, neither can be defined for an organism in abstraction from a particular niche.

This is a very strong claim, and one that directly opposes the conventional assumption, adopted by cognitivists, that an organism can be taken from its natural habitat into the laboratory, characterized, and returned to its habitat with the assumption that generalizations based on its behavior in the lab will be true of its behavior in the wild (Neisser, 1976; Gibbs, 1979). Cognitivists assume, in effect, that the contribution of the environment to perception can be "factored out" in much the same way that the contribution of friction can be factored out of a kinematics problem. In kinematics, such factoring is possible because friction is independent of the nature, and even the magnitude, of the force that causes motion. In the case of perception, cognitivists view the environment and the organism as contributing independent information to the process. Ecological realists deny this, claiming that the information contributed by the environment is information about affordances, and that affordances can only be specified relative to a particular organism. The information contributed by the environment is, ecological realists claim, in effect, about the perceiving organism itself. The assumption that the organism-niche system is the smallest possible unit of analysis for psychology forces psychology to adopt a natural-historic research methodology. This methodology does not prohibit experimental manipulation of the organism or the niche, but it does require that experiments be controlled for all factors that could, in the natural setting, influence the outcome. Such factors are generally not known in advance; considerable research is often required to discover them.

Ecological realists argue for the claim that the organism-niche system is the minimal unit of analysis by appeal to evolutionary considerations (Michaels and Carello, 1981). The argument can be summarized as follows. All organisms have co-evolved with particular niches. The niche evolutionarily constrains the organism, and the organism evolutionarily constrains the niche. This is what Michaels and Carello mean by "organism-environment synergy." Removing the organism from the niche, therefore, removes the very constraint that its behavior patterns, and the cognitive mechanisms supporting them, have evolved to cope with. The laboratory is, therefore, a relatively unconstrained environment, not the highly constrained environment that traditional approaches take it to be. As will be shown below, this argument forms the basis for the ecological realist's proposed solution of the problem of specifying the information contained in signals in perceptual media, the ecological realist analysis of perceptual illusions, and the ecological realist arguments against cognitivism.

The organism-niche system contains, minimally, three components: an organism, a set of propertied objects, and a perceptual medium. The objects in the niche interact with the perceptual medium, imparting to it a certain structure. For example, objects interact with the visible part of the electromagnetic field to produce a structured optic array with characteristics that can, at least in principle, be predicted from knowledge of the characteristics of the objects and of the incident light. Similarly, vibrating objects interact with the surounding medium (air or water) to produce ambient arrays of sound, and objects that release chemicals contribute to the ambient array for olfaction. Organisms interact with perceptual media by means of sensory receptor organs such as eyes, ears, or nostrils. By interacting with a medium, an organism can detect the affordances of the objects that caused the signals.

The principle empirical claim of ER is that, within a given niche, the properties of the ambient perceptual arrays specify the affordances of objects in the niche to the organisms in the niche. Examples of such affordances include things that can be run on (e.g. the ground for humans, thin branches for squirrels, and calm water for water striders), things that can be eaten (e.g. wood for termites, pizza's for humans), things that can be mated with (i.e. appropriate conspecifics), things that are inanimate but dangerous (e.g. fires and visual cliffs), things that are animate and dangerous (e.g. predators), etc. This "specification claim" implies that proximal stimuli are not impoverished, at least when it comes to affordances; if they were, affordances could not be specified by proximal stimuli. The perception of affordances is, moreover, claimed to be direct, i.e. not mediated by images, sense-data, internal representations, or any other "proximal surrogates" of the affording object itself (Turvey et al, 1981, p. 239)[11].

If perception is to be direct, the problem of channel equivocation cannot be solved by inference to a "most likely" solution, as it is in cognitivism. If perception is to be direct, channels must be unequivocal, i.e. each signal must unequivocally specify either a particular affordance for a perceiving organism, or no affordance at all.

Channels are, however, generally equivocal. It is possible, at least in principle, for any perceptual signal, i.e. any modulation of light, sound, chemical composition, etc. in a perceptual medium, to be produced artificially, i.e. by some state of affairs other than the interaction of the medium with whatever usually produces the signal. The possibility of creating a hologram, or a synthesized sound-wave pattern, shows that this is true for the perceptual modalities of sight and audition. This is just the "argument from illusions" used, e.g. by Gregory (1970) to argue that perception requires inference from an ambiguous, and therefore impoverished, proximal stimulus to a perceptual belief with unequivocal contents.

Ecological realists argue, however, that the in principle possibility of channel equivocation is of no direct concern to empirical psychology. What concerns ER is the possibility of *effective* channel equivocation, i.e. the possibility that a channel is equivocal for a particular organism in a particular niche. If signals in perceptual media are to specify affordances for organisms, then *in the organism's niche*, signals cannot be equivocal; i.e. a particular signal can only result, in the niche, from the interaction of the affording object with the medium in question.

The claim that channels are effectively unequivocal in natural settings clearly requires some supporting argument. Ecological realists claim that the problem of specifying affordances by signals in perceptual media is solved evolutionarily, i.e. that the co-evolution of organism and niche allows the organism to be sensitive only to properties of the objects in its niche that are transmissible by the media to which it is sensitive, or that are nomically related, in its niche, to such properties, and forces the organism to be sensitive, i.e. *attuned* to those properties of its niche that are affordances for it. Ecological realism thus depends very heavily on the notions of evolutionary selection and adaptation to support its solution to the affordance specification problem (Turvey et al, 1981; Michaels and Carello, 1981).

Because the ecological realist solution to the specification problem depends on the co-evolution of the particular niche of interest, the ecological laws linking the perception of affordances with particular actions are niche-specific. Ecological laws govern the organism-niche *system*, and cannot be "transferred" to a different organism or to a different niche; in particular, they cannot be expected to hold in the laboratory without the proper controls. The job of the ecological realist is to determine, for each niche of interest, what ecological laws hold in that niche. This is primarily a natural-historic endeavor, although experiments may be used to test proposed ecological laws.

As an example, consider the behavior of the marsh periwinkle, a tidal- zone snail described in detail by Turvey et al (1981). The niche of the periwinkle contains the stems of marsh plants that project above the water surface. The snail exhibits two distinct behaviors with respect to these stems. If the tide is advancing, the snail climbs the nearest stem, apparently to escape being caught on the bottom and washed onto the dry zone of the beach. At all other times, the snail avoids the stems as it does other obstacles to locomotion.

Turvey et al (1981) explain these facts by claiming that there are two ecological laws that apply to the snail-stem-tide system. They are, first, that if the tide is advancing, the snail perceives the stems as affording support, and climbs them, and second, that if the tide is not advancing, the snail perceives the stems as obstacles to locomotion, and so avoids them.

The snail perceives the stems as having affordances by "resonating" to the "higher-order invariants" in the structure of the optic array that encode the properties of the stems (roughly, that they are tall and rigidly attached to the bottom in the first case, and that they are in the line of motion in the second) that are relevant in the given situation (Michaels and Carello, 1981).

The ecological laws noted above can therefore be reformulated, from the point of view of the snail, as follows: if the tide is advancing, resonate to invariants encoding the property of affording support; if the tide is not advancing, resonate to invariants encoding properties relevant to obstacle avoidance. These laws subsume the behavior of the snail with respect to plant stems in its niche.

Resonance is taken to be a process that does not involve inference. This is at least plausible; when a tuning fork, for example, resonates to a certain frequency component of an impact, it does not appear to do so using any inferences. Similarly, when a radio receiver picks up a signal by resonating to it, it does not use inference to do so (Gibson, 1966; Michaels and Carello, 1981). Indeed, ecological realists often gloss "resonance" as "information pickup" (Gibson, 1979; Michaels and Carello, 1981).

If resonance does not require information processing, it does not require a memory to hold information to be used in information processing. Ecological realists, therefore, claim that "memory does not have a place in the ecological approach to perception" (Michaels and Carello, 1981, p. 76).

The general form of ecological realist explanations is shown in fig. II-2. The output of the resonator is a specification of an affordance that is used by whatever process determines the behavioral response of the organism to perceptions of the affordance in question. This latter process depends not only on the affordance perceived, but also on the state of the organism at the time

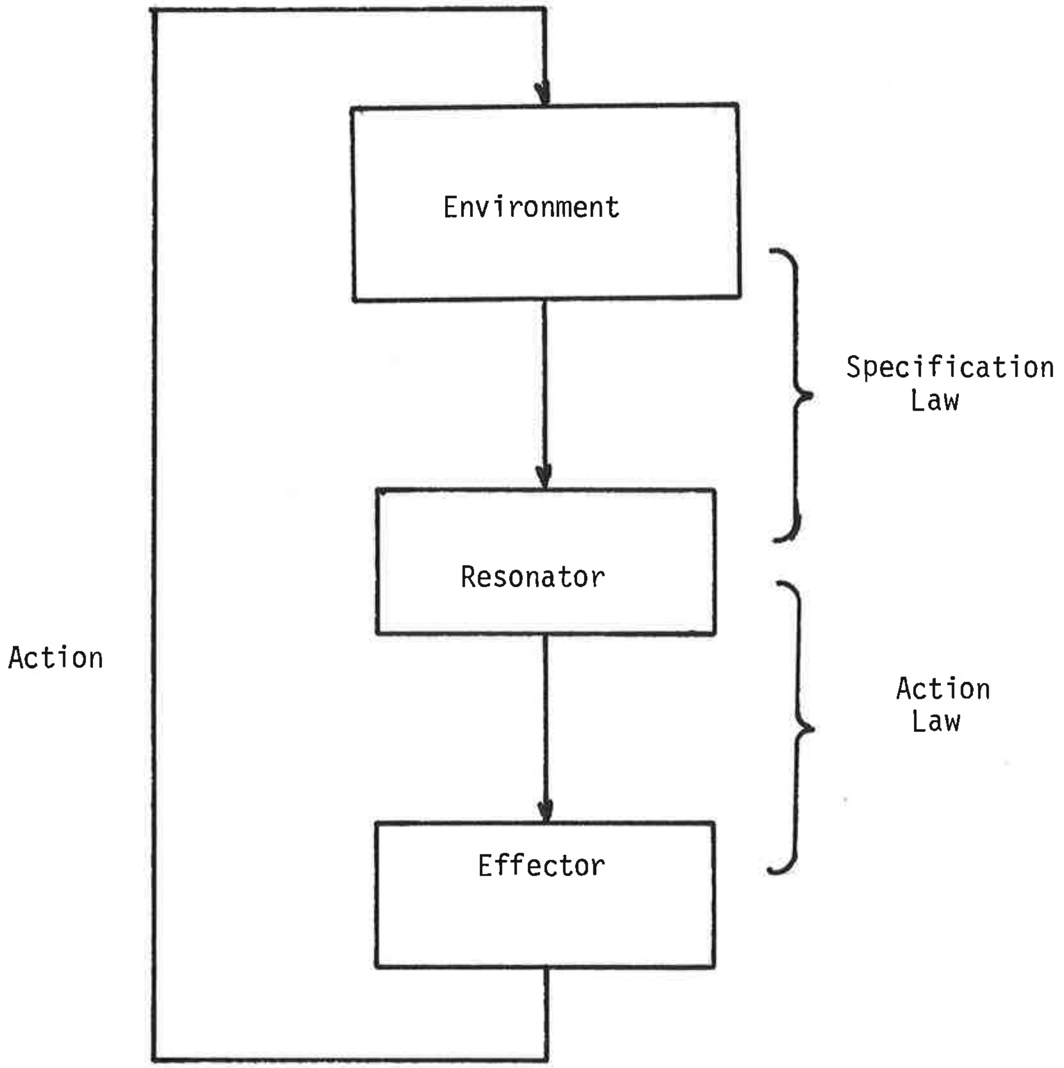

Fig. II-2: Schematic showing the major components of an ecological realist explanation. Affordances of environmental objects are detected by the process of resonance. Specification laws govern the specification of affordances by signals; action laws govern the responses of animals to detected affordances. Together, specification and action laws constitute ecological laws.

(Michaels and Carello, 1981; see also Ch. IV).

Behavioral plasticity is only possible in ER if the abilities or inabilities of organisms to resonate to specific properties can change with time and circumstances. To pursue the analogy with resonance in a tuning fork, changes in an organism's abilities to resonate are analogous to changes in the length of the fork. The resonant properties of organisms can change due to the action of either evolutionary processes or learning. Learning is viewed, however, not as the assimilation and storage of information, but as a change in the ability to resonate to encodings of a specific affordance. The results of evolutionary and learning processes are, therefore, taken to be the same. Michaels and Carello (1981) summarize the position of ER as follows: "We should conceive of the role of experience in the same way that we conceive of the role of evolution. Both lead to a new animal that is better able to cope with its environment" (p. 77). How the mechanism of learning could be like the mechanism of evolution is, however, left entirely unclear. This question will be considered in depth in Ch. IV.

At first glance, ER appears both sketchy and oversimplified. Michaels and Carello (1981) attribute this appearance to the youth of ecological realism, and to the concern of ecological realists with the philosophical questions dividing them from cognitivism. They are confident, however, that the ecological realist position will receive support eventually from both neuroscientists and evolutionary biologists, and that psychologists will eventually be convinced as well. The fact that several eminent cognitivists, such as

Neisser (1976) and Shepard (1984), appear to have been convinced lends some credence to this expectation.

Prima Facie Differences

As pointed out in the introduction, cognitivism and ecological realism appear to agree on very little. The disagreements concern both factual and methodological issues.

On the methodological side, cognitivism and ER disagree, first, on what the proper unit of analysis for perceptual theories is. Ecological realists claim that it is the niche-organism system; cognitivists claim that it is the animal. This disagreement leads to a disagreement about experimental methods, the most obvious manifestation of which is the debate about illusion paradigms. Cognitivists interpret perceptual illusions, such as the Muller-Lyer or Necker cube figures (fig. II-3), as supplying evidence for perceptual inference (Fodor and Pylyshyn, 1981; Rock, 1983). For example, humans incorrectly infer, they claim, that the two arrows in the Muller-Lyer figure are of different lengths. Ecological realists, on the other hand, claim that, in each case, the "illusory" figure has properties that cause it to interact with the medium in a certain way, and that, in the proper niche, anything having these properties would indeed have the affordance in question.

As an example, Turvey et al (1981) consider the research of Kalmijn (1971; 1974) on electrical field detection in sharks. A shark of the appropriate type will dig in the sand whenever it

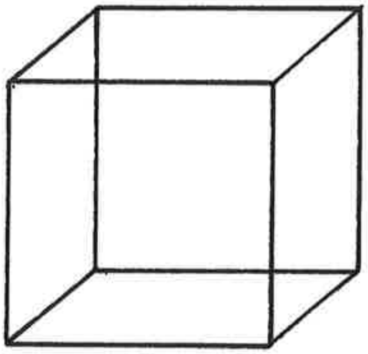

The Necker Cube

The Muller-Lyer Figure

Fig. II-3: Two illustrations known to cause visual illusions (e.g. Gregory, 1970).

detects an electric field of the type caused by the presence of a buried flatfish. If an antenna generating the correct field pattern is buried in the sand, the shark will dig for it. The cognitivist interpretation of this phenomenon is that the shark infers (in this case, incorrectly) that a fish is present when it detects the field. To do so, the shark uses the belief that the correct field being present indicates that a flatfish is present. The ecological realist interpretation is that, in the shark's niche, the field pattern nomically specifies the presence of a fish; the shark has no need to infer it, or to store any beliefs to be used in inferring it.

On the factual side, ecological realism and cognitivism disagree, first, about the specification of the properties of distal objects by signals in perceptual media. Cognitivists claim that signals are impoverished specifications of the properties of distal objects; ecological realists claim that, if the domain under consideration is restricted to the organism's proper niche, the signals are not informationally impoverished.

Cognitivists claim that, since signals are impoverished, organisms must infer the properties of objects from signals, and that knowledge stored in memory must be used to do so. Ecological realists deny that signals are impoverished, and so deny that inference or memory are necessary.

Lastly, ecological realists claim that perceptions and actions are nomically correlated. Cognitivists deny this claim, arguing instead that the correlation between perception and action, and indeed, between perception and belief, is non-nomic.

These disagreements are not minor; they cover almost all of the questions with which perceptual theories are concerned. Perception is, however, a phenomenon that is now understood in considerable detail. The debate between these two theories should, therefore, be swift and decisive. It is not. Both sides have advanced arguments against the other, and, as is typical in philosophical disputes, all of them have some claim to validity. The next two sections present the principal arguments and replies.

Cognitivist Arguments Against ER

There are three major cognitivist arguments against ecological realism: the illusion argument, the constraint argument of Fodor and Pylyshyn (1981), and the explanatory power argument of Ullman (1980) and Marr (1981). This section examines these arguments, and their force against ER.

The illusion argument (e.g. Ullman, 1980; Fodor and Pylyshyn, 1981) is the most common and straightforward cognitivist argument against ER. It attacks the claim that affordances are specified by proximal stimuli by pointing out the existence of laboratory-induced illusions. The fact that illusions are possible, it is argued, shows that proximal stimuli cannot specify distal situations, and in particular, cannot specify affordances.

The ecological realist answer to the illusion argument is now familiar. Ecological realists admit that affordance specification may fail in the laboratory, but claim nonetheless that it succeeds in the niche. This claim is based on the argument that the niche provides constraints on specification that the laboratory setting does not; in particular, the claim that signals are constrained to specify affordances in the niche, but not in the laboratory. The ecological realist response is, in effect, a challenge to cognitivism to demonstrate failures of specification in natural settings. This challenge will be taken up in Ch. IV.

One can also read the illusion argument as a claim by cognitivists that "factoring" the organism-niche system into organism and niche, and concentrating on the relation of the organism to proximal, not distal, stimuli yields an explanatorily superior theory. Ecological realists deny this claim, arguing that cognitivism, and indeed all theories that employ such a factoring, cannot satisfactorily describe the organism-environment relation, and therefore, cannot "do justice to the practical success of an organism's 'everyday' behavior" (Turvey et al, 1981, p. 238).

The second major cognitivist argument against ER is the constraint argument of Fodor and Pylyshyn (1981). This argument has two stages. In the first stage, Fodor and Pylyshyn argue that ecological realists must claim that properties such as the property of being a shoe, or of being Grandma's favorite metal, are directly perceived. In the second stage, the argue that ER can provide no account of how such properties could be directly perceived.

For the first stage of the constraint argument, consider the property G(x): x is Grandma's favorite metal. Fodor and Pylyshyn argue first, that this property is perceptible by humans, i.e. that humans can see that something is G if they know G, and second, that it is completely implausible to assume that G is nomically specified by any property of ambient light. They then argue that, since G(x) is perceptible but not specifiable, perceivers must infer that objects have G(x).

Two options are open to ER in response to this argument. It can either deny that G(x) is perceptible, or else claim that G(x) is, in fact, specifiable. If it denies that G(x) is perceptible, it must, however, still provide some account of how people can come to believe that something is G on the basis of perception. Since, as Turvey et al (1981) admit, beliefs can be fixed either by perception or inference (p. 254), the only possibility is that the belief that something is G is inferred. This, however, will not do; ER cannot countenance inference. ER must, therefore, claim that G(x) is perceptible.

The second stage of the argument is considerably more complex. If G(x) is directly perceptible, Fodor and Pylyshyn claim, then it must either be a transmissible property of light, or it must be nomically correlated with such a property, i.e. it must be quasitransmissible. G is not, however, a projectible property of physics; G cannot, therefore, directly affect the way that an object having G interacts with light. Therefore, G cannot be transmissible. Therefore, G must be quasitransmissible.

If ER admits that G is quasitransmissible, however, it appears that the notion of quasitransmissibility is completely unconstrained, i.e. it appears that it might turn out that <u>any</u>

property is quasitransmissible. ER must, therefore, provide some account of what properties, or what sorts of properties, are quasitransmissible. This account must, however, be non-circular. The detection of quasitransmissible properties is supposed, in ER, to explain affordance perception; therefore, quasitransmissible properties cannot be defined as detectable, on pain of circularity. ER must, therefore, provide an a criterion for quasitransmissibility that does not appeal to notions such as "detectable" or "perceptible." Fodor and Pylyshyn examine four possible ways of doing this drawn from Gibson (1979), and find them all to be inadequate. They conclude that the claim that perception is the direct pickup of information is unconstrained, and therefore that it has no explanatory power.

Fodor and Pylyshyn first consider the possibility that "only the ecological properties of the environment are directly perceived" (p. 144). Ecological realists clearly accept this claim, together with its equivalent, i.e. that all directly perceived properties are ecological. Fodor and Pylyshyn reject this possibility, however, claiming that humans can perceive all manner of properties that are not obviously "ecological," such as G(x), the property of being more than a meter from the Eiffel tower, or the property of needing ironing. They point out, quite correctly, that if 'directly perceptible' and 'ecological' are interdefined, ecology cannot constrain direct perception.

It was pointed out above, however, that ecological realists hold that the specification problem is solved by evolutionary selection. They must, therefore, accept the claim that the directly perceptible properties of objects are those that the organism is selected to detect. Moreover, Michaels and Carello (1981), in their discussion of learning (see above), at least hint strongly that learning is to be viewed as an individualized form of "selection." They claim explicitly, at any rate, that learning, like evolutionary selection, changes what organisms are attuned to. The ecological properties for a given organism are, therefore, the properties that the organism has been selectively attuned to by evolution or learning. This is an interpretation of 'ecological property', however, that Fodor and Pylyshyn, and apparently Gibson himself overlook. It is, however, exactly the interpretation employed by Turvey et al (1981).

If ecological properties are defined in this way, i.e. as properties that the organism is selected, by either evolution or learning, to detect, then it at least looks as if the notion of ecological property has some constraining power. Biological evolution, at any rate, does not act at the atomic scale; therefore such properties as 'has N atoms' are not directly perceptible. In general, biological evolution will only be able to select organisms to detect properties that are either transmissible or quasitransmissible by media in the organism's niche. In the case of learning, however, this is not nearly so clear; in fact, it at least appears that some organisms - humans - can learn to detect, or at least recognize, any property at all.

Whether the fact that humans can recognize many seemingly arbitrary properties constitutes an argument against ER is, however, not completely clear. Ecological realist discussions of learning, to date, are quite sketchy and incomplete. If it does nothing else, however, the first gambit of the constraint argument establishes the burden of proof on ER. If it cannot be shown that the ecological realist notion of learning is constrained in the same way that the notion of evolutionary selection is constrained, i.e. if it cannot be shown that learning can proceed without inference, ER falls prey to the constraint argument.

Fodor and Pylyshyn next consider the possibility that "the projectible properties of ecological optics are directly perceived" (p. 145). The projectible properties of ecological optics are the properties of things that are, according to ER, specified by the structures of ambient perceptual media [12]. Affordances are therefore projectible properties of ecological optics, and are the ones of interest to ER.

Fodor and Pylyshyn consider as candidate affordances such properties as edibility, and argue that edibility is not projectible on any reasonable account of the relation of edibility to the properties of things that determine how they structure light. This account, however, ignores the fact that affordances are, in general, both niche- and organism-specific. 'Edible' does not specify an affordance, while 'edible by O in niche E', for a specified organism O and niche E, might well do so. Without taking this distinction into account, Fodor and Pylyshyn's claim that affordances are not projectible properties carries no weight.

Again, while it is not obvious that looking a certain way to O will always be nomically correlated, in O's niche, with being edible (i.e. it is not obvious that edible-by-O-in-E is quasitransmissible, even in E), it is at least plausible. The survival of most organisms would be very hard to understand if it were not true that the way things look to most organisms in most niches nomically correlates, in the niche in question, with whether they are edible. However, ER must argue for this claim. So far, detailed arguments have not been forthcoming.

Neither of Fodor and Pylyshyn's first two arguments, therefore, does more than shift the burden of proof. There appears to be something to the claim that directly perceived properties are "ecological," but the presentation of the notion of selection by ecological realists is neither detailed nor clear enough to be sure. There appears to be very little to the objection that "ecological optics" projects the wrong properties, but again, it is not clear what properties it does project. Both of these questions will be considered in detail in Ch. IV.

Fodor and Pylyshyn consider and reject two other possible ways of constraining the notion of direct perception: the claim that phenomenological properties are directly perceived and the claim that whatever properties "perceptual systems" respond to are directly perceived. As Fodor and Pylyshyn point out, the latter claim is clearly circular in the absence of a non-psychological way of picking out perceptual systems. Fodor and Pylyshyn's suspicion that ER cannot provide one appears to be correct; ER appears to have no way, at least at present, of relating physiological or anatomical facts about an organism's sensory equipment to its putative ability to resonate to properties specifying affordances, let alone of relating facts about sensory equipment to the affordances specified. Claims about perceptual systems, therefore, appear to have no constraining power.

The former claim, i.e. that phenomenal properties are directly perceived, can provide no constraint in the case of organisms, such as the marsh periwinkle, for which there is no method of deciding what properties, if any, might count as "phenomenal." Ecological realists, moreover, generally do not explicitly consider the question of phenomenal properties (Michaels and Carello, 1981; Turvey et al, 1981). Phenomenal properties appear to be, in this discussion, simply a red herring.

In summary, the strength of the constraint argument of Fodor and Pylyshyn (1981) is unclear. ER appears to be able to provide a sketch of a constraint on direct perception: roughly, an organism can directly perceive property P in its niche if evolution or learning has provided it with a mechanism for doing so. If selection can only provide mechanisms for detecting certain sorts of properties, the class of properties directly perceptible by a given organism is limited. The extent to which this requirement constrains direct perception is unclear, however, without a considerably better explication of what is meant by "selection" than is currently available.

The constraint argument does, however, amount to a fairly powerful burden-of-proof argument. ER is, in effect, challenged to provide an adequate notion of selection, and an adequate account of what properties are quasitransmissible in niches. It must, in fact, show that any property that an organism can come to know about can be learned about without inference. If it cannot do this, it fails.

Elsewhere, Pylyshyn (1984) provides a variation on the second gambit of the constraint argument. He argues that it is a mistake to attempt to explain phenomena by appeal to laws, such as the ecological laws of ER, that are not expressed in terms of the vocabulary of physics: "if we were to ignore the vocabulary of physics, and described the physical events that caused cognitive state changes using some other set of terms, say, *only* terms that refer to *perceived* properties such as Gibson's 'affordances,' we would lose the only coherent way we have of talking about *all* aspects of the physical world" (p. 169, his italics). That sciences can safely talk about their domains in languages other than that of physics is, however, not just the conclusion of ecological realists (e.g. Michaels and Carello, 1981), but also a standard anti-reductionist claim of computationalists such as Fodor (1975). Moreover, it is quite clear that psychology does need to talk about properties such as affordances.

The major concern of Pylyshyn's argument, however, is apparently the problem of inter-niche comparisons. Such comparisons can be made in the language of physics, but it is not clear how they could be made in the language of affordances. It is not clear, in particular, how any but the language of physics can be used to say how, e.g. the same affordance is specified in different niches. The language of affordances, in particular, is prima facie insufficient not only for this job, but for the job of saying how any affordance is specified in any niche. This is, however, less of an argument against ER than a claim that ER must provide more information. This question will also be taken up in Ch. IV.

A third argument has been levelled against ER by Ullman (1980), Marr (1981), and Rock (1983). They argue that, while the claims of ER that organisms directly perceive the affordances of their environments may be true, they are not explanatory, or at any rate, not explanatory enough. Marr (1981), who is perhaps the most sympathetic critic of ecological realism, provides a case in point. He argues that Gibson, and earlier direct realists such as Austin (1962), were misled by the apparent ease with which humans accurately perceive quite complicated features of their environments into thinking that perception was a relatively simple cognitive task, i.e. a task not requiring complicated information processing, and in particular, not requiring inference, for its execution. Therefore, Marr surmises, Gibson rejected the claim that perception involves information processing. Rejecting information processing, Marr claims, is rejecting any hope of providing an account of how perception works, i.e. of how affordances are detected.

Marr's argument carries considerable weight. If it restricts itself to appeals to ecological laws (as Turvey et al do), it is not at all clear that ER can explain why organisms satisfy the ecological laws that they do. The is precisely where cognitivism appeals to inferential mechanisms: cognitivists explain the fact that an organism behaves in a certain way in a certain environment by appeal to the inferential mechanisms that it instantiates. Ecological realists, in contrast, claim that organisms are not inference engines, but rather tunable <u>resonators</u>. Resonance is claimed to be a non-inferential way of extracting information from perceptual media. Marr (1981), as well as Grossberg (1980), however, argue that the only way to extract information from perceptual media is through the use of some form of information processing. This is because, Marr and Grossberg argue, perceptual channels are <u>intrinsically</u> noisy and equivocal. In particular, the large size of the sensory detectors of organisms with respect to the size of possible variations in the media to which they are sensitive suggests that detection always involves the loss of information. Information processing is, therefore, required to replace the information that is lost.

Ecological realists do not appear to be aware of this difficulty, presumably because they have not seriously explored mechanisms of resonance. One can speculate, however, that many ecological realists would find information processing of the sort described by Marr (1981) or Grossberg (1980), to be completely acceptable. Ecological realists reject mentalism: the claim that organisms have beliefs, desires, etc. and that they must infer the properties of the world by supplementing the information in "impoverished" stimuli with such beliefs and desires (Michaels and Carello, 1981). The information processing that is postulated by Marr and

Grossberg does not, however, appear to involve explicit propositional attitudes on the part of perceiving organisms; it rather involves neural networks that perform such tasks as contrast enhancement (e.g. Grossberg, 1980). In particular, it is not at all clear that the information processing postulated by Marr and Grossberg amounts to inferring the affordances of the environment from information about affordances stored in memory.

Ullman (1980) offers a variation on this argument. He points out that many of the neural mechanisms involved in early visual processing are known (e.g. Ottoson, 1983), and that these mechanisms instantiate the information-processing functions proposed by Marr (1981) and by Grossberg (1980). ER, he argues, can offer no explanation of how what these neural mechanisms are doing relates to perception, i.e. of how what these mechanisms are doing relates to "resonance." In other words, Ullman argues that ER has no account of how its theoretical constructs relate to the facts of sensory physiology.

The neurological data on which the Ullman-Marr argument is based appear to be quite sound (Ottoson, 1983; Sterling, 1983; Gilbert, 1983). Ecological realists must, therefore, accept some information processing; in particular, they must accept the information processing functions, such as contrast enhancement, known to be instantiated by neural mechanisms. Whether this information processing counts as "cognitive" processing, and whether it counts as "inference" is, however, debatable. Pylyshyn (1984), for example, claims that it does not (p. 215).

In summary, the common cognitivist arguments against ER have only mixed success. The illusion argument amounts to a challenge, which ER presumably must answer by appeals to empirical data. The constraint argument is, as currently formulated, primarily a burden of proof argument. It shows that ER must supply details of its proposed mechanisms for learning. The Ullman-Marr argument shows that the detection of information involves information processing. However, it is not clear that ER must reject its conclusion.

Ecological Realist Arguments Against Cognitivism

The arguments of ecological realists against cognitivism are mainly concerned with the plausibility of lower organisms using complex information processing to solve everyday problems. There are two principal arguments, a constraint argument, which questions the ability of cognitivism to place any constraints on the notion of internal mental states, and the "recognition regress" argument. Different forms of the latter argument, which questions both the evolutionary plausibility of cognitivist information processing and the plausibility of cognitivist appeals to memory, are presented by Turvey et al (1981) and by Michaels and Carello (1981).

The constraint argument employed by ecological realists (the "ER constraint argument") attacks the cognitivist notion of a mental state by claiming that cognitivism postulates mental states, and inferential processes, willy-nilly, with no way of constraining either the function or the content of its constructs. As such, it is a descendent of traditional arguments against mentalism in psychology (e.g. Skinner, 1953; see also Block, 1980). All ecological realists utilize some version of this argument.

The ER constraint argument embodies two distinct concerns. The first concern is with the attribution of inferential mechanisms and mental states in general; ER claims that such attributions are unnecessary, and therefore, that they cannot be justified in any particular case. The second concern is with constraining particular inferential explanations. ER claims that there is no principled way of choosing the "right" inferential explanation from a set of candidates, i.e. that one such explanation is as good, or bad, as any other.

The ER constraint argument, like the cognitivist constraint argument, is best viewed as a burden of proof argument. It mounts a challenge to cognitivism to prove that it can somehow constrain appeals to mental states, i.e. to prove that models involving mental states are not completely arbitrary.

Pylyshyn (1980; 1984) provides the principal cognitivist reply to this challenge (see also Fodor, 1984; Cummins, 1983; Dennett, 1978b, part II). He claims that an empirical constraint on mental constructs can be developed on the basis of the phenomenon of cognitive penetrability.

Pylyshyn (1980) defines a cognitively penetrable function (CPF) as a cognitive function that can be affected by "such purely cognitive factors as goals, beliefs, inferences, tacit knowledge, and so on" (p. 111). In other words, a CPF is a cognitive function, the output of which can vary as a function of the semantic contents of the beliefs, and other propositional attitudes, of the system that instantiates it. By introducing the notion of a CPF, Pylyshyn distinguishes processes that are cognitively penetrable from those that are not, i.e. that are invariant with respect to changes in beliefs, goals, tacit knowledge, and so forth.

At least some of the processes involved in perception in humans are clearly cognitively penetrable (see Pylyshyn, 1984, for an extensive discussion). For example, our beliefs, knowledge, and goals strongly influence how we see the environment. If a person knows how to read English, tokens resembling 'cat' in relevant ways will be seen as meaning cat. If a person is looking for a place to sit down, a stump, a box, or a car's fender may be seen as having the affordance "sit-onable." If a person is afraid while walking at night, any motion may be seen as embodying a threat. In general, what we know about a thing, or about the uses to which a thing may be put, influences how we see the thing, or at least our perceptual judgements about the thing.

Other aspects of human perception are cognitively impenetrable. For example, a straight stick half immersed in water looks bent, even if we know that it is not. The Necker cube illustration (fig. II-3) looks three-dimensional, even though we know it is planar. The moon looks bigger when it is near the horizon than it does when it is near the zenith, even though we know that it does not change in size. Similarly, the line with inward-pointing arrowheads in the Muller-Lyer figure (fig. III-3) looks longer, even when we convince ourselves by measurement that it is not.

Pylyshyn (1980; 1984) claims that cognitive penetration provides a natural constraint on mental models. If a process is cognitively penetrable, Pylyshyn argues, it must be represented as a process defined over mental representations, i.e. as a process taking as input, and producing as output, beliefs, plans, goals, and the like. Cognitive penetrability is, therefore, proposed as a sufficient condition for representational modeling. As will be shown below, however, it is not necessary. If a process is not cognitively penetrable, it is not necessarily representable as a transducer (Fodor's parsing modules, for example, are not penetrable, but do perform inferences (Fodor, 1983)).

Moreover, Pylyshyn argues, if a process can be represented without reference to mental constructs, then it should be so represented (the "explanatory priority assumption"). These claims provide the basis for Pylyshyn's (1984) rejection of mentalistic explanations of the behavioral capacities of lower animals.

The cognitive penetrability criterion, if coherent, provides at least a partial response to the challenge posed by the ER constraint argument. It provides a way of determining whether mentalistic explanations are needed (by cognitivism's lights, at least), while the accompanying explanatory priority assumption enjoins cognitivists not to use mentalistic explanations where they are not needed. It does not, however, by itself provide a way of selecting a unique mentalistic explanation in every circumstance (Fields, 1983b). Both the coherence of, and possible refinements of the cognitive penetrability criterion will be considered in Ch. V.

The second major argument raised by ecological realists against cognitivism is the recognition regress argument. The major premise of the recognition regress argument is that cognitivism, and indeed, any information processing theory, must claim that "for an organism to perceive property x, it must have the concept of property x" (Turvey at al, 1981, p. 285). The argument takes two paths from this premise. Turvey et al concentrate on the origin of the putative concept of x, while Michaels and Carello concentrate on how such a concept could be used.

Consider the origin problem first. Turvey et al (1981) argue that, if organisms are to infer the affordances of objects from impoverished proximal stimuli, then "evolution must ... provide living things with the conceptual basis needed to make correct inferences .. [which] .. must include concepts that stand for environmental properties and concepts that stand for how those properties structure energy distributions in media" (p. 246). In other words, evolution must provide organisms with two knowledge bases, first, knowledge of what to expect to find in the environment, and second, knowledge of how signals in media relate to properties of things. Turvey et al ask what the origin of these knowledge bases could be, i.e. how evolution could go about providing them. They answer that cognitivism must claim that such knowledge is inferred on the basis of experience, but that, as such inferences would require these very knowledge bases, this answer leads to an infinite regress (or dualism).

If, for example, a marsh periwinkle is able to see stems as climbable because it has the concept of a climb-upable thing, and a concept linking, e.g. tall straight things with climbability, one can ask where it got these concepts. The answer, clearly, is that it either learned them, or was born with them. Turvey et al use the familiar Fodor argument (Fodor, 1975) to show that the periwinkle could not learn about climb-upability. Learning is assumed to be by generalization over instances. The periwinkle could only generalize over instances of climbable things, however, if it could group those things together under a concept. If the concept used in the grouping is "climb-upable," one has an infinite regress. If the objects in question are grouped under a different concept, how does the periwinkle infer that things grouped under the concept in question are climbable? Such an inference also requires a concept of climbability, and so leads to an infinite regress as well. The marsh periwinkle's concept of climbability, like a human's knowledge of language (Fodor, 1975), must therefore be innate. Indeed, the argument, if sound, shows that all concepts must be innately representable, even in humans. This result is, of course, only fuel for the ecological realist fire.

If this is the case, Turvey et al argue, then evolution must have provided the marsh periwinkle with the concept climbable, and must have provided us with, among others, concepts such as that of a neutron star. Where do these concepts come from? Turvey at al assume that if a conceptual ability is innate, then it must have been inherited from some ancestor [13]. The question of origin is then asked for the ancestor, with the same, regressive, answer.

Turvey et al conclude that, in order to explain concept origin, "the representational medium must exhibit preadaptive foresight, being able to represent all relevant states of affairs, be they extant or future ... there is no sensible scientific story to be told about such foresight" (p. 251).

Michaels and Carello (1981) concentrate on the mechanism by which instances are to be comprehended under concepts. If concepts are stored in memory, how is a perceived instance connected with the right concept? This is an instance of the general pattern recognition problem: if shapes are recognized by fitting them to templates (general patterns), what mechanism recognizes the best fit of template to shape? The standard answer, e.g. for the recognition of geometrical shapes, is that the input is examined by a set of detectors for specific shapes that work in parallel, and that emit signals proportional to the quality of the fit between the input and the shape they detect (Lindsay and Norman, 1977 provide a particularly clear discussion).

Michaels and Carello ask how such a system could detect affordances. Presumably, one could imagine a set of stored "affordance templates" that compared incoming data with certain characteristics that indicated the presence of an affordance. Given the formality condition, the comparison would have to be between the form of the input representation and that of the template. The form of the input representation must, therefore, contain enough information to specify which template is the right one, i.e. to guarantee a unique match. Producing the right form for an input representation, i.e. the form that conveyed the right information to the pattern matcher, would, Michaels and Carello argue, require that the form-producing system knew what information the form had to convey. This, however, is tantamount to the claim that the form producing system knows the relevant affordances, and could recognize the affordance indicated by each input. Otherwise, they claim, there is no way to understand how the form generated by the form-producing system could be the right one, i.e. the one that the pattern-matcher matches the right pattern to. This, however, also leads to an infinite regress.

Given the results of this argument, its premise bears examination. Is the premise that concepts are required for property identification plausible? Is it entailed by cognitivism?

One of the more familiar and flagrant appeals to conceptual knowledge in an information-processing system is that of Fodor (1968). Fodor explains the ability to tie one's shoes roughly as follows. In one's head is a motor executive with a library of programs. When one wants to tie one's shoes, the motor executive takes down a book containing a shoe-tying program, and reads and follows the instructions in the program (cf. Michaels and Carello, 1981, p. 135-137). When the instruction, END is reached, the executive puts the book away and waits for the next instruction from whatever process decides what one wants to do.

In Fodor's account, the book contains knowledge of how to tie one's shoes. The executive, on the other hand, contains knowledge specifying what to do, e.g. what muscles to move, when each instruction is read, and knowledge of how to read the book. The account thus appeals to three sets of tacit (Fodor's term) conceptual knowledge to explain shoe tying.

One can tell an analogous story about the perception of plant stems by the marsh periwinkle. The periwinkle contains two executives, a percept executive and a motor executive. The motor executive has a library similar to that envisioned by Fodor; it contains books with programs describing how to climb and avoid. The percept executive is a pattern matcher. It has a similar library, containing books describing features of perceptual media, and the relations of such features to features of the visible world. It associates each set of features with an affordance. When the percept executive, observing the signals coming in through the eyes, sees something that, according to the list of criteria in the book, specifies the affordance "climbable," it signals the motor executive to take out the book labelled "climbing" and to execute its program.

This program for the marsh periwinkle requires three knowledge bases: sets of criteria for identifying affordances, programs for carrying out actions, and knowledge of how to read the books. Cummins (1982b; 1983) has called the first two knowledge bases "internal manuals," and models of this form "internal manual models" (IMM's).

The analysis of Cummins (1982b; 1983) shows that IMM's are <u>never</u> required to reproduce the I/O properties of an information-processing device. Simpler models that do not involve internal manuals are always possible.

The major premise of the recognition regress argument, that property identification requires stored conceptual knowledge, is therefore true, at least if the knowledge is stored as it is in standard cognitivist IMM's. The recognition regress argument shows that if this premise is true, infinite regresses follow. Cognitivism, therefore, involves infinite regresses as long as it appeals to IMM's. The argument may, however, be unsound when applied to information-processing theories that do not involve IMM's. This possibility will be considered in Ch. V.

Summary

In summary, the arguments of the ER-Cognitivism debate are all somewhat unsatisfying. The constraint argument establishes that ER will be false if it cannot provide a coherent notion of selection that has constraining power, but does not establish that it cannot do so. The Ullman-Marr argument establishes that ER must contenence a particularly low-level form of information processing, but does not show that it cannot do so. While ecological realists have often claimed that they reject all forms of information processing, it is not clear that they have considered the form discussed in the Ullman-Marr argument.

The ER constraint argument challenges cognitivism to provide a way of showing that mentalistic models are not arbitrary; however, it largely ignores standard arguments (e.g. Cummins, 1983) that they are not. The cognitive penetrability criterion, in particular, meets this challenge, at least partially, by providing a way of saying when mentalistic models are not necessary. This partial success suggests, at least, that cognitism is constrainable.

Similarly, the recognition regress argument establishes that a cognitivism based entirely on IMM's would lead to infinite regresses. However, it does not show that cognitivism must be based on IMM's. The work of Cummins (1983), in fact, shows just the opposite.

The debate between ecological realism and cognitivism appears, therefore, to be a standoff. The arguments given so far certainly do not justify Fodor and Pylyshyn's (1981) claim that ER is "fundamentally wrong" (p. 195) or Turvey et al's (1981) claim that cognitivism is not "scientifically acceptable" (p. 238).

Breaking this theoretical standoff requires, at the very least, seeing whether either theory can meet the challenges posed by its opponents without sacrificing its fundamental assumptions. To this end, chapters IV and V critically examine, and reconstruct the two theories in their own terms. The debate is then re-opened as a contest between the reconstructed theories in Ch. VI. Before proceeding with this, however, it is useful to examine in more detail the goals, both tacit and explicit, of the theories in question.

NOTES - CHAPTER II.

1. Computationalists, as well as ecological realists, have proprietary views about what the performance of a natural system, and in particular, of a natural perceptual system, _is_ in any particular case. This question will be considered in detail in Ch. III.

2. The claim that a computer running a program that instantiates an algorithm describing an information-processing capacity of a natural system _duplicates_, as opposed to merely simulates, the capacities of the natural system is quite tendentious (e.g. Searle, 1980). It will not, however, be considered here; see Cummins (1983, Ch. III), and Pylyshyn (1984, Ch. 2).

3. The computational level of description is the level of description of an interpreted program (see also Ch. V). Fodor (1980) and Pylyshyn (1980; 1984) refer to this as the "semantic" level of description; Newell (1982) calls it the "knowledge level."

4. The taxonomy presented here is, of course, only one of many. Dennett (1984), for example, divides the field into two general groups: those who espouse the MIT version of cognitivism ("high church computationalists" or "east coasters," e.g. Fodor or Pylyshyn) and everyone else ("west coasters"). The latter designation includes connectionists such as Anderson and "neuronalists" such as Patricia Churchland. It also includes those who accept the notion of information processing, but reject the formality condition (Fodor, 1980; see also below), and hence reject computationalism. The latter group, whom Dennett christens "zen holists," includes diverse philosophers and psychologists, among them Nagel (1974), Dreyfus (1979; 1981), Haugeland (1978; 1979; 1981a), Searle (1980; 1981; 1983), and Heil (1981; 1983), all of whom share a distrust of computational approaches and a concern with describing experience. This latter group will not be discussed further.

5. The term, "folk psychology" was introduced by Lewis (1966). Pylyshyn (1984) argues that folk psychology is the proper conceptual starting point for cognitive science, and that the conceptual basis of cognitive science should remain fairly close to that of folk psychology; Stich (1983) and Paul Churchland (1984) argue that folk psychological concepts will have no place in a mature cognitive science.

6. This claim is meant to be taken literally. The interpreta-

tion of computationalism as putting forward a computer "metaphor" or "analogy," which is popular with ecological realists (e.g. Michaels and Carello, 1981) and zen holists (e.g. Heil, 1983), is a misrepresentation. This point is made clearly by Haugeland (1981a): "The basic idea of cognitive science is that *intelligent beings are semantic engines* ... people and intelligent computers ... turn out to be merely different manifestations of the same underlying phenomenon" (p. 31, his italics). See also Pylyshyn (1984, p. 43).

7. It is worth pointing out that the view that perception requires inference is much older than computationalism, going back at least to Locke's *Essay Concerning Human Understanding*, and in psychology, to Helmholtz's *Treatise on Physiological Optics*. That perceptual input is impoverished is, indeed, plausibly included in folk psychology.

8. A transducer is defined as a device that translates information from one medium to another (e.g. Pylyshyn, 1984, p. 151). The loudspeaker provides a good example; its function is to translate from modulated alternating current to modulated sound. The translation must proceed without processing, i.e. no information can be added to the signal being transduced. A transducer can, however, be noisy, or its output can be equivocal (see Shannon, 1948); in such a case, information is lost during the transduction.

9. Many information-processing psychologists, e.g. Fodor (1984), agree that a mentalistic psychology cannot apply across the board. They reserve mental state attributions for complex organisms, and describe simple organisms with the same vocabulary used to describe the functions of transducers. Even Fodor (1984), however, admits that it is an open question whether any particular lower animal can be characterized as a transducer.

10. If the distinction between being able to see things a certain way and actually seeing them that way (on all occasions of contact) is ignored, then knowing something in the ecological realist sense is just having a certain ecological law apply, one specifying that a certain thing will be seen a certain way, and that a specific behavior will result. In that case, affordances (see below) are just causes of actions. Ecological realists avoid this conclusion, as will be seen below, by introducing a concept of attention as the "control of detection" (Michaels and Carello, 1981, p. 69). If an organism is not "attending to" an object, it will not detect its affordances, even though it "knows" how to detect them. Behavior is, therefore, determined not by the affordances themselves, but by the epistemic perception of affordances. The question of attention will be discussed in detail in Ch. IV.

11. The term 'direct', while historically motivated, is unfortunate, as it gives rise to various misconstruals of the issue. Searle (1979), for example, has argued that perception cannot be

explained in terms of representations because representations are not the objects of perceptual experience. This may well be an argument against sense-data theories that claimed that what perceivers experience are not objects but sense-data, such as the claim that "it is always sense-data that are directly perceived" (Ayer, 1956, p. 97), but it is not an argument against cognitivism, the intended target of Gibson's introduction of "direct" perception. Both cognitivism and ER, correctly or not, concern themselves with perceptual information transfer independently of perceptual experience (Fodor and Pylyshyn, 1981; Turvey and Carello, 1981). Ecological realists themselves, however, in casting the debate in terms of what organisms *see*, seem sometimes to confuse cognitivism and sense-data theories. The "direct perception" debate is not about experience, not about what organisms see, and not about the effects of perceptual media. It is about whether signals are informationally impoverished with respect to the perceived properties of distal objects. Cognitivists claim that they are, and that the organism must, therefore, add information to that available through perception. Ecological realists claim that signals are not impoverished, and therefore, that the organism need not and does not add information to that available through perception. The direct perception debate is, therefore, about whether perceptual input is inferentially augmented. If organisms do not add to the information available in perceptual media, then they can be assumed, at least with respect to perceptual information transfer, to be law-instantiating "black boxes." Proponents of ER claim that this conception of the organism is sufficient in psychology; cognitivists claim that it is not.

12. "Ecological optics" is Gibson's term for the study of the specification of affordances in perceptual media (Gibson, 1979; Reed and Jones, 1982). If something is a projectible property for ecological optics, then it is a property that can be specified by a signal in a perceptual medium.

13. They do not consider the possibility that conceptual abilities arise from selected mutations. This is, however, at least a possible explanation (Fields, 1983a). It will be considered further in Ch. IV.

CHAPTER III

THE EXPLANANDA

In order to compare cognitivism and ecological realism, it is first necessary to have a specification, if not of what perception is, at least of what the explananda of perceptual theories are, that is neutral with respect to the cognitivism-ecological realism debate. It is not obvious, however, just what the task of a mechanistic theory of perception is. In most cases, the immediate result of the perceptual process, i.e. the output of the perceptual system, is not experimentally observable [1]. What counts as perception is therefore, to a large extent, determined by perceptual theory itself. That is, part of the task of a perceptual theory is to say what perception is by saying how the input/output function of the perceptual system should be characterized [2]. Both cognitivism and ecological realism include such specifications of the nature of perception.

It will be argued in this chapter that an adequate definition of perception that is neutral in the required way is not likely to be forthcoming. This difficulty can be circumvented, however, by identifying certain explananda that are common to ER and cognitivism. Perception can then be identified as *whatever* process explains these phenomena [3].

Attempts to Define Perception

A theory of perception must define perception in such a way that the processes involved in perception are clearly distinguishable from other processes, such as digestion or lactation. In particular, a theory of perception must be able to distinguish perception from other psychological processes (at least, it must do so if it countenances other psychological processes). Even a theory, such as ER, that insists that perception and action be treated together (Michaels and Carello, 1981, p. 47-48) must be able to distinguish the process of perception from the process of action.

Cognitivists draw what appears at first glance to be a straightforward line around perceptual as opposed to general cognitive processes with the claim that perception is whatever process eventuates in perceptual beliefs. This claim, unfortunately, immediately runs into trouble even within cognitivism: it is not clear what differentiates perceptual beliefs from other kinds, or how a theoretical criterion to do this job could be applied experimentally [4].

The characterization of perception in terms of beliefs does not, moreover, provide even a useful starting point for a specification of the task of perceptual theories that is neutral with respect to the cognitivist-ecological realist debate. As seen in the last chapter, ecological realists flatly reject appeals to beliefs, or to mental states of any kind (e.g. Gibson, 1979). Instead, they typically characterize perception as whatever process provides the information about the environment necessary for the coordination of behavior (Michaels and Carello, 1981).

This characterization is, however, no more neutral than that assumed by cognitivists. As Michaels and Carello point out, it allows episodes of perception to be of arbitrary length, even to last the length of the organism's life, and allows organisms to do things, e.g. explore their surroundings, in the course of a perceptual encounter. It therefore begs the question against the cognitivist's claim that the effective stimuli for belief fixation are typically both momentary and informationally impoverished (Fodor and Pylyshyn, 1981) [5].

The prima facie characterizations of both approaches are, therefore, question-begging. Indeed, since cognitivism and ecological realism disagree precisely on whether perception requires inference, and on whether there are perception-action laws, any definition of perception that takes a stand on either of these issues will be question-begging. A more general way of defining, or otherwise characterizing, perception is therefore needed for the purposes of analyzing the ER-cognitivism debate.

A somewhat more sophisticated characterization of perception can be constructed using Dretske's (1981) notion of information transfer. A system S being in state S(i) carries the information, in Dretske's sense, that a system S' is in a state S'(j) if and only if the conditional probability $P(S'(j)|S(i))$ is unity. Consider an organism O that perceives that a distal stimulus S has some property P. One can ask what state of O, if any, carries the information that S is P. If such a state can be found, one can characterize as perception whatever process produces, for a given O and P, such a state in such situations. This account of perceptual information transfer is quite general. It is, therefore, interesting to see if the Dretskean account can be used to generate a characterization of perception that is neutral between ER and cognitivism.

The Dretskean account, unfortunately, runs afoul of the perceptual illusions. Given sufficient technological control of perceptual media, e.g. given the ability to produce flawless holograms of any visual scene, one could, at least in principle, break the one-to-one correlation between the distal situation and the organismic state required by Dretske's account in any particular case. If this were true, however, no organismic state would carry the information that any S was P, for any non-disjunctive P, i.e. for any P that does not include as a disjunct the property, "is an illusion." Hence no process would count as, e.g. seeing that snow is white, or as seeing a red car go by. What we see, however, are things like snow and cars, not things like (snow or illusion) or (car or illusion) [6].

As seen in the last chapter, cognitivism and ecological realism interpret the problem of illusions in diametrically opposed ways. Cognitivists interpret the possibility of illusions as proof of the claim that perception requires inference (Gregory, 1970; Fodor and Pylyshyn, 1981; Rock, 1983). They therefore argue that organismic states carry the information that distal objects have particular properties only because organisms infer that the distal objects have such properties. Ecological realists, on the other hand, interpret the possibility of illusions as proof of the claim that ecological laws are niche-specific. They argue that organismic states carry the information that certain affordances are present not because of any inferences carried out by the organism, but because the notion of information transfer must be relativized to the niche.

Therefore, while the Dretskean account does not beg any questions against either ER or cognitivism, it cannot provide any common ground between the two theories, either. As an alternative strategy, one can ask for a characterization of the explananda of perceptual theories that does not require perception to be defined in advance. Whatever else might be included in their domains, perceptual theories certainly must explain the fact that perception can function in the coordination of complex behaviors such as speech, locomotion, etc. The coordination of behavior is therefore the most obvious place to look for an independent characterization of the explananda of perceptual theories.

Perceptual Constancies

Humans have the so far unique capacity to report to an experimenter what they perceive. Perception reports, therefore, constitute an important class of explanada for perceptual theories. Such reports, however, demonstrate a striking lack of correlation between either proximal or distal stimuli and what is reported as perceived. To take a familiar example, human subjects who view a coin on a flat table report seeing a round coin, even when the viewing angle is oblique enough that the proximal stimulus - the projection of the surface of the coin on the surface of the retina - is elliptical. An even more puzzling case is provided by the perception of alphabetic characters. A very large variety of shapes can be readily identified by literate humans as tokens of letters such as 'a'. This is true even though, for example, the proximal stimulus corresponding to 'a' has a shape much more similar to that of the proximal stimulus corresponding to 'o' than to that corresponding to the 'A'.

The fact that coins are perceived to have constant shape when viewed from different angles, or that marks with very different shapes can be identified as tokens of the same letter, are examples of the phenomenon of _perceptual constancy_ (e.g. Lindsay and Norman, 1977). The perceptual constancies are systematic departures from strict one-to-one correlations between stimuli (proximal or distal) and responses. It is possible that, in some cases, the constancies represent correlations between specific unknown features of stimuli and responses. This is, however, not true in general; it the case of letter perception, for example, what letter is expected in a particular context significantly influences letter perception (Lindsay and Norman, 1977). Perceptual constancy is, therefore, one of the reasons why the theortical analysis of perception is so difficult.

Perceptual constancy effects are not restricted to humans. The behavior patterns of animals across the phylogenetic spectrum exhibit effects that can be attributed to perceptual constancy.

Consider the perceptual guidance used by an animal when running through thick woods. In such a case, perceptual guidance must be quite accurate; otherwise the chance of collision with a tree will be large. The proximal stimulus, as e.g. determined by the patterns of light and dark projected on the retina, is, however, constantly changing. The changes in proximal stimulus may be correlated with small changes in the behavior pattern, e.g. with swerves to the left or right; however, the main features of the behavior, such as average speed and direction, may remain constant. Moreover, it may be possible, e.g. for a animal with good night vision, to execute the behavior pattern under various conditions of illumination, even though changing the illumination conditions significantly changes the progression of light-dark patterns on the retina. A natural explanation of these facts is to say that what the animals sees while running remains essentially constant, i.e. to claim that there is a perceptual constancy associated with proximal stimuli of the type encountered when the animal runs through the woods.

Consider again the case of the marsh periwinkle described in the last chapter. The periwinkle behaves as if it sees all plant stems as the same. Here again, there is an independence of particular details of the proximal stimuli that suggests perceptual constancy.

Given the evident lack of one-to-one correlations between proximal stimuli and behavior, the fact that the behavior of most organisms generally makes sense demands explanation [7]. The concept of perception constancy allows much of this behavior to be organized in a coherent way. A theory of perception is faced with the task of explaining why this way of organizing stimulus-response pairs is so useful. A theory of perception is, in other words, faced with the task of explaining the perceptual constancies. In what follows, the perceptual constancies will be taken to be the prima facie explananda of the theories being considered. Perception can be thought of as whatever process makes perceptual constancies constant. This implicit definition of perception will be adopted in the discussion that follows; asking whether ecological realism or cognitivism can explain perception will be taken to be asking whether they can explain the constancies.

Both perceptual theories under consideration are parts of global theories of cognition. They therefore have, in addition to the goal of explaining perception, the goal of advancing the cognitive theories of which they are parts. This larger goal is most obvious in the interests of proponents of both theories in changes in perceptual abilities. The version of the recognition regress argument advanced by Turvey et al (1981), for example, is driven by a concern with the mechanisms of perceptual change. The need to incorporate the theory of perception into a general theory of cognition, therefore, is a significant constraint on the development of the perceptual theories being considered. The effects of this constraint will be especially apparent in Ch. VI.

NOTES - CHAPTER III.

1. There is, in the case of humans, always introspection. Introspection is not very useful, however, for the purposes of cognitive science (but see Shanon, 1984 for an opposing view). There are at least two reasons for this. First, we do not know what phenomenal states are like, or even if phenomenal states exist, in lower organisms. Introspection could, therefore, only be of any use with humans. There does not, however, appear to be any principled relationship between what humans can introspect and the outputs of perceptual processes. For example, Marr (1981) points out that humans can, by defocussing their eyes, become aware of the "raw primal sketch," the brain's initial representation of the visual world. Humans are not, however, aware of the processes involved in natural language parsing, even though this is a higher level (although encapsulated) process (Fodor, 1983).

2. This can be compared with the case of learning. In order to be plausible, theories of learning must rule out not only cases of maturation and injury, but also such phenomena as sunburn as cases of "learning." This is done by adopting a theoretically proprietary definition of learning. For example, cognitivism distinguishes rational cognitive learning from such processes as operant conditioning, a distinction that behaviorists cannot make (e.g. Lindsay and Norman, 1977).

3. This is not an unusual situation in science. It is arguable, for, example, that the theoretical terms of physics, such as 'electromagnetic field' are defined implicitly by the models, such as Maxwell's equations, that are constructed to explain certain antecedently characterized explanada, such as light and radio waves (e.g. Nagel, 1961).

4. The cognitive penetrability criterion of Pylyshyn (1980; 1984), and Fodor's work on informationally-encapsulated modules (Fodor, 1983) may be viewed as attempts to generate such a criterion (see also Ch. II, V).

5. The ecological realist characterization, moreover, makes no distinction between perception and cognition in general. Consider, for example, the relation between perception and action that occurs when one follows instructions for an activity such as cooking. Cognitivists explicitly separate such activities into a perceptual component and a non-perceptual cognitive component that includes, e.g. such processes as remembering where the measuring cups are

stored. By ruling out such a distinction, the ecological realist characterization of perception begs a second important question against cognitivism.

6. This is true pretheoretically, at any rate. It is this pretheoretical description of perceptual information transfer that must be accounted for by an adequate theory.

7. The fact that behavior patterns can be invariant under changes in both proximal and distal stimuli is often, somewhat misleadingly, referred to as the "stimulus independence" of behavior (e.g. Newell, 1980; Pylyshyn, 1984). Newell (1980) takes stimulus independence to be one of the criteria for a device to be a "physical symbol system," i.e. a general symbol-manipulating system, the type of system that cognitivism attempts to explain.

CHAPTER IV

ECOLOGICAL REALISM RECONSTRUCTED

The cognitivist arguments presented in chapter II pose three problems for ER. The problem of illusions will be assumed, for the present purposes, to have been answered by the ecological realist claim that descriptions of affordance perception must be relativized to a particular niche. The problems raised by the constraint argument and the Ullman-Marr argument cannot, however, be answered so quickly. The constraint argument requires ER to clarify its solution to the specification problem, and in particular, to clarify the role of selection-by-learning in that solution. Second, the Ullman-Marr argument requires ER to make some response to the claim that the detection of signals always requires information processing, and to offer some explanation for the neural structures that appear to carry out that information processing.

This chapter considers the explanatory structure, and the factual claims, of ER in light of these problems. It shows, first, that the explanatory power of ER resides in the concepts of resonance and selection, not in the concept of ecological law. Resonance, recall, is the proposed mechanism of information detection in ER, while selection, by either evolution or learning, is the mechanism by which organisms become competant to resonate to a particular affordance-encoding property of signals in a perceptual medium.

The explanatory power of these two notions is then examined. It is argued, first, that the radio (Gibson,1966) and planimeter (Runeson, 1977) analogies employed by ecological realists to explicate and defend the notion of resonance are inadequate, and that reliance on these analogies leads ecological realists to underestimate seriously the difficulty of using resonance as a theoretical notion to explicate the perceptual guidance of behavior. A realistic model of the processes involved in resonating to information bearing signals is presented. It is argued that resonance always requires the use of stored information. ER must, therefore, countenance the storage of information by organisms if it is to appeal to resonance as a detection mechanism. The debate between ER and cognitivism, therefore, cannot be about whether organisms store information; it must be about what information is stored, and how it is stored. This result is independent of the question of the impoverishment or richness of perceptual information; organisms must use stored information to detect affordances even if affordances are fully specified in perceptual media.

Once resonance itself has been understood, the control of resonance must be characterized. It is argued that ER can only explain the control of resonance, and hence perceptual plasticity, by carrying out a functional analysis of the organism. ER cannot, therefore, adopt the view of an organism as a law-instantiating black box; it must instead view the organism in much the same way that cognitivism does, as composed of interacting functional subsystems. ER must, in other words, adopt the analytic, as opposed to subsumptive explanatory strategy. Ecological laws are, therefore, not the explanans, but the explanada, of ER.

Evolution and learning are described as mechanisms by which the control of resonance is effected. It is argued that, given this analysis of resonance and its control, changes in an organism's capacities to detect affordances must be due to evolutionary selection, not learning.

Finally, the question of inference is discussed. It is shown that ER must accept explanatory appeals to inference in a wide sense, i.e. in a sense that does not require appeal to explicitly represented propositions serving as premises and conclusions. This question, too, is independent of the question of impoverished input. The inferences required do not serve to compensate for stimulus poverty; they rather allow the extraction of coded information from overly _rich_ proximal stimuli.

The reconstruction of ER advocated here shares the general framework and guiding assumptions of standard ER. In particular, it accepts the claims that perception must be considered relative to a specific niche, and that organisms do not infer the affordances of their environments from impoverished proximal stimuli. It is, however, freed of the adherence to the purely subsumptive explanatory strategy that trammels standard ER. The use of the analytic explanatory strategy allows the reconstructed "analytic" ecological realism presented here to explain phenomena, such as the mechanism and control of resonance, that cannot even be addressed in standard

ER. The strength of the reconstructed theory in the ER-cognitivism debate will be assessed in Ch. VI.

The Explanatory Power of Ecological Laws

The Ullman-Marr argument begins with the claim that ecological laws, by themselves, have very little explanatory power. It is, therefore, necessary to see what ecological laws do explain, and how they do it.

As shown above (fig. II-2), an ecological law relating the perception of a particular affordance to an action is a composite of two laws, a specification law and an action law [1]. The specification law relates the affordances of the perceived object for the perceiving organism to the properties of the signal by which it is perceived by the organism in question. The specification process must guarantee, within a niche, that affordances for particular organisms are uniquely specified by signals detectable by those organisms. Signals specifying affordances are detected by means of resonance. If an organism is attuned to a particular affordance, it will resonate to a signal encoding it, and thus detect it. A specification law, therefore, relates an affordance to a signal, <u>and</u> relates the signal in question to the resonating abilities of the organism. The action law relates the detection of a signal specifying an affordance to an action.

Ecological laws are, therefore, restricted in their range of application not only to a specific niche, but to a specific state of attunement on the part of the organism. One must ask what effect these restrictions have on the explanatory power of the laws.

The traditional, Logical Empiricist view of subsumptive explanation is that the explanatory power of a law increases with its generality (e.g. Nagel, 1961). For example, Newton's laws of motion are more explanatory than Galileo's because they are more general; the explanatory domain of Newton's laws properly contains that of Galileo's laws. This view of subsumptive explanation will be adopted here.

A straightforward consequence of this view is that a law-like, i.e. counterfactual supporting statement with little or no generality has little or no subsumptive explanatory power. As the generality of such a statement decreases, i.e. as its range of applicability decreases, the pressure on those appealing to the law to explain why it holds in those, and only those, circumstances increases. In the limit of no generality, when the putative "law" holds of only one thing in one circumstance, saying that the thing obeys the law is just saying that it behaves in its own idiosyncratic way; <u>nothing</u> is explained by saying that it does what it does because it obeys its (completely idiosyncratic) "law."

Such, however, seems to be the case with ecological laws in ER. Their application is restricted not only to a particular niche, but also to a particular state of the organism. Therefore, there is considerable pressure on ER to explain why particular laws apply where they do, and not elsewhere.

The extent to which the domains of ecological laws are significantly restricted can, however, be expected to vary considerably with phylogeny. It is probably safe to assume, for example, that a wide range of individual sets of environmental conditions are ecologically equivalent, i.e. constitute the same niche, for many lower organisms such as worms, insects, or fish [2]. It is also safe to assume that at least most of the organisms of many species, or perhaps even of more distant relation, in many niches are in equivalent states of attunement. For example, all members of a large school of herring are probably attuned quite similarly to each other, and to the members of any other school of herring.

These assumptions are, however, quite unsafe in the case of higher organisms. All humans, for example, cannot be expected to be attuned to the same features of their environments, even though the classes of things to which they are attuned may have considerable overlap. Moreover, quite similar sets of environmental circumstances are notoriously non-equivalent in terms of what they afford a particular human. Slight changes, for example, in emphasis or tone of voice can allow a single grammatical structure in a natural language such as English to afford quite different meanings to a competent hearer. In the case of humans, therefore, one can, in general, assume neither the equivalence of individuals nor the equivalence of niches.

Ecological laws may therefore pick out anything from interesting behavioral generalizations (in the case of the herring) to idiosyncratic behavioral quirks (in the case of human language understanding). In the first case they would have, on the traditional view, at least some subsumptive explanatory power, while in the second case, they would have none [3].

This analysis of the explanatory power of ecological laws is not likely to go over well with ecological realists. Turvey et al (1981), for example, simply assume that if a law, i.e. a lawlike statement, can be found that describes a situation, then that statement explains the situation. This assumption embodies the Logical Empiricist claim that explanation is merely nomic description, i.e. subsumption.

The job of psychology is, however, to explain not only why a particular organism that perceives A does B, or even why any organism that perceives A does B. Psychology must also explain why, or better, how perceiving A leads to doing B. As Dretske (1977) and Cummins (1978; 1983) point out, no universal statement of the form "if A then B" shows why or how, in any particular case, A causes, produces, or leads to B [4]. In order to explain how the perception of an affordance leads to the production of a behavior, one needs not a law, but a mechanism.

To illustrate this point, consider the ecological realist solution to the problem of illusions. In an illusion, an organism perceiving an object produces a behavior that would have been appropriate if the object had the affordance that it appeared to have. Ecological realists argue that illusions are not a problem because, they claim, the specification law that would apply in the natural setting does not apply in the laboratory setting.

Even if this account of illusions is accepted, however, there is still some explaining to do. The fact remains that the organism exhibits a particular bahavior B, both in the natural environment and in the lab. Given that the specification law, and ipso facto the ecological law that is supposed to explain behaviors of type B does not apply, *what explains the behavior*?

There are two possible answers to this question. The subsumptive answer is the claim that, in the laboratory, the organism's behavior is explained by appeal to a *new* ecological law that relates the affordance that the illusory stimulus actually has for the organism to B. In the case of the Nuller-Lyer figure, for example, the figure plausibly has an affordance something like "causes mistaken length-perception reports" for most humans. According to the subsumptive account of illusions, there is a law linking this affordance to the reports of test subjects in laboratory illusion studies. Turvey et al (1981) explain a change in an organism's behavior pattern in another context as due to its "moving in the context of one set of (nested) laws rather than another" (p. 299). This at least suggests that they would adopt the same subsumptive strategy when dealing with illusions.

This new "law," however, is not only not general; it is completely unmotivated. The organism did not evolve in the lab, nor (on the first trial), can it be said that it learned to respond to the actual affordance of the "illusory" stimulus with B. In short, the new "law" is arbitrary; it has no explanatory power at all. We are still completely in the dark as to why the organism did what it did.

The second possible answer to the question of why the organism exhibits B is that the mechanism by which it detects the affordances of objects in its natural environment is unable to discriminate between the natural object and the illusion. While this is a problem for the organism, it is not itself a problem for ER; the organism does not detect the affordance that the illusory object actually has, it can be claimed, because it was not selected to do so (but see below). The organism does "detect" an affordance that the illusory object does not have, but this can be explained by appeal to the presumably evolutionary selection of the organism to respond to objects that have the affordance in question. Michaels and Carello's (1981) analysis of the Necker cube illusion (p. 180) adopts this strategy.

Both of these strategies for explaining what the organism does appeal to the fact that an ecological law that holds in the natural environment does not hold in the lab. This claim is, in fact, the heart of the ecological realist analysis of illusions. Given its importance, it is essential that ER explain it. Why should a law that applies "naturally" not apply in the lab?

There are two trivial answers to this question. First, it can be claimed that the law does not hold in the lab because the lab differs from the natural environment. This is certainly true, but it tells us nothing about why the differences in the two settings have the effect that they do; it only tells us <u>that</u> some differences are important as a matter of brute observational fact. This fact, however, does not explain why the law holds in one setting but not the other; it is, indeed, precisely this fact that needs explaining.

Second, it can be claimed that, as the specification of the law limits its application to the natural environment, it is impossible for it to apply anywhere else. This response, however, does not tell us why what the law says is true is only true where the law says it is true. Again, it does not tell us why the law applies in one setting but not the other.

A third answer is given by Michaels and Carello (1981): laws hold because the organism-environment relations that they describe are selected by either evolution or learning. Ecological laws can, however, hold in niches in which they were not selected, as long as the relevant properties of both niche and organism are preserved. If this is to be a substantive constraint on ecological laws, ER must provide some way of saying what the "relevant properties" that determine whether a law holds in a given niche are. The theory of resonance, described below, is an attempt to to do this. It provides an answer to the question of what properties an organism must have in order for an ecological law to hold. However, as will be shown below, this response to the question posed by the illusion problem abandons the black-box characterization of the organism, and makes essential use of a functional analysis of the organism's internal states.

In summary, ecological realists must explain both how the organism perceives the affordances that it does, and why ecological laws apply in, and only in, the situations in which they in fact apply. The first of these problems concerns resonance; the second concerns selection.

Resonance

When a guitar string is plucked, it selectively resonates to the component of the plucking motion that excites it at its natural frequency. This sort of resonance is a common feature of both natural and artificial energy-transfer systems. It is most efficient, and often only possible, to transfer energy from one subsystem to another at the resonant frequency of the receiving subsystem.

Ecological realists propose that perception be viewed as a process in which an organism resonates to information present in its environment (Gibson, 1966; 1979; Michaels and Carello, 1981). Michaels and Carello (1981), for example, claim that "biological systems resonate ... to information that the environment 'broadcasts'" (p. 83). Resonance is, therefore, the proposed mechanism of direct perception (see also Ch. II). Resonating to information is not, however, prima facie the same as resonating to energetic input of a particular frequency. While the latter process is understood in great detail, the former is not.

The importance of the notion of resonance in ER requires that it be explicated and defended with care. The arguments that have been advanced so far by ecological realists in support of the suitability of the concept of resonance as a description of affordance detection are, however, quite weak. Michaels and Carello (1981), who provide the most detailed treatment of resonance to date, employ two analogies, the "radio metaphor" of Gibson (1966), and the "planimeter analogy" of Runeson (1977), in an attempt to explicate the claim that organisms resonate to information, and thus to provide support for the claim that resonance provides a mechanism for direct perception. Due to the weight that is placed on them, these analogies are worth examining in some detail.

Michaels and Carello, to begin with, make much of the fact that the radio receiver, unlike a computer, does not store any of the information that it receives: "the radio metaphor captures the essence of an information-detecting machine. It is also a useful metaphor with which to contrast direct perception with the more usual storage or library metaphor of brain (sic), for in radios, a record is not stored, while in libraries it is" (p. 63). The radio is thus claimed to be an information-detecting device that needs no memory because it does not process information.

This fact, however, is merely an artifact of the usual use of radios. If every radio were built with an internal tape recorder, e.g. hooked up in parallel with the speaker, this fact would not change the operation of the radio receiver at all. If the detection of signals by the radio without the internal tape recorder is an analogy for perception, surely it is also an analogy if the radio has an internal tape recorder. Whether the radio stores what it receives in a memory is simply irrelevant to its operation.

What Michaels and Carello need to show, in order to prove that the radio does not infer information from the signal it receives, is not that it does not make a record of the transaction, but that it does not use any stored information. What it does with its output is, therefore, completely irrelevant; what is in question is how it detects its input. It is, therefore, useful to digress briefly to consider the operation of radios.

Radio stations transmit information by broadcasting a "carrier" wave of a particular frequency in which the information is encoded by amplitude (AM) or frequency (FM) modulation, i.e. by making small changes in either the amplitude or the frequency of the carrier that correspond one-to-one with the amplitude or pitch of the sound being transmitted. Radio receivers detect these signals by resonating with the carrier wave. The carrier wave induces an oscillatory current, with frequency equal to that of the carrier, in the antenna, an hence in an oscillator circuit. By tuning the receiver, one changes the frequency with which the oscillator circuit will preferentially oscillate. Whichever station transmits with a carrier frequency equal to the frequency to which the oscillator preferentially oscillates, i.e. to which the radio is tuned, will be picked up.

When a signal is detected, the oscillating current in the oscillator excites a "demodulator" circuit, which detects the small modulations in the carrier. The demodulator does not resonate to the carrier (it is "transparent" to the carrier); it rather "skims off" the modulations, and sends them to an audio amplifier, where they are amplified and sent to the speakers. The speakers convert the signal to acoustic waves that can be detected as sound. These functions are shown schematically in fig. IV-1.

Michaels and Carello provide a somewhat simplified account

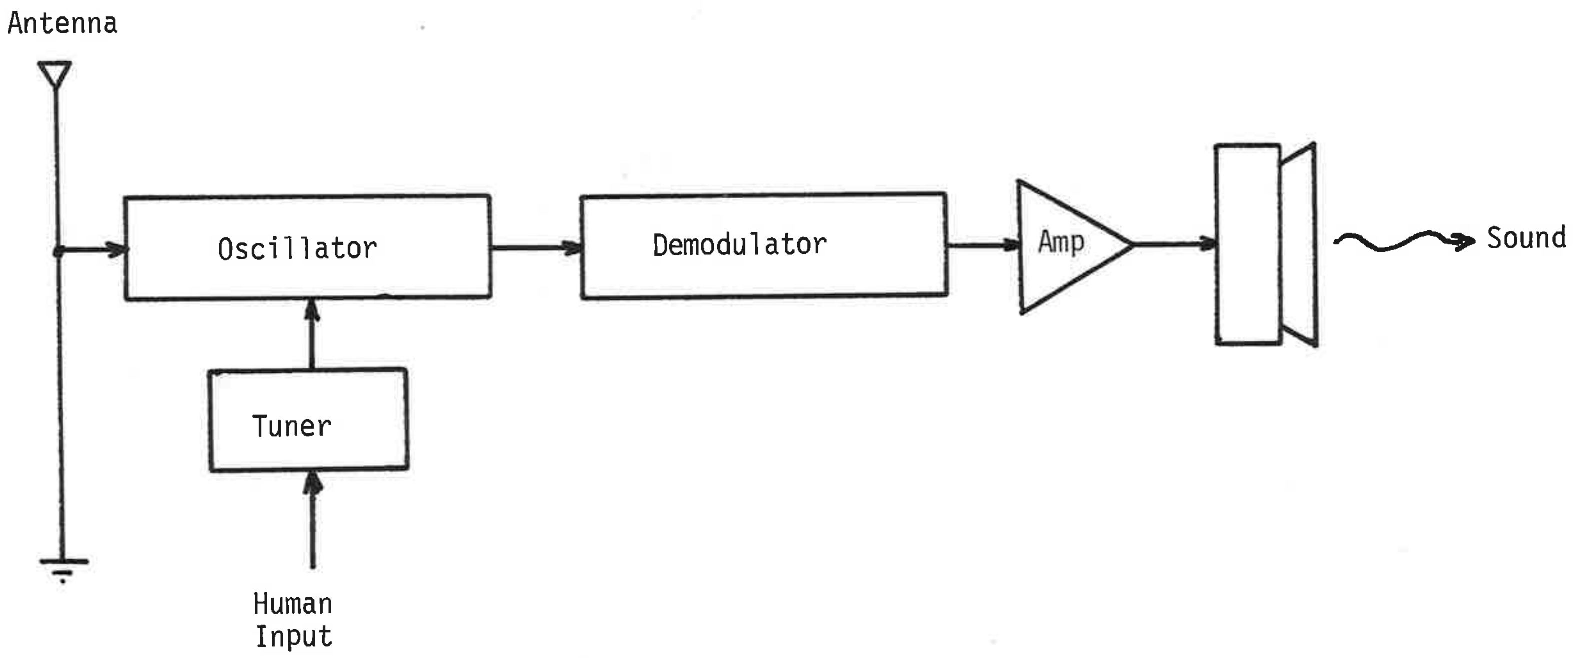

Fig. IV-1: Schematic of a radio receiver circuit. For details of such circuits, see Salm (1966), Rosenthal (1965). Tuning the oscillator is accomplished by human input, i.e. by someone operating the tuner control.

of transmission and reception (p. 63-64), and then introduce the analogy with vision as follows: "The parallels to the theory of direct perception are fairly obvious. In the case of vision, electromagnetic radiation (light) is modulated by reflection...The peripheral sensory organs, like the receiving antenna, must be transparent to the carrier frequency; that is, both must let the signal pass through. Finally, the information must be 'tuned in'" (Michaels and Carello, 1981, p. 64).

Unfortunately, this analogy breaks down in several places. The most obvious problem is that animals are supposed to resonate to affordances, i.e. to meaningful properties of objects in the niche. The meanings of radio transmissions are, however, lost on the radio receiver; it is completely insensitive to the semantics of the broadcast. It is only the listeners to whom the semantics make any difference. Moreover, the receiver is not "attuned," in any meaningful sense, to the <u>information</u> in the signal; tuning a radio selects not the information in a signal (the modulations), but only the carrier. Thus the radio not only does not "understand" the information in the signal in any way, it does not even resonate to the modulations that encode the information; it resonates only to the informationally-irrelevant carrier wave. If organisms were truly analogous to radio receivers, they would "resonate" to light, but not to information encoded in the light. They could tell, e.g. when the frequency of the incident light changed, but would be completely insensitive to information about affordances encoded in the light. They would, in other words, do exactly what ecological realists take such pains to argue that they do not do.

Consider, for example, the marsh periwinkle. The periwinkle is supposed to resonate to signals encoding the affordances climbupable and collide-withable. If the periwinkle behaved the way that a radio does, however, it would resonate not to the affordance-encoding modulations, but to the carrier, i.e. to the light itself. "Tuning" the periwinkle would not be changing the affordances to which it could respond, but would rather be selecting the frequency range of light to which it could respond, e.g. by changing the light absorbing molecules in the retina. This, however, is not the ecological realist view; it is precisely the view that ER rejects.

The radio metaphor is therefore worse than useless; it is seriously misleading. If the claim that perception involves resonance is to be explicated, let alone defended, a better analogy is required.

Michaels and Carello move from the treatment of the radio metaphor to a discussion of Runeson's analogy between the "direct" detection of properties specifying affordances and the "direct" measurement of planar area by the polar planimeter, a common drafting tool. The polar planimeter is chosen as an analogy for direct perception because it is "a real-world device that registers a higher-order property without computation" (Michaels and Carello, p. 66).

The planimeter is a simple mechanical device, consisting of a fixed pole and two rotatable arms linked end-to-end, with which one can measure the area of any closed planar figure by tracing its boundary with the end of the far arm of the planimeter (Michaels and Carello, p. 67). As Michaels and Carello point out, the device does not perform well measuring length, and does not measure area by multiplying lengths.

The planimeter does not, however, strictly speaking, measure area. It rather allows the comparison of areas in the same way that a sensitive balance allows the comparison of masses. A dial on a planimeter reads area in square centimeters, for example, only because it has been calibrated to do so. The calibration is accomplished by, e.g. marking '1' at the position on the dial that is indicated when the planimeter is used to measure the area of a square previously known to be one square centimeter in area. Standard units of area (including the square centimeter) are, however, defined in terms of conventional length units. The area used as a calibration standard is known, therefore, only because someone, at some time or other, actually measured the sides of such a square and calculated the area. Therefore, while the planimeter itself does not calculate areas from lengths, it can perform its function only because something calculates areas from lengths. Planimeters cannot be direct area detectors unless there are some indirect, i.e. calculating, area detectors.

It can be objected at this point that the fact that areas are defined in terms of lengths is merely a fact about our culture. Surely a system of units could be devised in which area was fundamental, i.e. not defined in terms of any other units. Even if areas were taken as fundamental, however, area detectors would still measure area only relative to an arbitrary, conventional standard. This standard would be fixed by convention, not detected. In the absence of such a standard, the planimeter cannot detect area; all it would detect is the ratio, in arbitrary units, of the areas of two figures. While it might be claimed that the planimeter detects area ratios directly, it is simply false that it detects areas directly. Whatever is to be said about the general questions of what sorts of properties can be detected directly, or of the role of standards in property detection, it is clear that the planimeter analogy, as formulated by Runeson or by Michaels and Carello, does nothing to support the specific claims of ER.

The possibility remains open that some device directly detects an interesting higher-order property without being calibrated to do so by means of a procedure that requires the detection of lower-order properties. The planimeter analogy does nothing to defend this claim, however; it leaves open the alternative possibility that nothing can detect a higher-order property, except in the degenerate sense in which, e.g. a rickety barn detects the presence of a high wind by falling down, without being calibrated to do so in the way that the planimeter is. The planimeter analogy is, therefore, also of no use in either explicating or defending the claims of ER concerning resonance.

In summary, neither of the analogies offered by ecological realists explicates the notion of resonance to information, or supplies any evidence for the applicability, or even the coherence of the notion. One must conclude, therefore, that the ecological realist notion of resonance is prima facie question begging; ecological realists appear, at least, simply to call "resonance" whatever mechanism the animal uses to detect information. If the concept of resonance to information is to carry any explanatory weight, this trivialization must be avoided. In particular, ecological realists must show that what they have called "resonators" are not, in fact, inference engines.

A more profitable approach to the question of resonance to information is to ask what sort of device could detect information directly through the use of resonance. To see the challenge of this clearly, imagine a radio receiver that actually did resonate to the modulations that encoded the information of interest. For example, instead of FM channels, tuning such a device might be accomplished by means of buttons labelled 'news', 'classical', 'blues', and 'rock', or perhaps (on a later model), 'good news', 'bad news', etc. When one pushed a button, e.g. the 'news' button, the receiver would scan the broadcast bands until it found one that was broadcasting news, and would then lock in on it.

It is, of course, well beyond our current technological capabilities to build such a device; doing so would, indeed, require solving the pattern recognition problem, a major problem of perceptual constancy, in the case of language recognition. One can, however, speculate about how such a device might work. Consider the problem of resonating to news instead of commercials. The voice pitch and tone patterns used (presently) by news and commercial announcers are different; therefore, it is at least plausible that news and commercial broadcasts could be distinguished on the basis of different characteristic modulations in the carrier. Assuming that this is true, one can imagine that the demodulator of the receiver could resonate to this "higher-order invariant" of the modulation encoding the sound of the announcer's voice, and so select only modulations encoding the right voice inflections and tones. In other words, one can imagine that the receiver, by resonating to a property of the modulation specifying the type of broadcast of interest, could selectively detect only certain broadcasts.

This design for a receiver that resonates to information clearly assumes that some property of a news broadcast _other than_ the semantic properties of the information being transmitted indentifies the transmission as news. If this is not true, then the idea of resonance to information is in serious trouble, as there is no reason to suppose that the semantics of a news broadcast is transmissible, or quasitransmissible, in any medium. Indeed, there is every reason to suppose that the semantics, e.g. of tokens such as "Iran," is completely conventional. If ER is to employ the notion of resonance to information, it must be prepared to argue, in each case, that such non-semantic transmissible properties exist. This question will be considered at length in Ch. VI.

This ficticious example of resonance to information can be compared with what is known to happen in the brain. Shepard (1984), for example, points out that resonance has long been taken seriously as a _non_-analogical description of the dynamics of neural state transitions in the brain (e.g. by Hebb, 1949). In particular, Grossberg (1973; 1980) has developed a detailed theory of the resonant responses of groups of neurons in the mammalian visual system to various stimulus patterns of light and dark. These responses are "resonant" because patterns in the stimulus cause isomorphic firing patterns of groups of neurons whose firing patterns were initially chaotic, just as, in cases of mechanical resonance, such as the response of a child's swing to the force of pushing, an initially chaotic motion becomes coherent due to the action of a coherent force. Grossberg's mathematical model predicts, from a very general consideration of the functions the visual system must perform, many of the transformations in firing patterns that are postulated by Marr's (1981) computer-simulation study of mammalian early vision, transformations that have been observed in numerous electrophysiological recording studies (reviewed by Otto-son, 1983; Sterling, 1983; Gilbert, 1983) [5].

Resonance also appears to be one of the major mechanisms serving to reinforce certain neural connections at the expense of others during visual system development, and presumably during brain development in general (Lund, 1978). Many studies employing sensory deprivation during development have shown that neural networks that respond to particular geometric forms (e.g. edges) in the normal adult do not develop if the animal is reared in an environment lacking such forms. The standard conclusion from such studies is that the neural activity caused by visual contact with objects of particular shapes during brain development is necessary for the establishment and maintanence of the neural connections required to see those shapes (e.g. Hubel and Weisel, 1977). The processes by which sensory stimulation regulates neural development can, in general, be understood as driven by resonance; the non-random connections made in response to repeated stimulation are self-reinforcing (e.g. Grossberg, 1980).

These studies indicate that at least some subsystems of perceptual systems, which, unlike radios and planimeters, do detect information, can be described as resonating to properties of incident sensory (e.g. optic) arrays that have at least something to do with the information in question. One can now ask a central question: Is this resonant detection of information direct in the ecological realist sense, i.e. does it involve no addition of information to that present in the stimulus?

This question can now be seen to be dangerously equivocal. Consider again the modified radio analogy. The signal detected by the radio encodes information, e.g. a news report. What eventually comes out of the radio in the form of sound is that information, re-encoded in a different medium. Ideally, the radio adds no information to the news report. With respect to the news report, therefore, the radio is a transducer; its sole function with respect to this information is to convert modulated radio waves into modulated sound.

This is not, however, the whole story. Consider the operation of the demodulator itself. Its job is to resonate selectively to signals with a certain property, the property that indicates that a signal encodes a news report. The demodulator is like a template; only signals with the right "shape" can be "seen" through it.

Templates, however, encode information: the information that a particular shape is the right one for some function. The demodulator that detects news reports also encodes information: the information that a particular property of the modulation indicates that it encodes a news report. Indeed, <u>any</u> resonator encodes some information, namely, the information that its natural frequency, or natural modulation, is whatever it is.

It is important to see that this information is not encoded in a proposition "written on" the resonator, e.g. in the sense of the "brain writing" Fodor (1975) advocates and Dennett (1978a) denigrates. It is, instead, encoded in the causal operation of the device, and in the causal connections between the device and other devices. Let us call this "causally-encoded" information. Cummins (1982b) points out that causally encoded information can be interpreted propositionally; i.e. the states of the device can be interpreted as tokens in a language. However, causally encoded information does not serve as data that the device uses as input; it rather characterizes the action of the device on its input.

The demodulator uses the information that it causally encodes to detect the signal of interest. In other words, because, and only because, the device is constructed so as to causally encode certain information, it is able to detect the signal that it in fact detects. Without the stored information, detection could not occur.

This analysis can easily be extended to the neural circuits in the brain that resonate to certain sensory inputs (the "feature detectors", etc.). These neural circuits causally encode certain information, such as the information that a certain optic structure represents an edge, and they can perform their selective detection functions only because they encode such information.

In summary, the answer to the question of whether detection by resonance is direct has two parts. First, the information encoded by the signal is not altered by a resonant detector. In the case of the radio, the news report is transduced. In the case of the visual system, the information present in the optic array is re-encoded in the firing patterns of the cortical neurons responsible for generating Marr's "2 1/2-d" and 3-d "sketches" of the environment (Marr, 1981). With respect to this information, detection is transduction. If the signal is impoverished, a resonant detector will produce an impoverished output; it the signal is rich, a resonant detector will produce a rich output.

The detectors in both cases, however, must use causally encoded information to accomplish detection. This requirement is independent of the question of whether stimuli are impoverished or rich. ER must countenance this use of encoded information if it is to explain perception by appeal to resonance.

Resonance can, perhaps, explain the perceptual constancies; it cannot, however, by itself explain how resonance to particular properties of the environment relates to action. In order to explore this relation, it is useful to digress, first, to consider the mechanisms by which the properties to which an organism resonates change through time. The relation of resonance to behavior, and the relation of behavior to the control of resonance, can then be considered simultaneously.

Control of Resonance

Consider again the behavior of the marsh periwinkle, the marine snail discussed by Turvey et al (1981). The marsh periwinkle exhibits two behaviors with respect to plant stems in its niche. When the tide comes in, the marsh periwinkle approaches, under visual guidance, and climbs a plant stem to avoid being washed away. At other times, it avoids plant stems, while moving around on the bottom feeding. What makes the marsh periwinkle respond differently to visual contact with the plant stem in the two circumstances, i.e. what makes it resonate to different affordance-encoding properties of the light in the two circumstances?

Turvey et al describe the situation as follows. The plant stems have at least two affordances for the periwinkle: "climb-upable" and "collide-withable" (Turvey et al, 1981, p. 298). When the tide is coming in, the periwinkle perceives the former affordance when it comes into visual contact with a plant stem; at all other times it perceives the latter affordance. "The occasion of contact with the incoming tide plays the role of a state of affairs that <u>selects</u> (in the sense of <u>attunes</u>) a marsh periwinkle/niche relation" (p. 298, their italics).

There are two strategies for explaining the control of affordance detection with which this statement is consistent. One is to claim that, in the two situations faced by the periwinkle, two different ecological laws hold. The tide, in effect, selects the law that subsumes the situation, and explains the behavior. This is the subsumptive strategy; the behavior is explained by appeal to laws, and to the "selection" of laws. Turvey et al appear to embrace this strategy when they describe the situation as involving "an organism, on a given occasion, moving in the context of one set of (nested) laws rather than another" (p. 299).

This explanatory route is, however, quite unsatisfying; it suffers from all of the problems shown earlier to plague subsumptivism in psychology. Moreover, no hint is given as to how an event, the advance of the tide, can "select" a law, let alone a particular law. Changing an ecological law involves changing the state of attunement of the organism to which the law applies. Saying that a new law is "selected," however, tells us nothing about how the change in attunement occurs. Unless ER can explain how laws, and with them states of attunement, are selected, its account of affordance detection, whatever else it accomplishes, removes none of the mystery of how things can work one way in one situation and another way in another. This is, in essence, the same objection to subsumptive explanation that was raised earlier; appeals to laws in ER tell us nothing about how or where those laws apply.

Another strategy is available, however. One can claim, in the case of the periwinkle, that the tide acts directly on the snail to select its state of attunement. This form of "selection" has the distinct advantage that a mechanism could at least be imagined that utilized ordinary, physiological causation to carry out the selection. For example, sensory neurons that detected, e.g. changes in water pressure or current flow could communicate with the neural circuitry that determines the response to particular visual stimuli. It is therefore at least possible that this form of selection could be explained by appeal to the periwinkle's neural organization. Moreover, selecting the state of attunement automatically selects, for a given organism, the ecological law being satisfied.

An organism is attuned to an affordance if, were it to be exposed to a signal encoding the affordance, it would resonate to the property of the signal encoding the affordance. Changes in attunement are, therefore, changes in the abilities of the organism to resonate to properties of incident signals. The information-receiving radio provides an analogy: when it is tuned to news, it cannot resonate to, e.g. music, and so cannot receive music.

Explaining changes in attunement, therefore, requires explaining how an organism's ability to resonate to particular properties of signals can be changed. The processes that change an organism's abilities to resonate must occur inside the organism's skin; when such changes occur, nothing changes outside the skin [6]. ER cannot, therefore, accept a characterization of the animal as a "black box;" it must provide some account of what occurs inside the skin of the organism when its abilities to resonate to particular properties change.

The simplest account of the mechanism of changes in resonance is based on the assumption that each organism has available resonators for properties encoding different affordances, and that different resonators are "turned on," and thus different affordances "tuned in," in response to different environmental circumstances. In the case of the marsh periwinkle, for example, one can imagine that two resonators are present, one of which detects properties of light encoding the affordance "climbable," and the other of which detects properties of light encoding the affordance "collide-with-able." These resonators are "turned on" or "off" depending on whether the tide is advancing or not.

If this assumption is accepted, the model of the periwinkle shown in fig. IV-2 results. The periwinkle is assumed to have a mechanoreceptor that is specific to the pressure variations caused by the advance of the tide. When the tidal advance is detected, the resonator for the affordance "climb-upable" is activated. This resonator responds only to visual signals that have the property, presumably a vertical intensity discontinuity of some minimal width, that is correlated in the periwinkle's niche with the affordance. When a signal with this property is detected, the resonator produces an output that signals the motor system to initiate climbing behavior, which is assumed to be a fixed-action pattern (see Kandel, 1979 for a discussion of similar patterns in Aplysia, another marine gastropod).

When the tide is not advancing, the resonator for "collide-

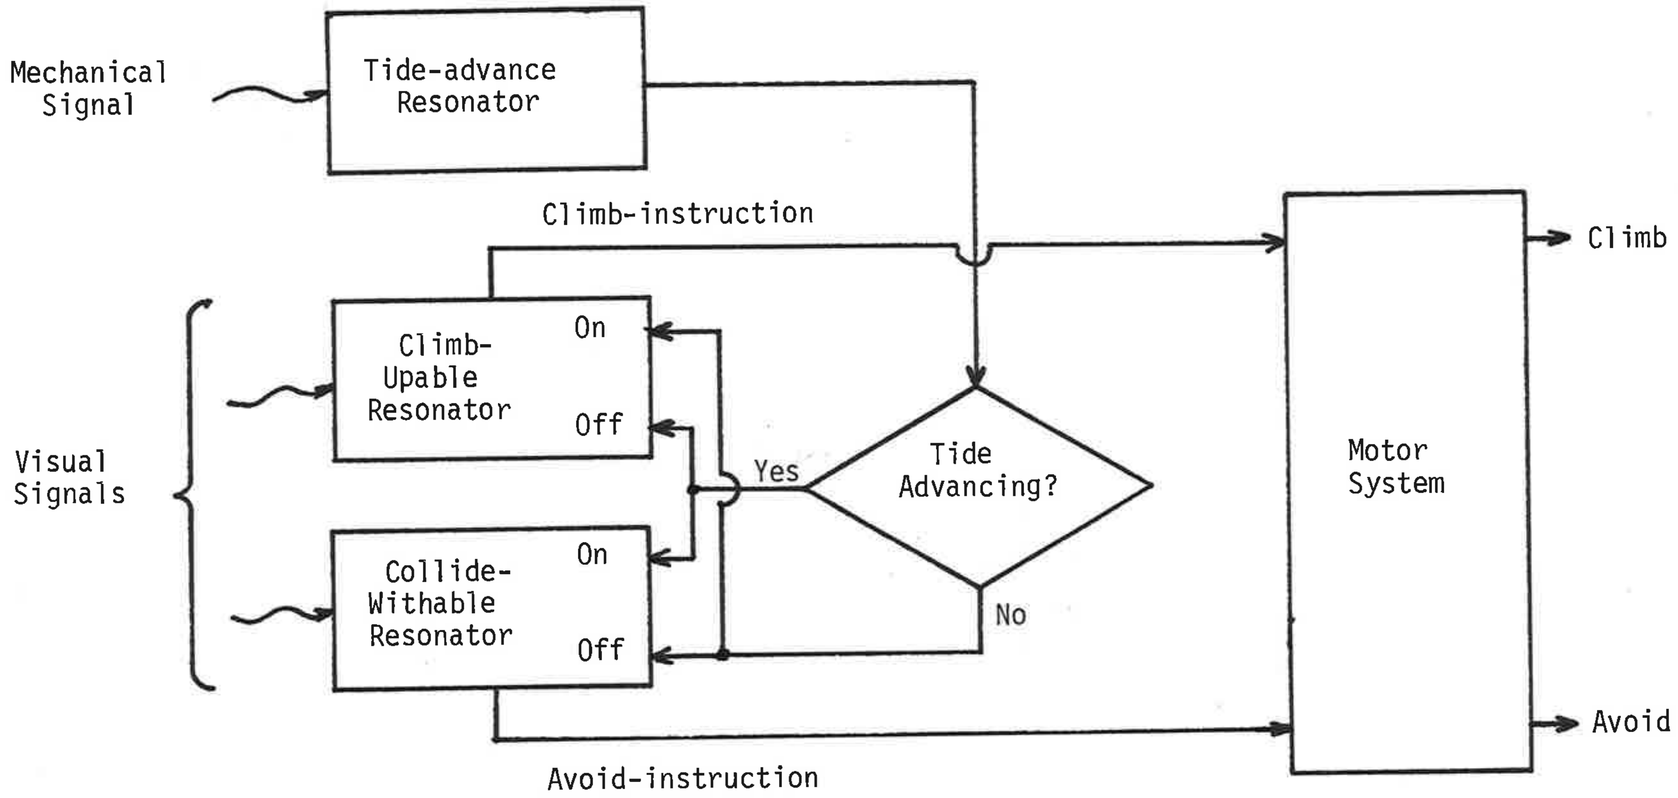

Fig. IV-2: Schematic model of the control of affordance detection in the marsh periwinkle. Outputs from resonators are assumed to be excitory.

withable" is active, and that for "climb-upable" is not. When this resonator detects signals with the property, presumably of occluding the projection of the surface of the seafloor, correlated in the periwinkle's niche with being an obstacle, the periwinkle avoids the obstacle [7].

The two visual resonators respond selectively to different properties of the same visual signal. They can, therefore, be viewed as feature detectors (cf. Lindsay and Norman, 1977). Michaels and Carello (1981) claim that ER rejects the existence of feature detectors: "the suggestion is that 'feature detectors' ... are artifacts" (p. 174). However, they do not consider the possibility that there are feature detectors for the very "higher-order" variables that they claim are directly detected.

In general, modeling organisms using schematic diagrams of the sort shown in fig. IV-2 is very much against the spirit of ER; ecological realists tend to reject talk of "internal states" altogether. Neither Michaels and Carello (1981) or Turvey et al (1981), for example, employ internal-state models of any kind. Such models of the internal states of organisms are, however, fully consistent with the fundamental assumptions of ER. Moreover, they are required if ER is to reject successfully the black-box characterization of organisms.

The model shown in fig. IV-2 assumes that the perception of affordances is direct. For each affordance to which the organism is sensitive, there is a resonator which resonates only to signals encoding that affordance. Once an affordance is detected, the behavior of the organism is determined; the organism therefore satisfies laws linking the detection of affordances to actions. Moreover, the control of resonance is law-governed. The organism is, therefore, fully stimulus driven; its behavior is not mediated by internal, inferential processes that use beliefs, etc. to infer what it should do from what it sees.

Armed with the general notion of resonators as feature detectors that can be turned on or off as the affordances that they detect are "tuned in" or not, one can consider the general question of how to characterize organisms as agents with a variable behavioral repertoire. In particular, one can consider the question of how an organism can control what Turvey et al call the "selection" of laws.

Shaw and McIntyre (1974) introduce the term, "algorist" to describe the organism as a perceiving agent, i.e. as an entity able to attend to particular features of its environment, and to act on the basis of what it perceives [8]. Michaels and Carello (1981) characterize the algorist as follows: "the first approximation to an algorist is the nonalgorithmic [in Michaels and Carello's sense: nonmechanistic] and noninformational constraints on perception...-[it] is better thought of as those aspects of the animal - the whole animal - that render certain algorithms cost-effective, certain environmental objects useful, behaviors as intentional, and so on" (p. 74-75).

The control of resonance by an organism itself, as opposed to by external factors, is clearly a function of the organism as algorist. Michaels and Carello (1981) define "attention" as the algoristic control of detection, i.e. of resonance (p. 69). Merely saying that organisms are algorists does not, however, explain either how the control of resonance works or why any particular organism controls resonance in a particular way. Once again, what is needed is an account of a mechanism. ER has, so far, failed to provide such an account.

One of the most obvious examples of an attentional behavior is exploration; when an animal explores part of its environment (e.g. visually), it looks at, i.e. attends to various things in the environment with varying degrees of interest. Exploratory behavior can often be described as goal-directed; the animal is looking for food, for shelter, for a mate, etc. Exploration therefore falls into the class of "algoristic" functions, i.e. of functions that can only be understood psychologically by considering the organism as an agent that is able to control and direct its own behavior. It therefore provides a good focus for further analysis of the problem of controlling resonance.

Visual exploration involves the coordination of behavior with what is seen. For example, if what appears in the visual field is interesting, the animal looks at it more closely, or perhaps moves so as to look at it from a different angle or range. If the object is not interesting, the animal focuses attention on something else, perhaps moving in the process.

In the language of ER, the animal is initially "set" to resonate to some affordance (the "interesting" property, e.g. edibility), i.e. to selectively respond to a property of the optic array that encodes that affordance. Motion changes the sample of the optic array to which response is possible. Motion therefore increases the sample size, and hence the probability of detecting the affordance of interest. If the affordance is detected somewhere in the environment, i.e. if the corresponding property is detected in some sample of the optic array, an affordance-specific behavior ensues. The nature of this behavior depends on the particular affordance detected, e.g. approach may be specific to the detection of prey, but flight may be specific to the detection of a predator.

The new behavior may also be accompanied by a change in the affordance to which the animal resonates. Flight, for example, requires that the organism be able to resonate to affordances such as "collide-withable" which were relatively unimportant during a period of relative calm (e.g. while exploring an area for grazable grass).

This kind of exploratory behavior can be performed by the system shown schematically in fig. IV-3. The system looks for affordance A with an A-resonator, and for A' with an A'-resonator. If it fails to detect A, it performs behavior B while continuing to look for A; if it detects A, it performs behavior B', and begins looking for A'. If it finds A', it does B''. It therefore satisfies three ecological laws: 1) If not-A, do B and look for A, 2) If A, do B' and look for A', and 3) If A', do B''. Two of these ecological laws not only relate the detection of affordances to behavior, but also relate them to the control of resonance.

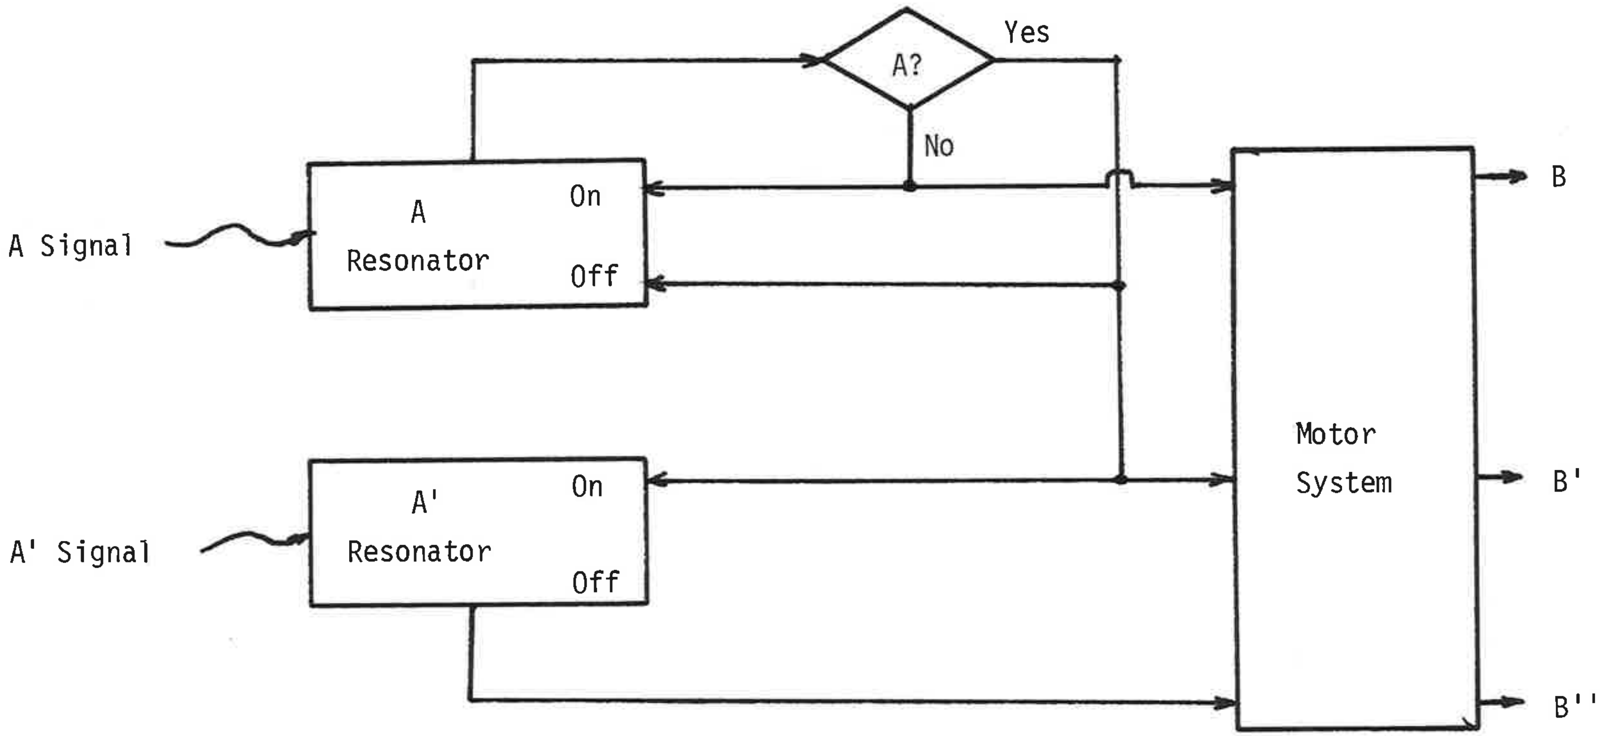

Fig. IV-3: Schematic model that generates simple exploratory behavior. The animal first does B while looking for A. When A is found, it does B' and looks for A'. If A' is found, it does B''.

In general, if ER is to explain exploration, it must either appeal to a new class of laws relating the detection of affordances to the control of detection, or it must appeal to explicit models of the organism, such as in fig. IV-3. To repeat the now-familiar criticism, the purely subsumptive strategy, which eschews functional analysis of the organism, is a dead end when it comes to explaining the interaction of affordance detection and attention. The special laws for doing this will explain nothing about how the process works. ER therefore must, if it is to explain attention, adopt the strategy of explicitly modeling the internal psychological mechanisms of the organism.

In summary, ER is faced with the problem of explaining how organisms are able to change the affordances to which they are attuned. The simplest way to account for this is to assume that organisms are equipped with resonators, e.g. neurons or neural circuits, that respond selectively to the properties of signals that encode affordances. Resonators are, therefore, feature detectors. Minimally, accounting for attunement requires the assumption that these detectors can be turned on or off, and/or the assumption that motor control can pass from one to the other. Without postulating some such mechanisms, ER cannot explain how the state of attunement of the organism can be changed or controlled.

Explanatory Power of Resonator Models

It has now been argued several times that the subsumptive explanatory strategy adopted by ER neither adequately explains the phenomena that ER attempts to explain, nor even addresses the questions of how or where any of the processes postulated by ER occur. The bulk of the preceeding discussion has, however, been an attempt to answer the latter questions, e.g. the questions of how resonance could occur, and of how it could be controlled, within the general framework of ER. The explanatory power of the answers that have been offered must, therefore, be assessed.

The model of fig. IV-1, for example, is an attempt to explain a behavioral capacity of the marsh periwinkle, the capacity to switch between a state in which plant stems are seen as climbable and a state in which they are seen as obstacles. The model appeals to a set of simpler capacities: the capacity to resonate to a property of incident light, the capacity to be turned on or off, the capacity to determine whether a signal is present in a neural connection, etc. All of these capacities are assumed to characterize neurons or groups of neurons in the marsh periwinkle.

The explanation of a complex behavioral capacity by analyzing it into a set of coordinated simple capacities is called "functional analysis" by Cummins (1975; 1983). Cummins (1983) argues that all successful explanations of complex capacities, whether in physics, chemistry, biology, or psychology, are explanations by analysis rather than by subsumption. The explananda of analytic theories are not events, but regularities between events. Ecological laws state putative nomic regularities between events of affordance detection and events of action. Ecological laws are, therefore, statements of the explananda for the analytic approach to the organism that ER must adopt if it is to have explanatory power.

Resonator models of the type shown in figs. IV-1 and IV-2 are examples of analytic explanations of the regularities described by ecological laws. The law, in each case, says what the phenomenon of interest is; the analytic model explains how the phenomenon occurs, and hence why the law holds of the organism in question. Ecological laws are tested by the experimental observation of behavior; resonator models can be tested by looking for biological subsystems that perform the simple functions postulated by the model, and by a variety of indirect behavioral tests (see also Ch. VI).

The analytic reconstruction of ER advocated here is an attempt to provide answers to the very questions that appeals to ecological laws cannot answer: the questions of how and why ecological laws hold, and the questions of how organisms can do what ER claims that they can do. If ER is to answer questions of this type, it must embrace the analytic explanatory strategy vis-a-vis the organism, and abandon the black box characterization. To this extent, it must concede a significant methodological point to cognitivism [9].

To see the power of the analytic approach, consider again the case of laboratory-induced illusions. The subsumptive approach was shown to face two problems when dealing with illusions: first, it could not provide an explanatory account of why the organism exhibits the behavior with respect to the illusory stimulus that it does; second, it could not provide an explanatory account of why the ecological laws holding in the natural setting did not hold in the laboratory setting.

The second difficulty cannot be addressed until the general issue of selection is discussed. An answer to the first difficulty can, however, be given immediately in terms of resonance. Suppose that the organism has a resonator that responds to a property of light P that is nomically correlated, in the natural environment, with an affordance A. An A-illusion is then any situation in which P is present but A is not. In such a situation, the organism will behave as if A were present, even though the ecological law that normally subsumes such behavior in the wild does not apply, since A is not present.

This phenomenon has quite a simple explanation. The organism will resonate to P if it is attuned to P and P is present. P is presented in the lab, resonance occurs, and the behavior is produced. In the laboratory setting, whether A is present is simply irrelevant. An approach that eschews reference to internal functions must forego this simple explanation.

Perceptual Error in Natural Settings

A much stronger test of the explanatory capabilities of analytic ER is provided by a problem that, thus far, has not been considered in depth by either ecological realists or their critics. This is the problem of *natural* circumstances that are "illusory" in the sense that organisms in them act as if objects had affordances that they do not, in fact, have. In such cases, organisms act in ways that are highly inappropriate given the affordances possessed by the objects in their environments.

As an example of such an error, consider the case of the unfortunate prey of the anglerfish. This is a case in which otherwise ecologically appropriate exploratory behavior is counterproductive, indeed fatal, for the exploring organism. Anglerfish are large, bottom-dwelling marine predators equipped with a long barbel with an enlarged fluorescent tip (Villwock, 1973). They typically remain stationary on the ocean floor, with the tip of the barbel dangling in front of their mouths. Small fish, the prey of the anglerfish, investigate and attack the tip of the barbel (the "bait") as if it were prey. In the course of such attacks, they come within the range of the anglerfish's jaws, and are themselves eaten.

The prey fish can be thought of as instantiating the model system shown in fig. IV-3, where A is the affordance 'is prey' and A' is the affordance 'is a predator'. B is then exploratory behavior (e.g. circling, attacking), while B' is escape behavior. The prey fish is engaged in exploration with the goal of finding food. It is, therefore, looking for food, not for predators.

The anglerfish is a successful predator because, presumably, the property of the optic array to which the prey-resonator resonates is ambiguous between specifying prey and specifying a predator. When the prey fish sees the "bait," it detects the affordance 'prey', and behaves accordingly. However, other features of the optic array, e.g. those resulting from the interaction of light with the large body and gaping jaws, etc. of the anglerfish, specify the affordance 'predator'. Why does the prey fish not detect this second encoded signal? The answer cannot be that the prey fish does not have a resonator for 'predator'; fish are generally very good predator detectors (Marshall, 1966). Assuming ER to be true, there are two possible explanations of this phenomenon. One is that, in this niche, the specification law relating the affordances 'predator' and 'prey' to detectable properties of light fails. The other is that, even though it could see the bait as signifying the presence of a predator, the prey fish does not detect 'predator' because it is not looking for it; the predator-resonator is turned off during the exploration process. In either case, the claim that signals always specify the affordances of objects to organisms is simply false. In either case, the signal can only be regarded as equivocal.

The example of the anglerfish and its prey is not unique; all cases of mimicry and protective coloration share the same properties. Explaining these phenomena appears to require, in general, the admission that signals can be equivocal in natural settings. Such explanations also appear to require feedback between affordance detection and the control of resonance of the sort shown in fig. IV-3 and hence, functional analysis of the organism.

Such misperceptions in natural settings present a serious problem for ER. They show that, for at least some organisms in some environments, perceptual signals can be equivocal. It is, therefore, simply false that organisms always perceive the affordances of their environments [10]. This problem shows quite clearly that ER must adopt functional analysis; without functional analysis, i.e. with no way of rationalizing exceptional cases, the possibility of misperception in natural settings is a straightforward counterexample to ER. For the present, it will be assumed that analytic ER can cope with this problem. The question of whether the problem amounts to a definitive argument against the entire ecological realist approach will be considered in Ch. VI.

Learning and Evolution

Given the above analysis of how resonance is controlled, one can ask how organisms come to control resonance in the ways that they do. This is part of the more general question of how organisms come to instantiate the ecological laws that they do, where an ecological law is now regarded as including not only nomic correlations between an affordance and a property of the medium and between affordance detection and behavior, but also as specifying a nomic correlation between affordance detection and the control of resonance. Can the analytic approach to ER contribute anything to the solution of this problem?

There are two obvious answers to this question of how organisms come to be able to detect affordances: evolution and learning. Both evolution and learning are claimed, in ER, to be mechanisms by which correlated changes occur in both organism and environment. The two mechanisms differ in the time required, and in the specificity of the changes that occur, but not in the final product. Both mechanisms produce "a new animal that is better able to cope with its environment" (Michaels and Carello, 1981, p. 77), and, since the organism is part of its environment, a new environment as well. While the claim that evolution and learning are examples of a single basic mechanism is somewhat counterintuitive, it is not wholly unreasonable. Indeed, this conclusion may be inferred from various systems-theoretic and thermodynamic analyses of the two processes (e.g. Bateson, 1979; Jantsch, 1980).

From the point of view of perceptual theory, both mechanisms increase the ability of animals to detect the relevant affordances of their environments; they also affect which affordances are relevant to a given animal. Michaels and Carello (1981), following Gibson (1966), characterize learning as "the education of attention" (p. 81), and claim that both evolution and learning "serve to make animals better able to detect the affordances" (p. 82).

According to ecological realists, evolution and learning are distinguished not by mechanism, but by end-product: "evolutionary learning and personal learning are thought to operate in an analogous manner" (Michaels and Carello, 1981, p. 82). Evolution is claimed to enhance only the detection of "universal" information, i.e. information that "all appropriate members of a species ... anywhere, and at any time, share the need to have the ability to detect," while learning is claimed to enhance only the detection of "local" information, i.e. information that is "unique to the particular organism- environment unit" (p. 79).

There are several problems with this approach. Prima facie, this way of distinguishing evolution and learning entails that organisms derive some of their perceptual abilities solely from the selective pressure of evolution, and derive the rest solely from individual learning. This is a strong nature-nurture distinction, and such distinctions are problematic in general (Lehrman, 1970). There are many cases in which discrete perceptual abilities develop as a result of both genetic and environmental influences. Consider, for example, the development of the visual system in the kitten, e.g. as reviewed by Lund (1978). Kittens reared in environments completely lacking features such as vertical lines will develop visual systems that are insensitive to the feature in question. Line detection ability is, therefore, not purely a case of "genetic preattunement," as Michaels and Carello call evolutionary perceptual modification (p. 78); individual experience in the rearing environment is critically important to the eventual state of attunement of the system. The ability to detect lines is, however, presumably an ability that all relevant members of the species need (e.g. for guided locomotion in a world containing vertical structures); lines are not "local" information. The nature-nurture distinction that ecological realists attempt to draw simply cannot be drawn in this case. The distinction between learning and evolution must be loosened somewhat, or else recast, to account for such cases. How to do this, however, is a matter of continuing controversy.

The identification of the subject of modifications in perceptual abilities presents a second prima facie problem. Michaels and Carello stress that both evolution and learning produce new organisms, not just smarter organisms. This leads to radically unintuitive descriptions of fairly commonplace phenomena, however. An organism cannot learn by exploration, for example; the correct description of the case must be that one organism begins an exploratory encounter, but that a different, smarter organism ends it. Any encounter that involves, or coincides with a change in resonant properties is subject to this difficulty. "Learning" is therefore best viewed as evolutionary selection at the individual level (Michaels and Carello, 1981).

How could this individual "selection" work? To answer this question, it is useful first to consider two other questions. First, what is selected, by either learning or evolution? Second, how does evolutionary selection work?

Selection, in ER, is a mechanism by which organisms are made "better able to detect the affordances" of their environments (Michaels and Carello, 1981, p. 82). Learning is the "education of attention," i.e. of the control of resonance. Evolution is, presumably, a similar process that works over a longer time-scale. These characterizations imply that selection can either change what the organism resonates to, or change how the organism's resonators are controlled. In order to answer the questions posed above, the mechanisms involved in such changes must be examined.

In the model of fig. IV-3, perception requires resonators to detect affordances, and controllers to turn these resonators on and off. There are, therefore, two questions about mechanism to answer in the case of either evolution or learning. First, where do novel resonators come from? Second, where do novel controllers come from? It is worth emphasizing that neither of these questions can even be formulated in classical, subsumptive versions of ER.

Resonators and their associated controllers are assumed to be instantiated in organisms as neural networks. Questions concerning the origin of novel neural structures are substantive biological questions. Not even the general mechanisms involved in these processes are well understood (see Raff and Kaufman, 1983, for a recent review). A few general arguments can, however, be advanced.

A possible evolutionary mechanism for the generation of novel resonators and control circuits is the duplication of an existing structure followed by divergence of the two daughter structures (Fields, 1983a). Thus, one "ancestral" resonator could, by means of an "error" during development, duplicate into two structures, each of which took on different functions. A second, presumably more difficult mechanism is the outright conversion of a previously-existing structure with some other function into a resonator

Neither of these mechanisms are, however, very plausible in the case of learning. Learning is generally distinguished from maturation, i.e. biological development. While this distinction is not completely clear at the neuronal level of description, a reasonable rule of thumb is that learning amounts to fine-tuning, i.e. that all large scale changes in brain structure are developmental. This is reasonable in light of the fact that the general architecture of the brain is fixed, and all neural reproduction and growth occurs, during development (Lund, 1977). It is, therefore, implausible that structures undergo radical changes in function, e.g. from something else to a resonator, let alone duplication and divergence, after development has stopped.

These two mechanisms are, however, the only obvious possibilities for generating new resonators. Among the prima facie possibilities, therefore, this argument from plausibility leaves only one: the possibility that learning does not produce any new resonators, i.e. that learning only affects the control of existing resonators. This hypothesis has, as a consequence, the claim that all of the affordances for organisms of a given type are fixed evolutionarily, i.e. that <u>no affordances-detecting capabilities are learned</u>.

The major difficulty for ER arising from the problem of illusions, and from the Fodor-Pylyshyn constraint argument, is that there is no obvious way of saying how learning constrains ecological laws. The hypothesis that no affordance-detecting capacities are learned removes at least most of this difficulty. All that remains to be explained is how the capacities for the control of resonance could be learned. Changes in the control of resonance, however, are not changes in the organism's behavioral repertoire, i.e. they are not changes in the set of ecological laws that the organism can satisfy. Such changes are only changes in the ecological laws that the organism satisfies under the particular occurent set of environmental conditions. From the point of view of the conflict with cognitivism, therefore, the hypothesis that no affordances are learned is quite attractive.

The hypothesis that no affordances are learned also solves the problem of distinguishing learning and evolution in a way that does not require a principle nature-nurture distinction. The role of learning can be confined to fine-tuning controllers; everything else is the business of evolution and maturation. It is, as noted above, at least plausible that a sensible theory of these processes will be forthcoming.

In summary, ER can either accept the hypothesis that no affordances are learned, which largely solves the prima facie problems associated with learning, or else provide an account of learning that solves the problems some other way. If the first route is taken, ecological realists must be prepared to argue that any affordance to which they appeal has been selected evolutionarily. If they take the second route, however, they must produce a new theory from a field in which there are no obvious candidates.

Inference in ER

Ecological realists, as shown in Ch. II, categorically reject the notion of inference as a feature of perception. The rejection of inference is based on the claim that proximal stimuli are not impoverished; therefore, their significance does not have to be inferred from data stored in memory. This claim can now be re-examined in light of the theory of resonators that has been developed in this chapter.

Resonators detect information by resonating to a modulation in a signal that encodes an affordance. As shown above, resonators causally encode the information that the affordance in question is specified by the modulation to which they selectively respond. They do not, however, take any propositions from stored databases as additional input.

Inference is traditionally viewed as a process that takes stored propositions as input, and produces more stored propositions as output (e.g. Pylyshyn, 1984). This will be called the "narrow sense" of inference. In the narrow sense, resonators do not carry out inferences, for they do not take stored propositions as input.

One can, however, also speak of a process as inferential in a considerably wider sense. Consider any process P with a characterized initial state A and final state B. Suppose that the states A and B can be given propositional interpretations I(A) and I(B). These interpretations may, for example, just be descriptions of A and B in some vocabulary. The process P will be called "inferential in the wide sense" if, for any allowed I/O pair A and B, I(B) can be inferred from I(A).

It is apparent that all manner of systems can be described as "inferential" in this wide sense [11]. It is, moreover, apparent that describing events such as the fall of a stone by appeal to the inferential relation between "the stone is released in a suitable gravitational field" and "the stones falls freely" will not add significantly to our understanding of stones. What, then, is the use in characterizing systems as inferential in this sense?

As Cummins (1983, pp. 40-44) points out, the critical question regarding such interpretive analyses is under what conditions it is useful, or even essential, for a particular explanatory task. The claim that a system carries out inferences does not itself explain the system's behavior; it rather formulates an explanadum for an analytic account of the mechanisms involved in the system's functioning. In the case of the stone, a mechanismtic explanation is readily available; in many other cases, and in particular, in essentially all cases of interest to psychology, mechanistic explanations are not currently available.

The process of resonance is inferential in the wide sense. Consider, for example, an optic array with a large number of properties encoding a large number of affordances. From the point of view of a system attempting to detect an affordance A, the modulation encoding A is the only one of interest. All of the other modulations are noise [12]. The system can only extract the information specifying A by separating the modulation that encodes it from the noise. How are we to explain the ability of the system to do this?

To perform this task, the resonator uses information, namely, the information that the particular modulation encodes the affordance, together with its ability to selectively resonate. As pointed out above, the system uses this informational content in the sense that, if it did not have the information, it could not do its job. The resonator can, therefore, be thought of as encoding the proposition "Signal S specifies A," whereas the medium can be thought of as encoding "S here now" when it is transmitting S. The output of the resonator then encodes "A here now," which is used by the rest of the system. "S specifies A" together with "S" entails "A". Resonance is, therefore, inferential in the wide sense.

This claim, like the earlier claims concerning information storage by resonators, is completely independent of the question of the impoverishment of proximal stimuli. Therefore, the claim that stimuli are not impoverished is no cause to reject inference in the wide sense as a description of what resonators do. Indeed, resonators must use information, and hence must carry out inferences in the wide sense, precisely because proximal stimuli are so rich. If every proximal stimulus encoded only one affordance instead of many, this type of inference, and indeed resonance, would not be necessary.

Summary of Analytic ER

This chapter has shown that traditional, subsumptive ER, the theory outlined in Ch. II, makes no attempt, and in fact, can make no attempt, to explain the mechanisms that account for how or why organisms do what they do. This theory is, therefore, quite unsatisfactory. It has, however, also shown that, if adherence to the subsumptive strategy is abandoned in favor of the analytic strategy, a theory with considerable explanatory power results.

The principle explanatory construct of analytic ER is the resonator. Resonators are systems that resonate to modulations in perceptual media that encode affordances. Resonators store information, and can be viewed as carrying out inferences in the wide, even if not the narrow, sense. Resonators are controlled by being turned on or off; this control is accomplished by controllers that provide causal links between events of detection by different resonators.

The theory of resonators allows explanations of various behavior patterns to be constructed. In particular, it allows the explanation of the ability to perceive the same object as having two different affordances (fig. IV-2), and allows the explanation of exploratory behavior (fig. IV-3). These explanations entail specific predictions concerning functional neural architecture; in particular, they entail the prediction that organisms possess resonators for each affordance to which they are sensitive, and the prediction that these resonators are connected by particular controllers. As will be discussed below, these are predictions that could be tested empirically.

Analytic ER, therefore, has far greater explanatory power then standard, subsumptive ER. It is, however, much closer, both methodologically and factually, to cognitivism. In particular, analytic ER agrees with cognitivism that the internal states of organisms are important, and that perception involves inference. It disagrees with cognitivism, however, in its claims that proximal stimuli are not impoverished, that resonance is the mechanism of affordance detection, and that affordance-detecting capacities cannot be learned. The strength of these disagreements will be assessed in Ch. VI, in which the fortunes of analytic ER in the ER-cognitivism debate will be considered.

NOTES - CHAPTER IV.

1. It is perhaps more accurate to think of ecological laws as specifying a relation between the perception of an affordance and the conditional probabilities for behaviors. Whether this refined notion of an ecological law is adopted, however, makes no difference to the present discussion.

2. These organisms are, of course, quite sensitive to the variations in certain environmental parameters, such as temperature, dissolved oxygen level, the presence of toxins, the availablity of particular kinds of food, etc. Within these constraints, however, a fairly wide variety of environmental situations may be equivalent. In general, whether two niches can be considered equivalent depends strongly on the explanatory task at hand. For example, two niches that differ only in one species may be equivalent if one is considering an organism that has no interactions with the species that is missing in one of the niches, but would certainly not be equivalent if one were considering, e.g. the prey of the missing species. Decisions about equivalence, therefore, require at least some knowledge of the ecological laws being investigated. They are, therefore, ultimately circular.

3. Another way of putting this is simply to say that, at least in the case of higher organisms, what are called "ecological laws" are not laws at all. The motivation for doing this is that they have no generality. Ecological laws are, however, law-like: the statement "X affords A to O" is (presumably) counterfactual-supporting; similarly, the statement "O will carry out B when it perceives A" is (presumably) counterfactual-supporting (Turvey et al (1981) discuss the problem of counterfactual support in detail). The terminology "ecological law" will be retained in the present discussion to avoid the appearance of begging the question of whether such claims are lawlike against ER; not even ecological realists make the opposite mistake of taking ecological laws to be universal.

4. It is important to note that this is true independently of whether laws are viewed as universal generalizations over particulars (the traditional, Logical Empiricist view) or as nomic relations between properties (Dretske's (1977) view, adopted by Turvey et al (1981). A nomic relation between properties cannot, by itself, explain why it is true, or how it is instantiated in any particular case.

5. Grossberg (1980), for example, shows that the on-center/off-

surround and off-center/on- surround systems found in the retina, superior colliculus, and primary visual cortex can be described in terms of resonance. He advances the theory that these systems developed evolutionarily as a solution to the problem of enhancing differences in nearly-random stimuli by means of resonant self-organization, i.e. the thermodynamic tendency for resonances in energetically- open systems to be auto-reinforcing (the central idea of this theory is that of positive feedback, except applied to the structure in question, not to its excitation; see Nicolis and Prigogine, 1979; Gierer, 1981 for general discussions).

6. There is, of course, one sense in which the external world does change, e.g. when someone learns to read: relative to *that person*, the world has a new affordance. Michaels and Carello's (1981) description of learning as a process that creates "a new animal" and, ipso facto, a new niche (p. 80) suggests that they would adopt this analysis (see also below). This "subjective" picture of the control of resonance is, however, explanatorily bankrupt. The question of *how* the state of attunement changes is left unanswered. This point will be discussed in greater depth when the ER analysis of learning is considered.

7. This model is clearly oversimplified. Presumably the resonator for "collide-withable" is used to guide approach to climbable things even after their climability is detected. This complication is, however, irrelevant to the present discussion. For a description of actual neural circuitry in *Aplysia*, see Kandel, 1979; Hawkins and Kandel (1984).

8. This terminology, which derives from Shaw and McIntyre's (1974) interest in action theory, is particularly unfortunate, as it invites comparisons of the algorist with the "ghost in the machine" of Cartesian theories (e.g. the soul in Eccles, 1980). It is introduced here for completeness' sake.

9. Analysis cannot, of course, go on forever. At some point in any analytic explanation, a level of description is reached at which appeals must be made to laws or brute observational facts. This level may, however, be well below that at which phenomena are considered psychologically interesting.

10. An especially committed ecological realist might still contest this argument, claiming that, as affordance perception takes time (e.g. Michaels and Carello, 1981), the proper analysis of the anglerfish case is that the prey fish is "still looking" when it is eaten. This is doubtless true. The problem is that the fish acts as if the bait affords, if not "prey", at least "worthy of further investigation." The fish dies investigating. In fact, the bait affords "get out of here as fast as possible," i.e. the message that a predator is hidden nearby. The fish does not perceive this affordance of the bait. What it does perceive, what it must

perceive, if its behavior is to be understood, is something like "worthy of further investigation." This affordance, however, the anglerfish clearly does not have.

11. Assuming a standard model theory, any system with a self-consistent I/O relation can be described as inferential in the wide sense.

12. Compare this situation with that of attempting to hold a conversation in a crowded party. In the latter case, all of the other sound in the room, even though it encodes information that could be of interest, is noise for the purposes of holding the conversation.

CHAPTER V

COGNITIVISM RECONSTRUCTED

Chapter II posed two challenges for cognitivism. First, cognitivists must show that mentalistic models are not arbitrary; i.e. they must show that there is some way to determine experimentally whether a given cognitivist model is true. Second, cognitivists must show that the recognition regress argument is unsound; i.e. they must show either that inferences can occur in the absence of stored concepts, or that inferences can use stored concepts in a way that does not lead to an infinite regress. This chapter proposes a reconstruction of cognitivism that meets these challenges.

The constraint problem in cognitivism is considered first. As shown in Ch. II, cognitivism entails the claim that cognition is computation. Thinking, for cognitivists, is instantiating a formal symbol manipulator, i.e. a program of some sort (e.g. Newell, 1980). Cognitivist models of mental processes are programs that can be instantiated on artificial computers. Solving the constraint problem, therefore, requires being able to say when a computer program duplicates a human cognitive process. Pylyshyn's (1980; 1984) cognitive penetrability criterion, introduced in Ch. II, is an attempt to do this.

In order to assess the success of Pylyshyn's criterion in meeting this challenge, one must be able to say clearly what the criterion is, and how it is to be applied in practice. The first part of the chapter, therefore, presents the formal tools, the notions of program instantiation, functional architecture, and virtual computation, necessary to answer this question. The theory of early visual processing developed by Marr (1981) is then discussed in terms of these notions. A theory of vision in the marsh periwinkle, based on Marr's theory of human vision, that explains the behavior of the snail with respect to plant stems is then constructed.

Once the necessary tools, and their applications, have been described, the cognitive penetrability criterion itself is considered. It is shown that the criterion can provide a constraint on cognitivist explanations only if a proprietary level of description is assumed in advance. The criterion does not, therefore, provide an empirical constraint; it rather highlights a set of phenomena that a cognitivist model must explain. Fodor's notion of informational encapsulation, also introduced in Ch. II, is shown to provide a general way of explaining cognitive penetration.

The recognition regress argument is then considered. It is shown that the this argument does, in fact, tacitly assume that all cognitivist models are internal manual models in Cummins' (1982b; 1983) sense. Cognitivist models can, however, always be constructed that do not involve internal manuals. The recognition regress argument is, therefore, unsound when applied to cognitivism in general.

The two principal ecological realist arguments against cognitivism having been dispensed with, the question of the role of appeals to inference in cognitivist models is reopened. It is shown that whether it is useful to characterize a process as inferential depends critically on the explanatory task at hand, and that the decision whether to employ inferential characterizations must be made prior to the construction of any cognitive model. The consequences of this fact for cognitivist methodology are then considered. Cognitivist models that incorporate this characterization of inference are compared with analytic ER models in Ch. VI.

Strong Equivalence

Block and Fodor (1972) argued, as part of their general attack on the then-popular behavioral, physiological, and functional-state identity theories of mental states, that more than a mere description of the input-output properties of an organism was required in order to type-identify mental states [1]. The argument consists in pointing out that a single I/O relation can be satisfied by any number of devices progressing through any number of sets of internal states. This is a formal result of automata theory, that had been used earlier by Nelson (1969) to argue against behaviorism in general.

Pylyshyn (1980; 1984) has recently revived this argument to support his claim that a description of the I/O relation of a device is insufficient to type-identify the computational processes that it instantiates. For example, multiplication can be performed by two different methods, successive addition and partial products; similarly, theorems in propositional calculus can be proven using different methods, e.g. truth tables and natural deduction. In each case, the I/O relations are the same; however, the methods, i.e. the algorithms used, are quite different.

Cognitive science asks not only what the I/O relations characterizing psychological processes are, but what the processes themselves are. In particular, cognitive science must ask, and answer, whether the processes implemented in a program constructed on the basis of experimental data are the processes carried out by the cognitive systems of the subjects who generated the data. If the program has a different I/O relation from the subjects when confronted with the same tasks, the theory that the program embodies must be wrong. If, however, the program and the subjects have the same I/O relation, the possibility remains open that they employ radically different cognitive processes. In such a case, the theory, which attempts to specify the processes used by the subjects, is still wrong. The nature of cognitive processes thus cannot, as Pylyshyn correctly points out, be determined from their I/O relations alone. Cognitive science must, therefore, have available some alternative set of distinctions for describing cognitive processes that allows them to be type identified, i.e. that allows a distinction to be drawn between different, but I/O equivalent processes.

Pylyshyn (1980; 1984) suggests that the proper criterion of similarity of cognitive processes is sameness of instantiated algorithm. Thus, if it can be shown that a theoretical model employs the same algorithm as a subject, one can say that the two are equivalent, and that the theory is correct. Pylyshyn calls this criterion the "strong equivalence" criterion. Two processes are, therefore, strongly equivalent in Pylyshyn's sense if and only if they can be described as instantiations of the same algorithm.

Pylyshyn does not, however, merely want a notion of similarity that distinguishes programs that instantiate the same algorithm from those that do not. He also wants the criterion of similarity to be incorporated into an empirical criterion that can be used to distinguish experimentally between systems that instantiate the same algorithm and those that do not. Strong equivalence is, therefore, a useful notion only if there is some way of telling whether two processes are instantiations of the same algorithm. In the case of computer programs, one can attempt to determine the algorithm used by looking at a specification of the program, e.g. a flowchart, or a printout of the code in some programming language. As Pylyshyn (1984, Ch. 5) points out, this is not always an easy task. If, however, it is not always an easy task when the program is written down, determining which algorithm a device instantiates is quite a formidable task when no linguistic representation of the program is available. This, however, is precisely the position of the experimental cognitive scientist.

Given this fact, it is clear that an experimental method of determining which algorithm a device instantiates would be very useful indeed. The cognitive penetrability criterion, as will be shown below, is meant to be at least an important component of such a method. Before discussing this criterion in detail, however, it is necessary to introduce the needed vocabulary; in particular, it is necessary to introduce the notion of a _functional architecture_.

Functional Architecture

Pylyshyn (1984) characterizes strongly equivalent programs thus: "two programs can be thought of as strongly equivalent or as different realizations of the same algorithm or the same cognitive process if they can be represented by the same program in some theoretically-specified virtual machine" (p. 91). Strong equivalence is thus characterized in terms of the notion of a virtual machine. A virtual machine defines a level of description of a computational process: "the formal structure of the virtual machine - or what I call its _functional architecture_ - thus represents the theoretical definition of, for example, the right level of specificity (or level of aggregation) at which to view mental processes" (Pylyshyn, 1984, p. 92, his italics). It is this ability to pick out a level of description, characterized by a functional architecture, that forms the basis for Pylyshyn's discussion of cognitive penetration. In order to assess the usefulness of the cognitive penetrability criterion, it is necessary to understand how this level of description is picked out. Doing so, however, requires understanding the notion of a virtual machine.

Consider a general-purpose computer, e.g. a VAX or an IBM-PC. Such a device has a certain hardware, composed of components that can be in various physical states (e.g. transistors that can be passing current or not, flipflops that can be "on" or not, etc.). When a computer runs a program, its hardware starts in one of the possible physical states, progresses through some set of intermediate states, and stops in some final state. The initial state is interpreted by the user as representing the input to the program, while the final state is interpreted as representing the output of the program. The intermediate states are interpreted as representing intermediate steps in the computation of the output from the input (Cummins, 1983) [2].

General-purpose computers are designed with many criteria in mind. Among them are that such machines must be versatile, i.e. able to run many different types of programs, and that they must be understandable, at some level of description, by their users. Computers meet these requirements through the use of programming languages. A programming language provides a way of writing programs using expressions that are about, e.g. arithmetic operations, sets, numbers, objects etc., as opposed to expressions that are about the states of the hardware. Programming languages thus allow programmers to avoid the question of interpreting the states of the hardware directly; all a typical programmer generally has to interpret are the referents of the names and the values of the variables contained in the program.

As might be expected, this massive reduction in interpretive effort is not free. Something still must interpret hardware states as, e.g. numerals, plus signs, alphabetic characters, etc. This job is taken over, in modern computers, by a set of additional programs.

This solution of the interpretation problem is, prima facie, regressive. If programs can only run because other programs interpret machine states, the way is open for an infinite number of programs to be necessary in order to run one program. This is avoided, in actual computer systems, by arranging programs hierarchically. "Low-level" programs, written in machine language, interpret states of the hardware. These low-level programs are, in turn, interpreted by higher-level programs, which are themselves interpreted by even higher-level programs, written in programming languages, which are interpreted by users.

The situation is thus as shown in fig. V-1. The states of the hardware are interpreted by progressively higher-level programs, each of which directly interprets the program below it. Each program is written in a language, the terms of which refer to the level below it. The highest-level program is interpreted by a user.

Each level in fig. V-1 defines a virtual machine (e.g. Shaw, 1974, Ch. 1; Rus, 1979, Ch. III). A machine language program, for example, defines a virtual machine for an intermediate program; the latter interprets the former in the same way that the former interprets the hardware. Cummins (1983) calls sequences of state transitions in virtual machines memory paths (p. 178 ff). A memory path in a virtual machine corresponds to a sequence of state transitions that occurs in the hardware when the program defining the virtual machine is executed.

When a high-level (i.e. non-machine language) program P is

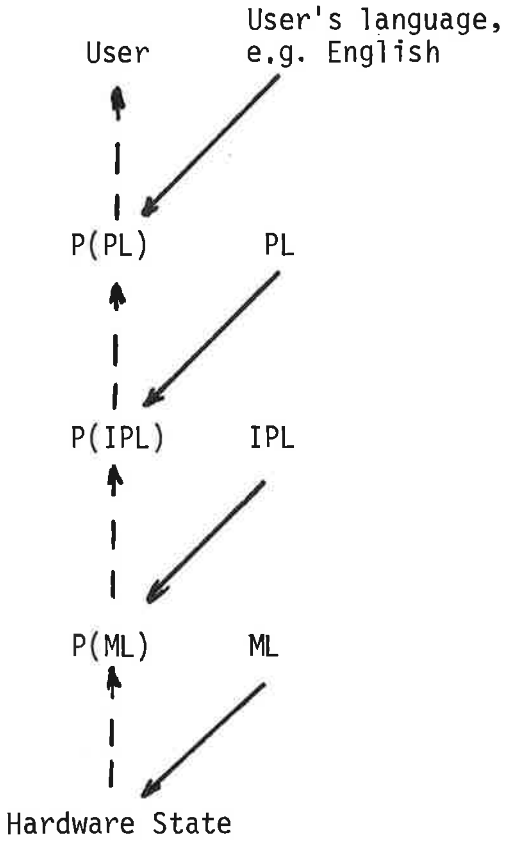

Fig. V-1: Hierarchy of computer languages: ML = Machine Language, IPL = Intermediate Programming Language (e.g. an assembler), PL = Programming Language (e.g. LISP). P(L) signifies a program in language L. Solid arrows represent linguistic reference; dashed arrows represent the relation "interpreted by."

executed with a certain input I, the hardware goes though a series of state transitions that can be interpreted, in the language L of the program, as corresponding to a memory path in the virtual machine VM defined by the program. In other words, the state transitions of the hardware, when viewed under the interpretation defining VM, imitates VM executing P for input I. High-level program execution, or virtual execution, is always imitation in this sense.

This imitation is principled, i.e. non-arbitrary, because each interpretation map in fig. V-1, except the user's interpretation of the "source" code P(PL), is fixed by the design of the languages ML, IPL, and PL. In particular, ML (i.e. the hardware itself) is designed so that, for each program P(ML) and its input I, the state of the hardware that initiates the sequence of state transitions that imitates V(P(ML)) executing P(ML) on I represents, in ML, the text of the program P(ML) together with the text of the input I. In general, at any level, the initial state for a memory path in the virtual machine is a representation of the program-input pair that interprets that memory path at the next higher level.

At each level, exactly one program, the "resident program" at that level, is actually executed. This program can be represented by a set of productions, (if S1, go to Si; if S2, go to Sj; ...; if Sk, go to S1; ...), whose members specify all of the computationally relevant allowed state transitions of the virtual machine at that level. Whatever the machine does, it executes this program. Episodes of execution correspond to memory paths, i.e. to specific sets of transitions, that are interpreted at higher levels as the "execution" of higher-level programs; i.e. episodes of execution of its resident program by a virtual machine imitate other, higher-level virtual machines.

The resident program at a level is what Pylyshyn calls the "functional architecture" at that level. From the preceeding, several facts about functional architectures can be derived. First, each level in the computational hierarchy corresponds to exactly one functional architecture. The hardware level corresponds to the hardware architecture. Virtual machines correspond to virtual architectures. Second, the functional architecture, or resident program, is never represented at its own level. It can, however, be represented at the next lower (as an initial state) or higher (as a name) level. Third, the functional architecture at a given level imitates higher-level functional architectures, and is imitated by lower-level functional architectures. The functional architecture at a given level may imitate other functional architectures, but never executes any other programs. It merely virtually executes them.

Given the notions of functional architecture and virtual execution, the explanatory structure, and assumptions, of cognitivist theories can be described in detail. Marr's theory of early visual processing provides a case in point. The next section describes this theory, and extends it to include a description of affordance perception. Cognitive penetration can then be examined within the context of this theory.

Marr's Theory of Early Vision

The theory of early visual processing in mammals developed by Marr and his colleagues, and described in Marr (1981), is perhaps the most detailed computational theory of vision to be produced to date (see also Marr, 1982; Poggio, 1982; Poggio and Fahle, 1983; Ullman, 1983). Marr's theory describes both the processes leading to the formation of perceptual images ("early vision"), and the processes involved in the identification of objects based on the information present in images. This theory can be described using the notions developed in the last section.

The functions performed by the early visual system, according to Marr's model, are shown schematically in fig. V-2. The eyes detect samples of a structured optic array that encodes the locations, shapes, colors, etc. of objects. Signals from these receptors are subjected to contrast enhancement to emphasize edges, color boundaries, and other discontinuities. This first stage of processing occurs in the retina and the superior colliculus of mammals. Its product is the "primal sketch," a 2-dimensional representation of the world in terms of intensity and color gradients [3]. The primal sketches of the visual fields of the two eyes are combined by the process of binocular fusion. This process, which effectively adds depth information to the two primal sketches, produces a "2 1/2-d" sketch. The 2 1/2-d sketch shows the surfaces of objects that face the viewer, and contains information specifying color, distance from the viewer, etc. The 2 1/2-d sketch is produced in

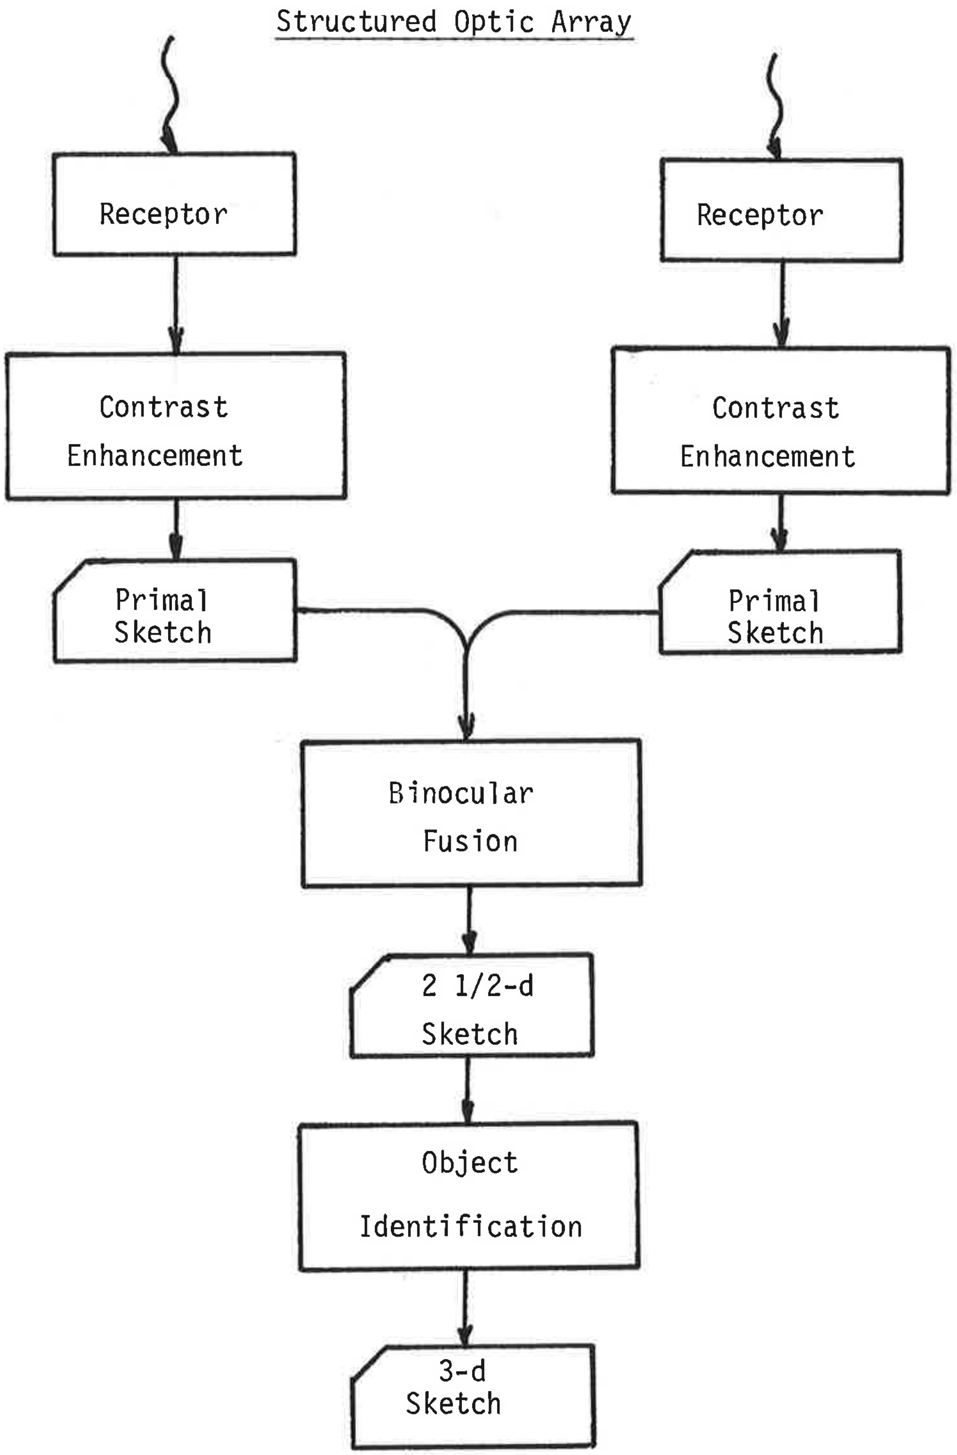

Fig. V-2: Functions performed by, and representations produced by, the early visual system in Marr's (1981) theory.

the primary visual cortex, where edge and angle detection, and stereopsis occur. This representation serves as input to the object recognition process, which presumably occurs in the association cortex (e.g. damage to the parietal cortex causes deficiencies in object recognition (Warrington and Taylor, 1973)). Object recognition uses information from memory about the shapes of typical things in the environment to construct a 3-d representation of the positions of objects in the scene being perceived.

The functions required at each stage of processing proposed by Marr's theory have been described mathematically in considerable detail (Marr, 1981). Moreover, many of these functions have been implemented in computer models. Marr's theory is, therefore, a serious contender, and perhaps the best contender, in the field of computational models of perception.

At least to the level of the 2 1/2-d sketch, which Marr (1981) characterizes as "the end, perhaps, of pure perception" (p. 268), Marr's theory amounts to a direct computational interpretation of events at the neuronal level of description. While Marr makes no attempt to model the behavior of individual neurons, the theory does attempt to model precisely the behavior of functional groups of neurons such as the retina. This attempt is highly successful in the case of the early stages in processing (see, e.g. the theory-data comparisons for retinal ganglion cells shown in Marr, 1981, fig. 2-17, p. 65). Marr attempts, therefore, to construct a program in a language quite close to the machine language of the brain that duplicates the functions involved in early vision.

The neural structures and processes involved in object recognition are much less well known than those involved in early vision (e.g. Ottoson, 1983). Marr's description of these processes is, therefore, given at a somewhat higher level. Marr (1981; 1982) assumes that object recognition is based on pattern matching between canonical elementary shapes stored in a "library" and the shapes specified by the 3-d, and on the front surface, by the 2 1/2-d sketches. This matching process involves "the deployment of gradually more detailed stored 3-d models during the process of recognition-derivation" (Marr, 1981, p. 327). This process is progressive, since more and more detailed models are used, and hierarchical, since the models used initially specify only the large-scale features of objects, while the later models specify smaller-scale features. Hypothesis generation and testing occurs at each stage until a satisfactory match is obtained (e.g. as described by Rock, 1983). The 3-d sketch is, therefore, generated from the 2 1/2-d sketch by a process that requires successive iterations, with knowledge input at each iteration.

Marr's description of the processes involved in object recognition is not in machine language; it makes no mention of specific neural structures. It is, rather, in a language, the terms of which refer to interpretations of neural structures at, perhaps, several removes. For example, the program interprets the structures in Area 17 of the primary visual cortex as feature detectors. Marr's object recognition program, therefore, defines a virtual machine that groups of neurons in the association cortex imitate.

Although Marr limits himself to the consideration of shape perception, one can speculate that his theory could be extended to account for affordance perception as well. Presumably, in a theory such as Marr's, higher organisms, at least, would infer the affordances of perceived objects from general knowledge of the characteristics of objects. One can imagine an affordance-recognition system scanning the 3-d sketch in much the same way that the object recognition system scanned the 2 1/2-d sketch, and inferring the affordances of objects by a process of successive approximation.

As in ER, however, it is simpler to consider affordance perception in the case of lower animals, where recognition of an affordance, by whatever mechanism, leads to the production of a fixed action. In such organisms, one can imagine that light is detected, and subjected to some minimal amount of processing to improve contrast and extract depth information. Such processing would produce an analog of a 2 1/2-d sketch. This sketch would serve as input to a set of feature detectors, each of which would respond only to sketches that contained features indicating the presence of an object of a particular type. For example, in the case of the marsh periwinkle, one might imagine that it had a feature detector that responded preferentially to things shaped like plant stems. The detection of certain combinations of objects could then be assumed to trigger certain fixed behavior patterns, such as climbing or avoiding.

Figure V-3 shows a model of the marsh periwinkle based on this reasoning. The periwinkle is assumed to have two transducers,

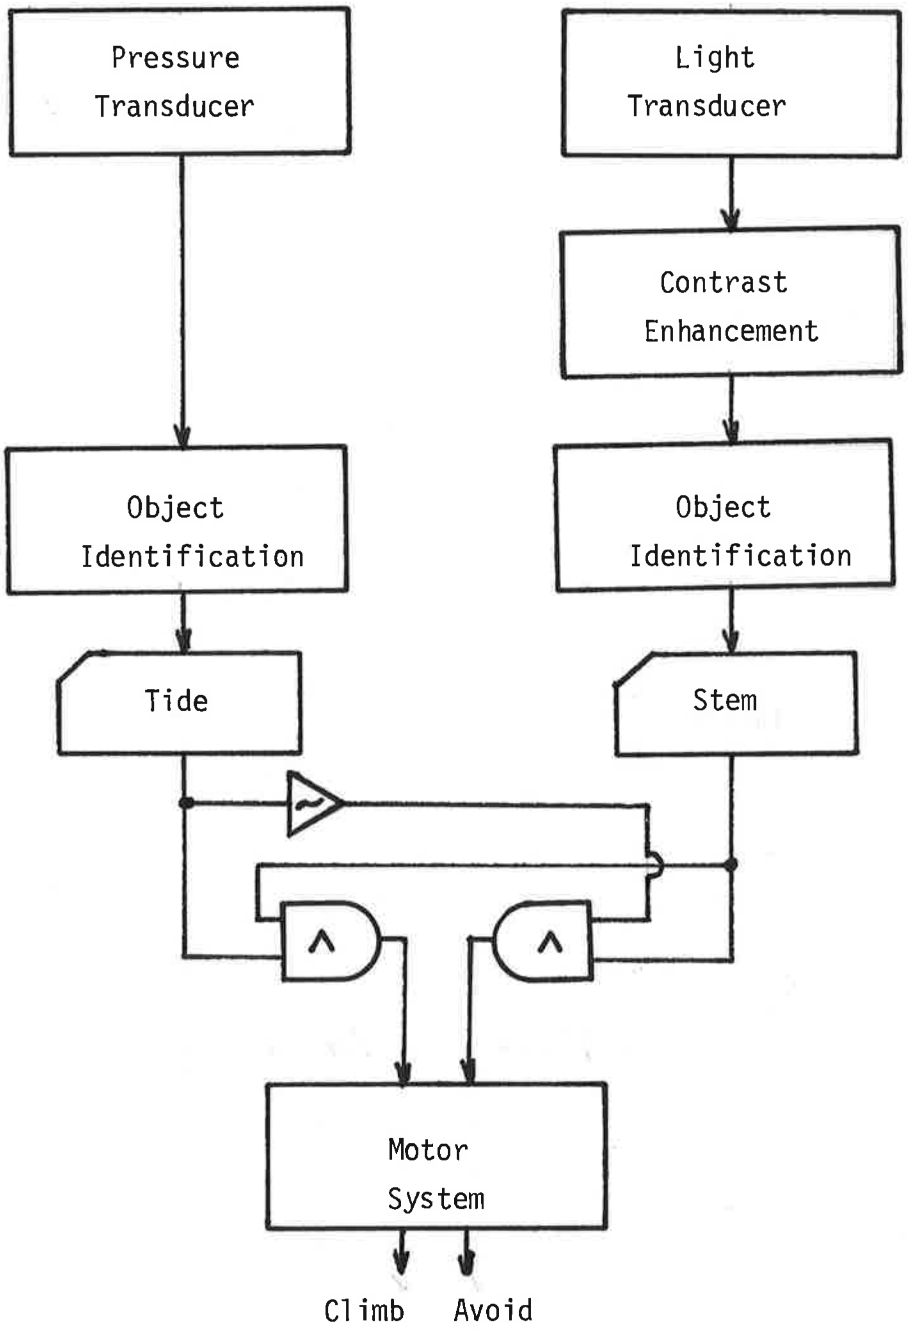

Fig. V-3: Schematic model of the marsh periwinkle's climb/avoid behavior based on the model of Marr (1981).

one for light, and the other for pressure, and two object identifiers, one for identifying the advance of the tide, and the other for identifying plant stems. The behavior of the periwinkle then depends in a simple way on whether it identifies a stem with or without also detecting the advance of the tide.

The functional units shown in figs. V-2 and V-3 are, in Marr's theory, informationally encapsulated in Fodor's (1983) sense. That is, each unit has access to only a particular knowledge base. In the system shown in fig. V-2, for example, contrast enhancement is carried out by an input module that has access only to the input signal, and to causally-encoded information specifying the functions that are to be applied to the input (described by Marr, 1981; Grossberg, 1980). This module has no access to, e.g. information describing the functions used for binocular fusion. The functional unit that carries out binocular fusion, on the other hand, has access to this latter information, again in a causal encoding, but has no access to information about contrast enhancement. In general, the functional units composing the perceptual system in Marr's theory are organized as shown in fig. V-4. Information flows from left to right, and from databases to processing units only. Data transfer from right to left, or from processors to databases is forbidden, except "within boxes" [4].

Marr's theory of perception, as applied to both lower animals and mammals, provides an example of a cognitivist theory of perception. The cognitive penetrability criterion, and the recognition regress argument, can now be discussed with respect to this

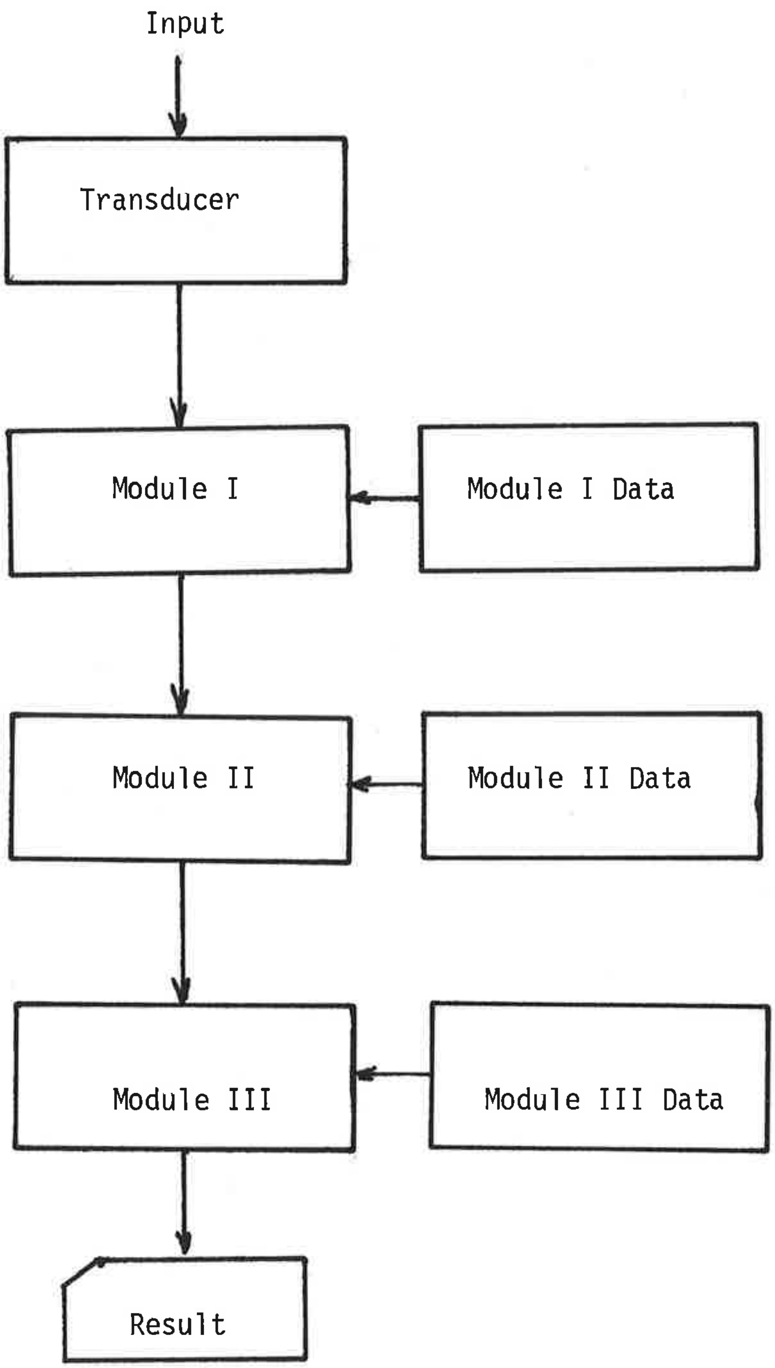

Fig. V-4: Modularity of functional units in Marr's (1981) model of early vision. No information is transferred in the direction opposite that of the arrows, or from any unit to its respective database.

theoretical model.

Cognitive Penetrability

Cognitive penetration, defined by Pylyshyn (1984) as "the rationally explicable alterability of a component's [i.e. cognitive function's] behavior in response to changes in goals and beliefs" (p. 133), appears to occur in humans after the construction of the 2 1/2-d sketch, but before the construction of the 3-d sketch. It therefore appears to occur during the process of object identification. This is by no means surprising; humans can clearly learn what new things are, or that things are not what they were originally taken to be, but it is not at all clear that things ever appear different as a result of learning (Fodor, 1984).

Pylyshyn (1984) is careful to distinguish cognitive penetrations, which he classes as phenomena that "must be explained by appeal to semantically interpreted representations" (p. 130), from other sorts of phenomena that may be indistinguishable, by their effects alone, from cognitive penetrations. For example, hitting someone on the head may cause them to report seeing stars, or seeing a bright flash, but this phenomenon does not count as "cognitive" penetration. The effect, in this case, is not mediated by the subject's beliefs, desires, tacit knowledge, etc. True cognitive penetration only occurs in cases in which the I/O relation of a cognitive process is changed due to a change in the contents of the subject's mental states.

Pylyshyn introduces the notion of cognitive penetration in order to use it to develop an empirical criterion for distinguishing what he calls "cognitive" processes, processes that require appeals to "semantically interpreted representations" for their explanation, from other kinds of processes. His argument runs roughly as follows. First, he introduces the notion of a functional architecture, claiming: "processes carried out in the functional architecture are processes whose behavior requires no explanation in terms of semantic regularities - that is, in terms of rules and representations" (p. 130-131). He then argues for the explanatory priority assumption, the assumption that, if a process can be described in such a way that the interesting regularities are captured at a low level of description, then any higher-level descriptions of the process are gratuitous (p. 131-132). Descriptions in terms of semantically-interpreted representations are, for Pylyshyn, and for cognitive scientists in general (e.g. Marr, 1981; Newell, 1982), higher-level than descriptions in terms of physical, chemical, or biological properties (the "device level" of Newell (1982)), or than descriptions in terms of the program being instantiated (the "symbol level" of Newell (1982)). From these three claims it follows that, since architectural functions need not be described at the semantic level, they should not be described at the semantic level. If a function is cognitively penetrable, however, then by the definition of cognitive penetrability, it must be described at the semantic level. Therefore, cognitively penetrable functions cannot be architectural; they "must be explained in terms of the semantic contents of beliefs [and other mental states]" (Pylyshyn, 1984, p. 134).

Two questions need to be asked about this argument. First, is it sound? Second, is its conclusion, that penetrable functions cannot be architectural, sensible? These questions will be considered in turn.

There are at least two difficulties with Pylyshyn's argument. The first is that the claim that "processes carried out in the functional architecture ... require(s) no explanation in terms of semantic regularities" is simply not true. As shown above, all the processes that any device carries out are carried out in the functional architecture; the program defining the architecture is the only program that the device executes. Therefore, if anything that any device does requires semantic explanation, then some architectural processes require semantic explanation.

This fact alone, however, does not show Pylyshyn's position to be untenable. Pylyshyn introduces the notion of functional architecture in order to distinguish functions that, for the purposes of a given explanatory task, can be regarded as instantiated as "low-level" processes from those that are considered "high-level." As pointed out above, it is always possible to draw such distinctions. Pylyshyn, however, wants the distinction to be _principled_; he wants it to simply be the case, as a matter of natural fact, that certain phenomena are cognitive and others are not. The argument for this principled distinction rests of the explanatory priority assumption.

The second difficulty concerns this assumption of explanatory priority. Consider the following case. An experimenter notes that periwinkles avoid stems in certain circumstances and climb them in others, and that humans avoid trees in some circumstances and climb them in others. What is common to these two cases? The most obvious commonality is that both periwinkles and humans treat objects (and fairly similar objects, at that) as obstacles in some circumstances, and as things to climb in others. This, however, is in both cases a semantic explanation, an explanation that depends on the notion of taking something to have a certain property. Hence, there is at least one case in which it is useful to view the marsh periwinkle from the semantic perspective, even though all of its behavioral regularities can be described at the program (symbol) level of description. Therefore, what counts as far as determining whether a certain level of description is appropriate is not whether the organism's behavioral regularities can be described at a lower level, but whether it is ever useful to describe them at a higher level. One way of asking whether higher-level descriptions are likely to be useful is to ask whether anything to which the organism in question might be compared can be adequately described at the lower level of description. However, any living thing might interestingly be compared to us; therefore, semantic descriptions of any organism might be appropriate [5]. Most organisms, and certainly marsh periwinkles, however, do not exhibit cognitive penetration. Therefore, while penetrability is a sufficient condition for the usefulness of the semantic level of description, it is certainly not necessary. The explanatory priority assumption, in its present form, is simply misguided.

Pylyshyn's argument for the cognitive penetrability criterion, therefore, suffers from at least two serious difficulties. One must ask, therefore, if the criterion is even sensible.

The most obvious way to patch up the criterion is to revise it to read, "if a function is cognitively penetrable, then it cannot be explained solely by appeal to processes at the level of description of programs; it must be explained in terms of semantic contents." This revised formulation recognizes that all processes are architectural, and that cognitive penetrability is a sufficient, but not necessary, condition for interest at the semantic level.

This revised formulation is quite sensible. If a function is cognitively penetrable, then it depends in a "rationally explicable" way on the semantic contents of beliefs, desires, etc. This, however, is surely interesting from the point of view of explanation. If a process depends on such factors, that dependence should be explained. Therefore, cognitive penetrability is a sufficient condition for semantic interest.

Is, however, the revised formulation useful as an empirical criterion? It is clearly not useful as a criterion for separating architectural functions from non-architectural functions: there are no such thing as the latter. Perhaps, however, it is a useful criterion for empirically deciding what set of functions to interpret as the functional architecture, and what to interpret as a program virtually executed on this architecture. This weaker use of the criterion is, perhaps, what Pylyshyn had in mind all along.

The revised criterion alone, however, will not perform even this function. The reason is that this formulation of the role of the criterion is circular. If it had not already been decided that the level of description of beliefs, desires, etc. was the one of interest, then cognitive penetration would not be defined in terms of beliefs and desires. If it has already been decided that the level of description of beliefs and desires is the one of interest, however, one hardly needs an empirical criterion to do the deciding. Pylyshyn evidently recognizes the direction of argument here, when he claims "the second [cognitive penetrability] criterion assumes that what I have been calling cognitive phenomena are a 'natural kind' explainable entirely in terms of the nature of the representations and the structure of the programs running on the cognitive functional architecture" (p. 113-114). However, if the line between virtual machine (i.e. functional architecture) and virtual programs has already been drawn at the semantic level of description, a criterion is not needed to draw it.

The fundamental problem with the revised criterion is that decisions about how to interpret the behavior of an arbitrary device, and about the level of description to be used to do so, are not empirical decisions (Cummins, 1977). Such decisions depend not only on the explanatory task at hand, i.e. on a prior decision as to what the phenomena of interest are, but also on such factors as explanatory coherence, generality, etc. For example, the states of a human brain could be interpreted as representing large integers, e.g. by numbering the neurons, assigning a "1" to every one firing and a "0" to every one not firing, and interpreting the result as the representation of an integer in binary notation. This integer would change as the firing pattern changed. The human brain could, thus, be interpreted as a large random-number generator. The problem with this interpretation is that it is not (or at least, not obviously) very useful. It does not serve any explanatory task, nor does it allow regularities between, e.g. humans and other animals, to be formulated that are interesting. These decisions have some empirical basis, e.g. they rely on the fact that no human science investigates random number generation by brains; they are not, however, empirical decisions. Such decisions may be revised on the basis of whether the investigation turns up any interesting regularities, but even the revised decision is essentially arbitrary.

If cognitive penetrability is not an empirical criterion for deciding what is and is not "cognitive," what is it? A reasonable alternative is to say that it is a statement of an explanandum for cognitive science. Some processes are cognitively penetrable, others are not. Why? Any cognitive theory worth its salt must answer this question. What Pylyshyn has done is to perform the service of bringing this question forcefully to our attention.

Fodor's (1983) notion of modularity provides the beginnings of a general answer to this question. A function is cognitively penetrable if, and only if, beliefs, desires, etc. are included in its knowledge base. The problem of cognitive penetration is not a question about the program being run, or the architecture that it is run on at all. It is a question about what data the program that defines the architecture has access to. Humans differ from periwinkles in, among other things, the fact that the database that serves the object identification function in humans contains all manner of beliefs and other mental states, the contents of which vary considerably as a function of experience. Humans can, therefore, see trees as all sorts of things. Marsh periwinkles, on the other hand, have a small. and apparently fixed database in the service of object identification, and see stems only as climbable things or obstacles. Changing how a periwinkle sees a stem would require the neural equivalent of re-soldering some connections in the physical substrate of the database. In the case of a human, it only requires saying, "look at it this way ..." with enough persuasive power.

In summary, Pylyshyn's cognitive penetrability criterion is not a criterion, and so does not provide any constraint on cognitivist explanations. The fact that some phenomena are interesting from the point of view of cognitive modeling is assumed, not discovered. Moreover, the way that a model is built depends not only on the facts of the matter, but also critically on the explanatory interests of the modelers, and on the task at hand. The phenomenon of cognitive penetration does, however, point out an important explanandum for cognitive science. Cognitivists, armed with the notion of modularity, can begin to explain this phenomenon. A function is cognitively penetrable if beliefs, desires, etc. are included in its knowledge base, and if these beliefs, desires, etc. can change as a function of semantically interpreted input. Whether ER can do as well toward explaining cognitive penetration will be considered in the next chapter. The question of constraining cognitivism will be readdressed there, as well.

The Recognition Regress Argument

The second major ecological realist argument against cognitivism described in Ch. II is the recognition regress argument. The major premise of the recognition regress argument is that, in cognitivism, but not in ER, "intensional description mandates conceptual ascription" (Turvey et al, 1981, p. 252). The question of what intensional description mandates in ER will be left to the next chapter. Does it, in fact, mandate conceptual ascription in cognitivism?

Consider again the model of the marsh periwinkle shown in fig. V-3. It will be assumed that the presence of transducers and simple logic circuitry does not mandate conceptual ascription; if it does, ER is in as much trouble as cognitivism, for it, too, requires that organisms contain transducers and logic circuits (see fig. IV-2). The question, therefore, is whether contrast enhancers or object identifiers require conceptual ascription.

Both of these processes use information. That information _could_, in each case, be encoded in a separate, propositional database that the relevant functional unit accessed as it worked. The information could, in other words, be stored, and used as input by the processes in question. As pointed out in Ch. II, this is the sort of information storage and use postulated by internal manual models (Cummins, 1982b; 1983), e.g. by the model of Fodor (1968).

The information stored by the marsh periwinkle does not, however, need to be used as input in this way. Alternatively, it can be causally represented in the sense described in Ch. IV; i.e. it can be encoded in the structure of the program instantiated by the functional unit.

Were the information used by these devices encoded in a separate database, and used as input to the process, it would be quite reasonable to call it "conceptual" knowledge. If the information is causally encoded, however, there is no obvious reason to call it "conceptual." There is, however, a very good reason for ecological realists, at least, <u>not</u> to call it conceptual. The reason is that ecological realist resonators also store information in this way; if this sort of storage is conceptual in cognitivism, it is conceptual in ER as well. One can therefore conclude, at least provisionally, that the recognition regress argument has no weight against the model of the marsh periwinkle shown in fig. V-3, or against similar models of other lower organisms. At least, ecological realists cannot use this argument against these models without using it against their own models, as well.

Let us consider, then, the model of human perception in fig. V-2. Is conceptual ascription required here?

The part of the model that concerns the production of the 2 1/2-d sketch does not require conceptual knowledge, for the same reason that the model of the periwinkle does not. All of the information used by these modules can be causally encoded. This leaves the information used in object identification, and, in the extended model, in affordance identification. In these cases, it looks as if conceptual ascription may indeed be required, if not in principle (Cummins, 1982b), then at least in practice. Marr, for example, talks about "libraries" of "canonical shapes" being used in object identification. It at least sounds as if these libraries are used as input to the program that identifies shapes.

If this is a case of conceptual ascription, however, it is a case of ascribing concepts to a _virtual_ machine, i.e. it is a case of interpreting the behavior of the hardware as the kind of behavior it would display _if_ it had concepts. The cognitivist (at least the cognitivist who respects the hardware/virtual machine distinction) nowhere need ascribe concepts to the hardware; concepts are only ascribed to virtual hardware, i.e. to _software_.

It is not at all clear that conceptual ascription to virtual machines will work in the premise of the recognition regress argument. The argument is about biological evolution, and, in Michaels and Carello's case, about whether a non-virtual machine could successfully access memory as a part of perception. Neither version of the argument even considers the question of the relation between the program a virtual machine runs and the program a non-virtual machine runs. It is not at all clear, however, that these two programs would share their memory or control structure; indeed, programs written in high-level languages do not, in general, share such features as memory or control structure with their machine-language counterparts. Unless the two programs share control structure, however, attributing concepts to the one is not attributing concepts to the other. If it is to use the recognition regress argument against the type of theory advocated by Marr (1981), therefore, ER must argue that, in the case of vision, concepts must be ascribed to both the hardware and the virtual machine. ER has not argued for this claim, nor is there any immediate reason to believe that it is true.

Inference in Cognitivist Models

The foregoing arguments are sufficient to defuse, if not fully refute, the standard ecological realist objections to cognitivism. They do not, however, fully answer the principal worry that the ecological realist arguments point out. Ecological realists accuse cognitivists of attributing to organisms cognitive capacities, and in particular, inferential capacities, that they do not need and could not acquire evolutionarily. How much weight does this objection really carry?

The traditional cognitivist view of inference, as pointed out earlier, is that cognitive systems infer conclusions from premises through the use of rules in much the same way that a beginning logic student might. Fodor's (1968) internal manual model of shoe-tying, in which the humunculus actually reads and follows instructions from a book, is a perfect example of this view of cognition. The worry that ecological realists raise is precisely the worry that actual cognitive systems do not possess anything analogous to the beginning instructions or the rule books in Fodor's model.

This traditional view of inference is applied to perception in the model shown in fig. V-5. The transducer produces a description of the world, such as a 2 1/2-d sketch, that serves as a premise for the cognitive inference engine. Inferences from such premises to the identities of objects in the world make use of facts about the world such as knowledge of the typical shapes of objects (as in Marr's model). Inference also requires rules, e.g. rules for matching general patterns with specific instances. The conclusion reached by the inference system serves, in turn, as a premise for similar systems that, e.g. decide what to do in the presence of particular objects.

The claim that cognitive systems employ stored rules is the heart of this traditional view. Many authors have argued, against the traditional view, that a system can <u>use</u> rules without storing them in a database of any sort (e.g. Fisher, 1974; McDermott, 1976: Dennett, 1978a; 1983; Cummins, 1982b; 1983). Cummins' critique of IMM's, in particular, is based on the fact that a system can run a program that instantiates rules for carrying out inferences without storing the program, or the rules anywhere.

If, however, it is possible to incorporate the rules postulated by the traditional model into the program instantiated by the inference engine, it is equally possible to incorporate all of the facts employed by the inference program as data into the program as procedural knowledge (McDermott, 1976). Assuming that the resulting system can be divided into a set of parallel, encapsulated modules (Fodor, 1983), a system such as is shown in fig. V-6 results. In

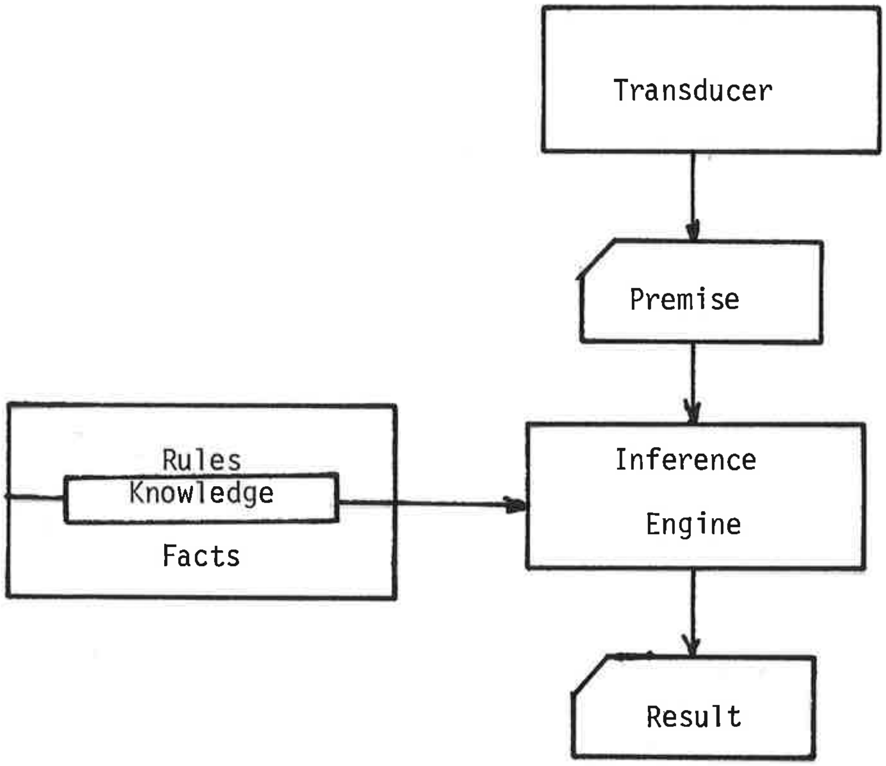

Fig. V-5: An internal manual model of an inferential process. All of the information used by the inference engine is stored in the database. Compare with the model of Fodor (1968).

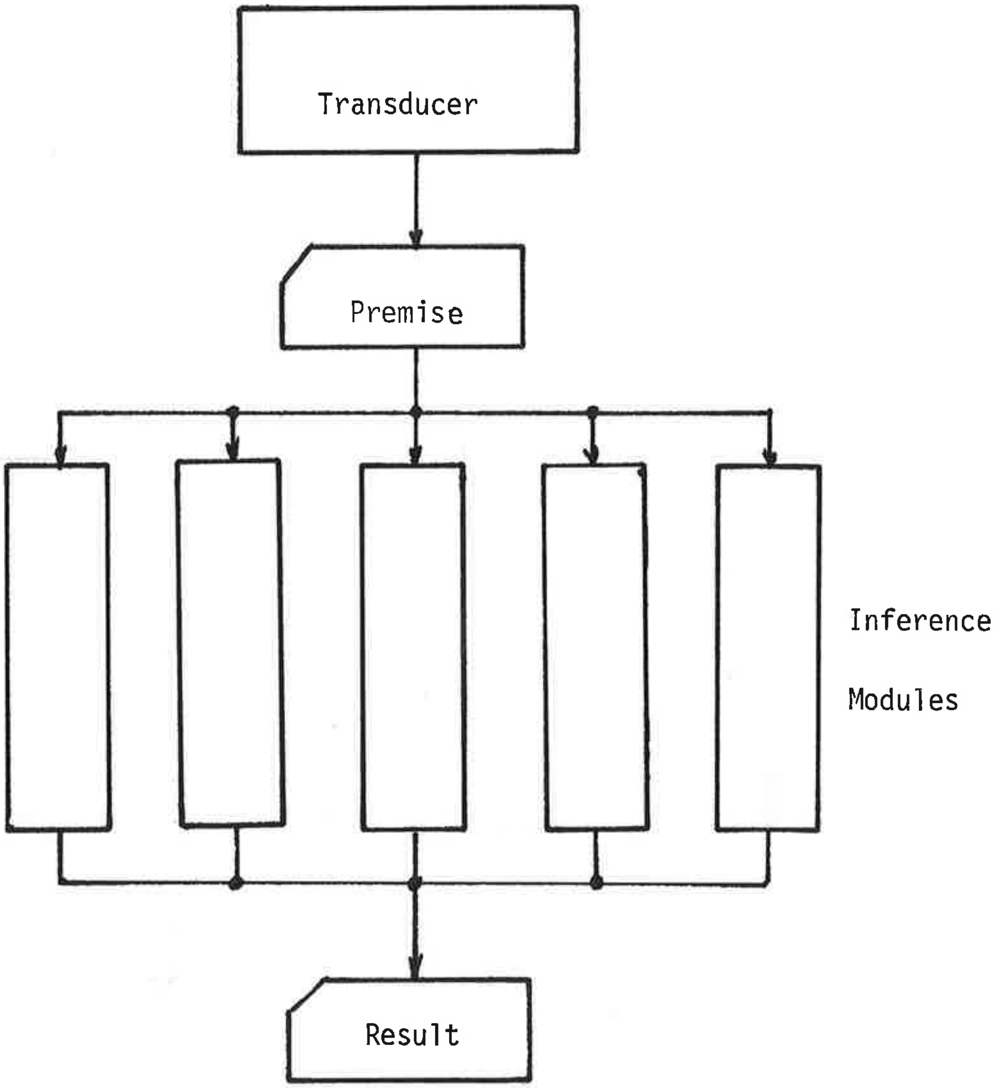

Fig. V-6: An inference system that uses fully-encapsulated inference modules (see Fodor, 1983). All of the knowledge employed by the modules is stored procedurally; there is no explicit database. The knowledge available to each module is limited; the modules are, therefore, presumably special-purpose.

this model, the propositional database has disappeared altogether. Production system architectures such as that of Anderson (1983), and connectionist architectures such as that of McClelland and Rumelhardt (1981) are of this general form.

The inference modules in fig. V-6 are inferential in exactly the wide sense introduced in Ch. IV. They are described as inferential, in other words, because it is useful to do so for the purposes of describing the system's behavior. Each module behaves as if it were a traditional inference system, with rules and data, of the sort shown in fig. V-5. The logical structure of the programs that they instantiate is, however, very different from that of the program instantiated by the inference engine of fig. V-5. The two systems may be I/O equivalent, but they are not strongly equivalent.

The traditional picture of inference, therefore, is unnecessarily restrictive. Like Pylyshyn's explanatory priority assumption, it rules out a wide variety of processes that may be usefully described as inferential in a completely ad hoc way. If this picture is abandoned, the need to postulate evolutionarily implausible databases in the support of cognitivist models disappears.

Which of the many possible I/O equivalent cognitivist models of a given cognitive process is correct is a matter to be decided on both empirical and explanatory grounds. In particular, models of the sort shown in fig. V-6 will predict a greater degree of encapsulation of object-recognition inferences than models of the sort shown in fig. V-5. The degree of encapsulation of many processes can be investigated experimentally (Fodor, 1983). On the other hand, a model that predicts a control structure similar to that of the neural network that instantiates the process being modeled will provide a fuller explanation of the process than one that predicts a completely dissimilar control structure. While these considerations are unlikely to pick out a single "best" model, they can be expected to narrow the field.

Conclusion

Animal brains are hardware systems very different from current digital computers. As Feldman and Ballard (1982) point out, animal brains are capable of performing very complex computational tasks, such as word recognition, in about one hundred computational steps. The task of cognitive science is to explain these capacities by duplicating them in programs.

The principal theoretical tool that cognitivism brings to bear on this problem is interpretive analysis (Cummins, 1983). Cognitivists attempt to construct virtual machines (programs) that emulate the brain, i.e. that imitate what the brain does. The models shown in figs. V-5 and V-6 are very abstract descriptions of such virtual machines. The goal of cognitivism is to find a virtual machine that runs a program that the brain can be described as running. Attempts such as that of Marr (1981) to construct models of cognitive processes that are simultaneously constrained by I/O and hardware considerations are the first steps toward determining the algorithms instantiated by the functional units of the brain.

Cognitivism has been reconstructed, in this chapter, in terms of the notions of virtual machines and functional architectures. As an illustration, the theory of vision of Marr (1981) is expounded, and extended, in terms of these notions. Marr's theory proposes that perception is the result of information processing by a number of informationally-encapsulated modules, each of which performs a specific function.

Pylyshyn's (1980; 1984) cognitive penetrability criterion for separating architectural functions from functions defined over semantically-interpreted representations was shown not to be a criterion after all, but rather a statement of one of the explananda of perceptual theories. Fodor's (1983) concept of informational encapsulation provides cognitivism with the beginnings, at least, of an explanation of this phenomenon.

The ecological realist recognition regress argument was shown to ignore both the possibility of causal information encoding and the virtual/non-virtual machine distinction. Unless it can be reformulated with these distinctions taken into account, it is unsound. There does not, however, appear to be a way to reformulate this argument that does not leave it cutting against ER as well as against cognitivism.

Finally, cognitivist appeals to inference are shown not to entail the representation of rules or knowledge in databases. While models that do invoke such storage are candidate cognitive virtual machines, they are by no means the only ones that satisfy the requirement of I/O equivalence with the brain. Alternative models exist that do away with the concept of a database altogether, thus obviating the principle ecological realist worry about cognitivist models. The fortunes of such models in the debate with ER will be considered in the next chapter.

NOTES - CHAPTER V.

1. Skinner (1953, Ch. I-III) provides a fairly clear statement of the thesis of the behavioral identity theory, i.e. that psychological states are behavioral states. Borst (1970) is a useful anthology of the main "in-house" arguments concerning the physiological identity theory, i.e. the claim that psychological states are to be identified with brain states. Putnam (1967) provides a clear example of the functional-state identity theory, the claim that psychological states can be identified with machine table states of a Turing machine that reproduces the I/O relations of the subject.

2. Program outputs are typically written, as files, to output devices such as disks or printers. The states of these devices must, therefore, considered as parts of the machine state for the purposes of this definition of program instantiation.

3. Marr (1981; 1982) emphasizes that the "sketches" produced by early vision indeed represent the world. Pylyshyn (1984), who reserves the label "representational" for processes and states that are cognitivly penetrable, disagrees, and in fact specifically mentions Marr's "sketches" as non-representational (p. 215). Pylyshyn's terminology seems overly conservative; there seems to be a clear sense of "represents" in which Marr's sketches, which have features that can be placed in a one-to-one correspondence with features of the visual world, do respresent the visual world. His conservatism, moreover, appears, as will be argued below, to result from some fairly deep confusions.

4. Indeed, one draws boxes, in Fodor's account, around the smallest processes that are encapsulated. Any arrows from right to left, or from a processor to its database, must be enclosed within a box.

5. A particularly striking example is provided by the fact that yeast (Alberts et al, 1983), and even bacteria (Morse, 1984) use sex pheromones to locate and attract mating partners. It is quite interesting to compare the use of pheromones in these organisms with their use in humans for the same purpose. If bacteria are, under some circumstances, interestingly viewed from the semantic level of description, then any organism is interesting from this point of view. This argument counters not only Pylyshyn's argument, by also Fodor's (1984a) argument that lower organisms, such as paramecia, should never be described at the semantic level.

CHAPTER VI

THE DEBATE REVISITED

The last two chapters reconstructed ecological realism and cognitivism "from the inside" in response to the debate presented in Ch. II. The reconstructions answer some, but not all, of the questions raised by the original debate. The purpose of this chapter is to see what is left of the debate, and to ask whether the conflict between ER and cognitivism can be resolved. It will show that ER and cognitivism are equivalent theories of the mechanism of transduction, but that they make different assumptions concerning the extent to which perception is transduction. Cognitivism assumes that only the initial stages of perception can be described as transduction, while ER assumes that perception is entirely transduction. This is a substantive empirical question that must be answered experimentally.

Before proceeding with this argument, however, it is useful to summarize the competing positions, and to see, in general terms, which questions left outstanding by the original debate have been answered by the reconstructions. The questions that remain will then be considered in the context of models of both lower- and higher-animal perceptual systems.

The ER-cognitivism debate is often presented (e.g. by Turvey et al, 1981) as a debate over whether perception is governed by ecological laws. Traditional ER is a subsumptive theory, in which subsumption under ecological laws is taken to explain perceptual phenomena. Cognitivists, in contrast, deny that there are any laws of the kind invoked by ER, i.e. any laws linking distal stimuli to percepts or actions. Indeed, the major focus of Fodor and Pylyshyn's (1981) critique of ER is an argument, the constraint argument, that attempts to show that ecological laws do not exist.

The issue of whether ecological laws exist was shown, in Ch. IV, to be something of a red herring. The important question is whether appeals to ecological laws can provide explanatory answers to the questions that cognitive science asks. The answer is that they cannot. In particular, the subsumptive approach to ER cannot provide any account of the mechanisms by which perception occurs. Following this reasoning, it was argued in Ch. IV that ER must adopt the analytic explanatory strategy, and must, in particular, countenence explanations that refer in essential ways to the internal states and processes of the organism. Only by appeal to such internal states and processes can ecological laws themselves be explained. The adoption of the analytic strategy is a major departure from the aims and methods of traditional ER.

The internal processes postulated by analytic ER include resonance, and processes involved in the control of resonance. The internal processes postulated by cognitivism, as described in the last chapter, include object and affordance identification. In ER, organisms are claimed to resonate to signals encoding affordances directly. In cognitivism, in contrast, they must infer the affordances of their environments from representations of the objects in their environments.

The second traditional characterization of the debate is that it concerns whether the perception of the affordances of objects in the environment is direct, i.e. unmediated by inferences (e.g. Michaels and Carello, 1981). Once ER has adopted the analytic explanatory strategy, however, this question also becomes something of a red herring. The important question, in this case, is not whether the mechanism of affordance perception is direct, but whether the proximal stimulus is impoverished, i.e. whether the proximal stimulus contains enough information to specify an affordance to the perceiving organism. Ecological realists claim that proximal stimuli are not impoverished, i.e. that they specify affordances, while cognitivists claim that they typically are impoverished, i.e. that they do not specify affordances.

The question of whether perception is direct is traditionally conflated with the question of whether perception requires the addition of information to the proximal stimulus. The analytic approach to ER shows that these two questions are distinct. Even in ER, perceiving an affordance requires the use of certain information, which is causally encoded by the resonator that detects the signal encoding the affordance. This information is used to decode the signal responsible for the proximal stimulation, i.e. to extract the information about affordances encoded in the signal. If, for example, the proximal stimulus encodes the fact that something is edible, the information contained in the resonator is used to decode the signal, but does not add any additional information about affordances to the message that the thing is edible.

Cognitivism and ER do not, therefore, disagree about whether information stored, or otherwise encoded, by the organism is used by the processes responsible for the perception of affordances; they both must assume that it is. Their disagreement, if it real, must concern either the nature of the information, or how it is added, or both.

In summary, there are three remaining issues in the ER-cognitivism debate. First is the question of specification: are affordances specified by proximal stimuli? Second is the question of what information is used to extract information about affordances from the proximal stimulus, and of how it is used. Last is the question of whether perception involves direct resonance to affordance-encoding properties of signals, or inferential processes of object and affordance identification of the sort proposed by Marr (1981). These three questions will be considered, in this chapter, in the context of several cases of perception by both lower animals, i.e. animals in which cognitive penetration does not occur, and humans.

Case I: Lower Animals

As an example of the problem of explaining perception by lower animals, let us consider again the case of the marsh periwinkle. ER and cognitivism propose different models of the perception of the affordances "climb-upable" and "collide-withable" by the marsh periwinkle. These models, shown originally in figs. IV-2 and V-3, are summarized in fig. VI-1. Which is correct?

The two models can be assumed to have identical I/O behavior. In each case, the response of the snail to a light array encoding the presence in the immediate environment of a tall, narrow object will vary depending on whether it senses the pressure variations encoding the advance of the tide. If the tide is advancing, the snail will approach and climb the object seen; if the tide is not advancing, the object will be avoided. These are assumed to be fixed action patterns, i.e. the system performing each is assumed to carry out a fixed set of behaviors whenever stimulated, and to do nothing when not stimulated.

Given these assumptions, the differences in the models are only those shown in fig. VI-1. How big are these differences, and how could it be determined which model is correct?

The most obvious difference between the models is that different devices are postulated by the two theories. ER postulates resonators, while cognitivism postulates transducers and object identifiers. Resonators and object identifiers are, however, both types of feature detectors (see Ch. IV, V). It is, therefore, not clear how great this difference really is.

In order to decide what differences separate the two models, it is necessary to see exactly how the two mechanisms work. Fig. VI-2a shows a possible model of the internal structure of a resonator. It consists of three functional units: a demodulator that resonates to the external signal, an analog-to-digital conver-

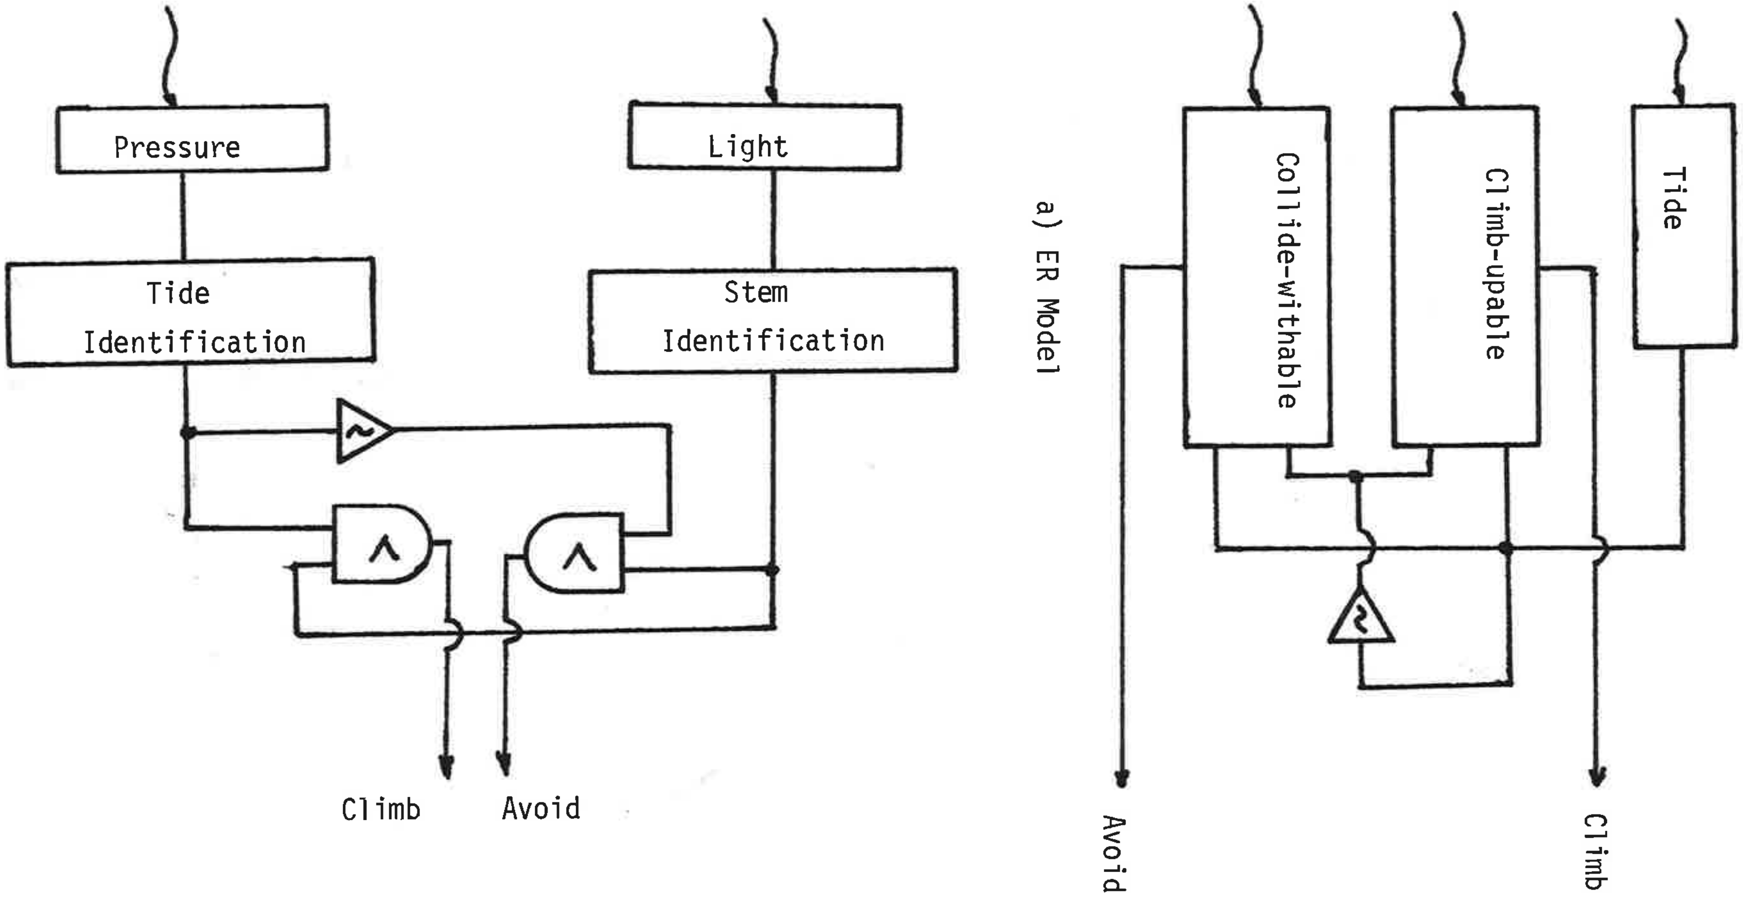

a) ER Model

b) Cognitivist Model

Fig. VI-1: Summaries of the cognitivist and ecological realist models of the marsh periwinkle.

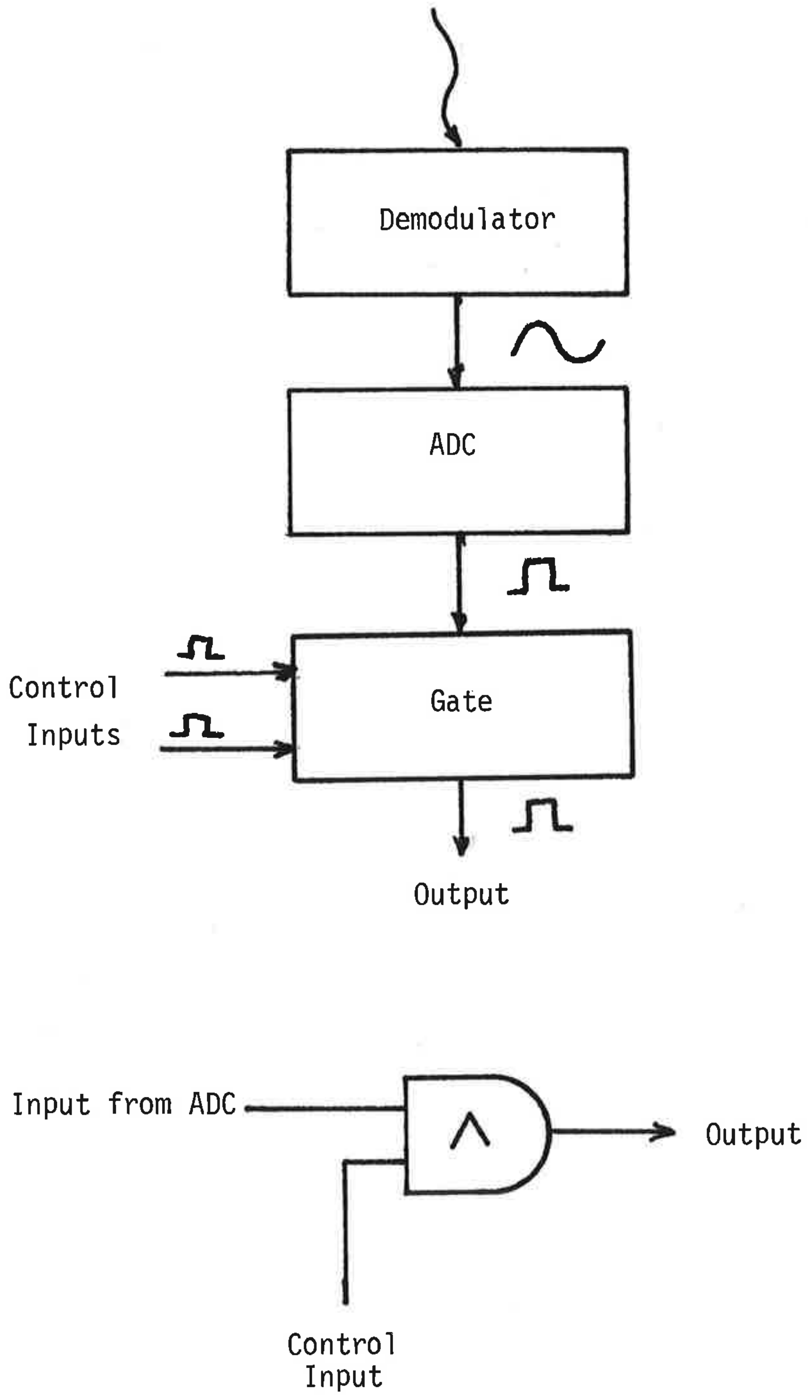

Fig. VI-2: Analysis of the resonator function into functions postulated by cognitivist models. Other functionally equivalent analyses are equally possible.

ter (ADC), and a gate that is controllable by other internal components, such as the logic circuitry shown in fig. VI-1. One can now ask how these components compare to those postulated by cognitivist models.

The demodulator component is a transducer, that is, it merely translates the component of the external signal to which it is attuned into an internal signal bearing the same information. The information in the component of the signal to which the resonator is sensitive is preserved exactly, in the absence of transducer noise, just as the information in the radio signal considered in Ch. IV is preserved exactly. The ADC is, likewise, a transducer; it translates its analog input into a digital output that can be used by the logic circuitry. Again, the two signals are assumed to carry exactly the same information. This component is postulated only so that the control of resonance can be via logic circuits. The controlling gate, on the other hand, is not a transducer; its output depends not only on its input from the ADC, but also on its control inputs, which come from other parts of the perceptual system.

This circuit can be compared with that postulated by the cognitivist model. In the cognitivist model, the object identifiers are transducers; they are not cognitively penetrable, so they have a fixed I/O function. The composite function of the input transducer and the object identifier in the cognitivist model is therefore analogous to that of the resonator-ADC complex in the ER model; the only difference is that the former responds preferentially to an object, while the latter responds preferentially to an affordance. The logic circuits with which these transducers communicate are, however, identical; as shown in fig. VI-2b, the controlled gate appearing in the ER model is functionally equivalent to the "and" gate appearing in the cogitivist model.

The *only* difference, therefore, between the two models is that the transducers respond to encoded affordances on ER, while they respond to encoded objects in cognitivism. However, it has been assumed that the property of the signal that encodes the affordance 'climbable' for the marsh periwinkle is exactly the same property that encodes the information that the object is a plant stem. The *functions* of the transducers, i.e. the I/O rules that they instantiate, are therefore identical. This is, however, just to say that, as far as the marsh periwinkle is concerned, responding preferentially to an object *just is* responding preferentially to an affordance.

This conclusion is surprising, but it makes sense in terms of the overall function of the periwinkle's cognitive system. Marr (1981) proposes that object identification, using canonical object-shapes stored in memory, is the stage of perceptual processing immediately following the construction of the 2 1/2-d sketch. Cognitivists, indeed, generally talk about inferring the affordances of an object from the properties *of that object*. Objects are taken to be detected through the action of feature detectors. What feature detectors actually detect, however, are not objects, but *properties* of objects, such as their shapes, colors, etc. It is these properties that are combined to construct the 2 1/2-d sketch. If the affordances of an object for the perceiving organism are nomically correlated with the properties of the object to which that organism's feature detectors are preferentially sensitive, then *feature detection is affordance detection*.

From the point of view of evolution, there is every reason to suppose that, at least in lower organisms, such as the periwinkle, that have fairly simple behavior patterns, feature detectors that are affordance detectors would be advantageous. Affordances are, by definition, the features of environmental objects that are relevant to the organism. The properties of the environment that exert selective pressure on an organism are, however, exactly relevant properties; they are, therefore, affordances. Detecting the relevant properties of things that are going to exert selective pressure is, however, generally advantageous; it means that one can recognize predators and other dangerous items in the environment. Ceteris paribus, one would, therefore, expect organisms with affordance detectors to enjoy a selective advantage.

In summary, if affordances are specified by signals in perceptual media, one can expect, on evolutionary grounds, that feature detectors of the sort proposed in Marr's theory will be affordance detectors. In this case, ER and cognitivism are simply equivalent; they propose equivalent models of the mechanism of affordance detection. The question of what information is specified by signals in perceptual media is, therefore, central to the debate. It was shown in Ch. IV, however, that there are cases in which affordances are not specified by proximal stimuli, even in natural environments. A consideration of these cases is therefore in order.

Perceptual Error

Recall the case of the anglerfish from Ch. IV. In this case, the prey fish acts as if it perceives an affordance that the anglerfish does not have, and as if it does not perceive the affordance that the anglerfish does have. Such cases show that the claim that the affordances of objects are always specified in proximal stimuli is, as a general statement, simply false. This claim is, however, central to traditional ER. What effect does this fact have on the debate?

The claim that affordances are specified in proximal stimuli is central to traditional ER for two reasons. First, traditional ER is concerned to explain the "practical success of an organism's 'everyday' behavior" (Turvey et al, 1981, p. 238). Traditional ER, therefore, tends to ignore exceptional cases. Second, traditional ER attempts to explain behavior subsumptively, by appeal to ecological laws. This means that traditional ER cannot appeal to failures of specific mechanisms to explain misperceptions; instead, it must appeal to failures of laws. As argued in Ch. IV, however, such appeals are completely ad hoc; they cannot themselves be explained. Traditional ER, therefore, has considerable difficulty accounting for unsuccessful behavior, such as that of the prey of the anglerfish (see also Ullman, 1980). Had ER also attempted to explain such unsuccessful behaviors, it might not have placed so much stock in the claim that affordances are always specified by signals.

Whether organisms are successful in their behavior is, however, largely irrelevant to the task of explaining how perceptual mechanisms work. One can describe the mechanism of resonance to a signal without worrying about whether the signal actually encodes the affordance that the organism appears to detect. The assertion that affordance detection works by the resonance mechanism is, therefore, independent of the claim that affordances are specified by signals. Analytic ER is, therefore, independent of this claim.

The specification claim is not, however, inconsequential for perceptual theory. One must assume that, in the natural setting, organisms "get it right" most of the time. This is all that is required for the evolutionary argument of the last section to go through. One does not, for this argument to be successful, have to assume that they get it right all of the time. This is an additional, unmotivated assumption of traditional ER. The existence of perceptual error in natural settings shows that it is false.

In summary, the claim that signals always specify affordances is false. Signals only specify affordances enough of the time for organisms to make it. The strength of this claim varies from organism to organism and from niche to niche, just as the evolutionary success of species varies. This claim must, however, be distinguished from the claim that affordance perception works by the resonance mechanism. The latter claim entails only that resonators respond selectively to signals that, in the proper niche, enough of the time, encode affordances. This claim is consistent with the existence of misperception in both the laboratory and the natural environment. Cases of error can be explained, in analytic ER, simply as cases in which the signal does not encode the affordance that it usually encodes. This explanation allows sense to be made of cases of both naturally occurring errors and of laboratory illusions within the general framework of analytic ER.

This solution to the problem of perceptual error can, clearly, also be adopted by cognitivists. Feature detectors respond to properties of signals that normally encode properties of objects. If, on a given occasion, the signal does not encode the usual object property, the feature detector responds anyway. These cases are illusions.

This account of perceptual error has the added advantage of decoupling the problem of illusions from the traditional question of perceptual inference. It makes no mention of inference. The organism does not indentify an object, and then infer an affordance; it responds to a feature of the signal that normally encodes a property of the object that may be an affordance. This process can be described, as Marr in fact describes it, as inferring the existence of the object from the existence of signal, but it cannot be described coherently as inferring the existence of the property from the existence of the object. The inferential path, in fact, goes in just the opposite direction. The existence of objects is inferred from the detection of their properties.

Changes in Perceptual Abilities

It has now been argued that cognitivism and analytic ER are equivalent theories of perception in the case of the marsh periwinkle, and by analogy, in the case of lower organisms in general. The problems posed by perceptual error do not differentiate between the two theories. It may, however, be possible to distinguish the two theories on the basis of their responses to the problem of changes in perceptual abilities. This possibility will now be explored.

Consider the case of the common tick (von Uexkull, 1957). As the tick matures, several changes in its behavior occur. Prior to sexual maturation, ticks live on or near the ground. After maturation and mating, female ticks climb trees and bushes, and then fall off when warm-blooded animals pass under them. The climbing behavior is due to a general phototropism, the release behavior is triggered by an olfactory cue, butyric acid. The mature female tick thus responds to two encoded affordances that it did not respond to when immature. How are perceptual theories to account for this change?

As argued in Ch. IV, analytic ER is committed to the claim that affordance-detection capabilities are not learned. The only other mechanism of perceptual change recognized by traditional ER is evolutionary selection. In the case of the tick, however, the most plausible explanation of the change is to say that it is *maturational*. The question of interest, then, is that of how perceptual theories are to account for maturational changes in perceptual abilities.

The resonator model shown in fig. IV-3 can account for maturational changes in perceptual abilities only if it is claimed that the changes in question amount to existing resonators being turned on. If this is assumed, then the only novel feature of the maturation case is that the signal that turns the mature resonators on originates inside, not outside, the animal [1]. Analytic ER can consistently assume the sort of internal clock that the "internal trigger" cases require, together with the multiple levels of control of resonators that such triggers involve. Traditional ER, of course, cannot do this.

Cognitivist theories of perception do not have to accept the problem of maturational changes in perceptual abilities as their own; unlike ER, cognitivism can relegate such problems to a separate general theory of cognitive processing. However, this secondary theory must face the problem with tactics much like those of analytic ER. As it was presented in Ch. V, cognitivism contains no provisions for creating new object identifiers; all changes in object-identificatuion abilities must, therefore, be due to changes made in the use of the outputs of existing object identifiers.

Neither of these responses to the problem of describing the tick's new perceptual abilities is, however, fully satisfactory. Physiologically, the tick does not turn on new receptors; it rather gains new receptors by a process of biological development. Maturational change presents a difficult problem for both ER and cognitivism because it involves changes in the hardware of the organism. Virtual machines can simulate such changes, but current artificial hardware cannot duplicate them. Models can, therefore, be at most I/O equivalent to the natural systems; they cannot be strongly equivalent. This problem, however, is faced by cognitive science in general; it does not distinguish between cognitivism and ER.

A second, similar problem is posed by the fact that the signal to which a resonator or object identifier responds is often fixed developmentally. Consider the nestmate recognition phenomenon in ants (Carlin and Holldober, 1983). Specific olfactory cues allow ants to recognize nestmates, regardless of their genotypes. These cues are different for different nests. The specific cue that indicates "nestmate" to a given ant is established at a particular point in maturation. Thus, at a particular time, the environmental signal that encodes an affordance is established that was not established previously.

This case cannot be explained merely by appeal to an internal signal that turns a resonator on or off. In this case, the resonator is not only turned on; the very feature of the ambient array to which it is to respond is established. Indeed, the signal encoding "nestmate" is just the first signal of the right type encountered by the detector when it matures.

This case provides a serious challenge to both cognitivism and ER. It is especially embarrassing, however, to IMM's. In an IMM, a description of the signal encoding "nestmate" must be stored in a database. However, this description cannot be so stored before the correct maturational stage, since the property encoding the affordance, and hence the content of the description, is not determined until the correct stage. The description cannot be encoded at the correct stage, however, since IMM's have no facilities for encoding data de novo (Fodor, 1975). An IMM, therefore, cannot account for this case [2]. Even in the case of such a simple animal as an ant, therefore, information must be encoded causally or procedurally.

The example of nestmate recognition shows more than this, however. It also shows that affordance or object recognition abilities must, in at least some cases, be instantiated in hardware of considerable plasticity. Neither ER nor cognitivism currently possesses the ability to describe such hardware.

In summary, considerations of typical changes of perceptual abilities in lower organisms do not clearly differentiate between analytic ER and cognitivism in terms of explanatory power. They do, however, effectively rule out IMM's, Moreover, it is fairly easy to find cases that neither theory can account for using its current explanatory apparatus. In order to deal with these cases, both theories must incorporate ways of describing hardware plasticity.

Case II: Human Perception

The foregoing has considered perception in lower organisms, in which it is plausible, at least, that feature detectors respond to properties of signals that normally encode affordances. Let us now consider the case of higher organisms, where the plausibility of this claim is much less clear.

Theories of higher organisms face an additional explanandum: cognitive penetration. Cognitivism has a ready explanation of cognitive penetration: it is the result of changes in the database used by a feature detector or other inferring system. If ER is to offer serious competition to cognitivism, it must provide some account of cognitive penetration. Unfortunately, ecological realists have not discussed cognitive penetration explicitly. The account that follows is, therefore, entirely an extrapolation from typical ecological realists claims.

Accounting for cognitive penetration requires, first, that the penetrable and impenetrable components of perception be separable. The first problem for an ecological realist account of cognitive penetration is that ER has no vocabulary for talking about non-epistemic perception. The first thing that comes into contact with a signal, in ER, is a resonator, and resonators are sensitive to affordance-encoding properties. Nothing in ER corresponds to the impenetrable "percept" or "image" that serves as input to the penetrable part of perception in cognitivism. ER has, therefore, no natural vocabulary for talking about the impenetrable components of perception that presumably preceed affordance identification.

It is not clear, however, that this is a very serious handicap. Consider the physiology of the perceptual system. The mapping of excitation patterns, in mammals, from the retina to the superior colliculus, and from the superior colliculus to the primary visual cortex, is well understood (Ottoson, 1983; Gilbert, 1983).

The pattern of stimulation of the cells of the cortex that receive afferents directly from the colliculus is, up to the effects of contrast enhancement, isomorphic to the pattern of stimulation of the retinal ganglion cells. The retina-cortex pathway is, therefore, a fixed transducer. Resonators, in mammals, must accept input from this pathway, i.e. resonators must be instantiated in or after the primary cortex. It is the output of this transducer, however, that Marr (1981) characterizes as the primal sketch, the first impenetrable "image" generated by the visual system. ER can talk about this "image" by talking about the output of this pathway. One can, therefore, imagine that ER could avail itself of the neural descriptions of perceptual systems in order to talk about impenetrable components of perception. For example, in describing "seeing" a face in the clouds, an ecological realist could appeal to processes before and after the occurance of a specific excitation pattern on the superior colliculus, in the same way that a cognitivist talks about processes before and after the construction of a primal sketch.

A more serious problem for ER is explaining how what things appear to be, i.e. what affordances they appear to have, can be highly penetrable. If a person comes to be able to see something in a new way, e.g. if someone who has always seen the vase-face as a vase comes to be able to see it as a pair of faces, s/he must come to resonate to a new affordance of the thing. In order to explain penetration, therefore, ER must explain how a person can come to be able to resonate to a new affordance of a thing. This is not all, however. A change in resonance is only due to cognitive penetration if it rationally explicable in terms of the contents of the subject's other internal states. For ER, this can only mean that the change must be rationally explicable in terms of the organism's other resonant properties, including control properties. In particular, if a change is due to cognitive penetration, it cannot be explained merely by appeal to changes in the hardware.

The simplest case of resonance change is one in which a person has a resonator for a certain affordance, and comes to be able to perceive, using that resonator, that a new type of object has the affordance. If the assumption, considered in Ch. IV, that new affordances cannot be learned is accepted, the harder case, that of a person coming to be able to detect an altogether new affordance, is ruled out. If it is to retain the assumption that affordances are not learned, ER must predict that people cannot learn to detect an altogether new affordance.

Let us take the simple case first. Consider a person A who, at t(0), comes to realize that objects of type X are edible, i.e. that they have affordance E. Before t(0), A does not perceive E when looking at an X, whereas after t(0), A does perceive E when looking at an X. Assume that this change is not maturational. Assume further that these facts are independent of attunement; before t(0), A did not perceive E when looking at X whether or not s/he was looking for something that was E.

This case cannot be explained in the way that changes in the behavior of the marsh periwinkle was explained, i.e. by appeal to the relevant resonator being turned on or off. A's resonator for E can be assumed to be always on (A is always hungry). The problem is not that A is not looking for something edible, it is that s/he does not see something that is edible as being edible. The case cannot, however, be explained by saying that the modulation to which the E-resonator responds changes. This is, in effect, changing resonators; a resonator is defined mechanistically in terms of the modulation of the signal to which it is sensitive. Changing resonators is, however, ruled out by the assumption that new affordances cannot be learned.

These, however, appear to be the only possibilities for explaining resonance change in ER. Either the on/off state of the resonator must change, or the property to which it resonates must change. Cognitive penetration cannot, however, be explained by appeal to either of these processes. Resonance change, therefore, presents a serious problem for ER. It challenges ER to come up with a theory of learning that is consistent with the rest of ER, yet rich enough to account for cases in which what things are taken to have a particular affordance changes.

It is at least possible that ER can can come up with a theory of resonance change that accounts for these cases. Cognitive penetration, however, poses other, more difficult problems. The changes in resonance effected by cognitive penetrations must be rationally explicable in terms of the contents of internal states. In order to account for cognitive penetration, therefore, ER must add to an account of resonance change an account of where the information that changes the resonant properties in question comes from.

Human beings are able to perceive many affordances that they are not born able to perceive, and that are not universal in Michaels and Carello's (1981) sense of being such that "all members of a species ... share the need and have the ability to detect [them]" (p. 79). The perception of many affordances of this type is cognitively penetrable; in many cases, one learns, e.g. through instruction, to perceive them. The most serious problems appear in domains where the "affordance" that is perceived is conventionally established. Language understanding is a case in point.

Ecological realists avoid the subject of language understanding, and it is not hard to see why. If the affordances of an item are only established by convention, it is hard to see how the relation between the affordance and the signal that encodes it could be nomic, even within a personally-defined "niche." Conventions may state rules, but they do not embody natural laws. Ecological realists cannot, however, take the out that language understanding does not involve affordance perception without admitting that ER cannot explain an important perceptual phenomenon. The problem must, therefore, be faced.

Language understanding appears, however, to be a case in which no physically-specifiable property of the proximal stimulus, at any scale, correlates with the affordance perceived (Rock, 1983). There are two reasons for this. First, the semantics of tokens are fixed conventionally; any token could, in principle, mean anything.

Second, there is an enormous perceptual constancy effect; humans can, for example, identify very different shapes as tokens of 'a', or very different sounds as tokens pronunciations of 'electric' or 'nuclear'. The properties that, e.g. identify marks as tokens of the letter 'a', or that identify spoken tokens of 'electric', do not appear to be transmissible. If this is the case, learning to, e.g. recognize a new token of the letter 'a' cannot be accomplished by any change, whether due to evolution or to some form of learning, in the property to which a resonator resonates.

The traditional ecological realist response to such challenges is to claim that all cases in which there are large perceptual constancy effects, or in which the encoding of an affordance appears to be conventional, are cases in which the proximal stimulus is underdescribed. This claim often takes the form of a bald assertion. For example, in response to a challenge to account for the perception of hedonic tone (humor, etc.) in spoken language, Michaels and Carello respond that "the invariants must be very higher-order (sic) indeed" (p. 179). Moreover, the quoted passage is the entire text of the response. If this is not a mere promissory note, nothing is.

In many cases, however, this response simply stretches credulity too far. It is at least conceivable that hedonic tone could be transmissible in the medium of spoken language. It is, also, at least conceivable that correct perceptions of hedonic tone could confer a selective advantage in evolution. In many other cases, however, the analogous claims simply cannot be believed.

Consider being able to "see" what one needs to do to get tenure within two years. This is clearly an affordance; tenure affords all sorts of things. Is this affordance transmissible? No one not in the grip of a theory could believe that it is. Similarly, no one not in the grip of a theory could believe, against all of the available evidence, that one can "hear" that a spoken sentence is grammatically correct, or even that a spoken sentence has the affordance 'being true'.

The most straightforward response to this case is to claim, with Dretske (1969), that what one needs to do to get tenured is simply not "seen" in any non-metaphorical sense, and that whether a sentence is grammatically correct is not "heard" in any non-metaphorical sense. This is the obvious response for cognitivists, who hold that, since affordances can be inferred from perceptual data, they do not need to be perceived. This option is not available to ecological realists, however; they do not believe that affordances can be inferred. They do not, in fact, countenence any non-perceptual cognitive processes that could infer affordances from any other characteristics of perceptual input. They must, therefore, claim that what will get one tenured is, like every other affordance, directly perceived. There is no reason to believe that this is even possible, let alone that it is true.

This case does not misrepresent ecological realist claims; it is, indeed, not even unusual. Michaels and Carello, in their discussion of "affordances in the human-made environment" (p. 54-56) claim, for example, that the perception of "mitochondria by microscope, or of nebula by telescope" is direct, and that the perception "that a limb affords setting is directly detectable by X-rays" (p. 55). These claims are justified only by the argument that they must be true in order to be consistent with the direct perception hypothesis. They are, however, simply false. The existence of mitochondria and nebula is inferred, and that with difficulty, often from very impoverished stimuli. The ability to perform such inferences, moreover, must be learned, and is highly cognitively penetrable. The same is true for the ability to diagnose broken bones.

In the latter case, interestingly enough, the affordance perceived is clearly universal, in the sense that it would useful for "all members of the species" to be able to perceive it. Michaels and Carello must therefore, to be consistent, argue that evolution prepared humans to analyze X-ray photographs. This simply cannot be believed.

In summary, not even analytic ER can explain cognitive penetration. There are two problems with penetrable perception. First, in the case of properties that can plausibly be claimed to be transmissible in a person's niche, such as "sit-onable," or perhaps "edible," ER can provide no account of the mechanism of cognitive penetration, i.e. of the way in which one learns to see new things as having old properties. Second, many important properties that the average human can easily identify perceptually are not transmissible. ER can give no account whatever of the perception of such properties, but it cannot, with cognitivism, allow that they are not perceived. ER must, therefore, be admitted to fail the task of providing an account of these human abilities.

ER as a Theory of Transduction

The perception of many properties by humans is, however, not cognitively penetrable, as shown in Ch. II. Shape, color, relative position, and relative size, the properties that are represented, in Marr's theory, in the 2 1/2-d sketch, are among these (Rock, 1983). These are properties that must be perceived correctly if such activities as walking and grasping, as well as language understanding, etc. are to be explained. ER may well be a viable theory of the perception of these properties.

The argument used in the case of the marsh periwinkle, however, applies to the perception of these properties. ER is a true description of the process, but so is Marr's theory; the two theories are equivalent in this domain. ER can, therefore, be regarded as a special case of cognitivism that applies in this domain.

How far this special case extends depends on whether ER can give an acount of resonance changes that do not result from cognitive penetration. If so, then ER can also account for the perception of transmissible properties in cases where cognitive penetration does not occur. If ER can supply a mechanism to account for this, whether the mechanism is actually used by perceivers will be an empirical question. It is, therefore, at bottom an empirical question how much of perception can be described in terms of resonance. Both physiological data and data on cognitive penetration are likely to be useful in this regard.

The Question of Constraint

The problem of constraining both ER and cognitivism can now be reconsidered. On part of the domain of discourse - the part dealing with transduction - the two theories are equivalent. They are, therefore, equally well, or equally poorly constrained. There may be a question of constraint here, but there can be no constraint argument that drives a wedge between ecological realism and cognitivism.

In the domain where transduction is not an adequate description of perception, ER is not poorly constrained, it is empirically false. In this domain, the constraint argument misses the boat. As argued above, the size of this domain is an empirical question.

The constraint argument must, therefore, be largely an artifact of the way in which the questions making up the traditional debate are posed. It must, in other words, be largely a matter of axe-grinding. In retrospect, one can identify the mistakes that lead to the formulation of the debate in terms of constraint. The ecological realist constraint argument, as shown in Ch. II, turns on the assumption that cognitivist models must invoke internal manuals. There is considerable historical validity in this assumption; many early cognitivist models, such as that of Fodor (1968) were IMM's. Moreover, even Pylyshyn's (1984) characterization of cognitive penetrability can be easily interpreted as involving IMM's. Cognitivist models do not, however, have to be IMM's, as shown in Ch. V. The constraint argument against cognitivism does not apply if IMM's are rejected; however, IMM's must be rejected on grounds internal to cognitivism. Similarly, the cognitivist constraint argument against ER is based primarily on a lack of appreciation for the explanatory power of the resonance theory. This is also understandable historically. By emphasizing subsumption under ecological laws, ecological realists themselves did not appreciate the power or the necessity of the resonator theory. By failing to ask questions that required analytic answers, ecological realists failed to see how to constrain appeals to ecological laws.

The constraint argument, and much of the traditional debate, thus turns on misrepresentations that were committed by the very theories being criticized. The traditional debate therefore served a useful purpose; it pointed out self-imposed weaknesses in the theories involved. It did not, however, point out, or in many cases even concern, the real differences between the two theories.

Conclusion

The conclusions of this analysis can be summarized in seven statements:

1. It was suggested in the introduction that the ER-cognitivism debate does not make contact with the real issues separating ER and cognitivism. This is true. The traditional debate focusses on surface issues, many of which turn out to be red herrings. It

misses the central question: can ER provide an adequate theory of the *mechanisms* of perception?

2. The best ecological realist account of the mechanism of perception is the resonator theory, which involves functional analysis of the organism.

3. In the case of the impenetrable perception of transmissible properties, the resonator theory and Marr's theory are equivalent in type; they postulate identical functions.

4. The problem of perceptual error in natural settings is serious, and unsolvable, in traditional ER. Such errors are merely an interesting explanandum for the resonator theory, and for cognitivism.

5. In the case of the perception of transmissible properties, ER must supply a new theory of learning if it is to account successfully for changes in resonant properties. What the mechanisms of such changes in these cases actually are is an empirical question. Therefore, the extent to which ER can provide a complete theory of transduction is an empirical question.

6. ER fails to provide an account of cognitively penetrable perception, since it fails to provide an account of the perception of non-transmissible properties. ER is, therefore, at best a spe-

cial case of cognitivism.

7. In summary, ER and cognitivism are not <u>diametrically</u> opposed after all. The historical debate is, to a large extent, an artifact of internal mistakes in the two theories.

NOTES - CHAPTER VI.

1. This assumes that maturation is not triggered by a single environmental event. In cases where this occurs, the model is somewhat simpler, as one need only postulate a resonator for the environmental event driving maturation.

2. This is, in fact, a particularly persuasive example of the recognition regress argument.

REFERENCES

Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J. D. (1983) Molecular Biology of the Cell. New York: Garland.

Anderson, J. R. (1983) The Architecture of Cognition. Cambridge, MA: Harvard University Press.

Anderson, J. R. and Pirolli, P. L. (1984) Spread of activation, Journal of Experimental Psychology: Learning, Memory, and Cognition 10 791-798.

Arbib, M. A. (1972) The Metaphorical Brain: An Introduction to Cybernetics as Artificial Intelligence and Brain Theory. New York: Interscience.

Arbib, M. A. (1981) Visuomotor coordination: From neural nets to schema theory, Cognition and Brain Theory 4 23-39.

Austin, J. L. (1962) Sense and Sensibilia. London: Oxford University Press.

Ayer, A. J. (1956) The Problem of Knowledge. Baltimore: Penguin.

Barwise, J. (1981) Scenes and other situations, Journal of Philosophy 78 369-397.

Bateson, G. (1979) Mind and Nature. New York: E. P. Dutton.

Block, N. (1980a) What is functionalism? In Block (1980b).

Block, N. (1980b) Readings in the Philosophy of Psychology, v. 1. Cambridge, MA: Harvard University Press.

Block, N. and Fodor, J. (1972) What psychological states are not, Philosophical Review 81 159-181. Reprinted in Block (198b).

Borst, C. V. (1970) The Mind/Brain Identity Theory. London: Macmillan.

Boyd, R. (1980) Materialism without reductionism: What physicalism does not entail. In Block (1980b).

Bronfenbrenner, U. (1979) The Ecology of Human Development. Cambridge, MA: Harvard University Press.

Carlin, N. F. and Holldobler, B. (1983) Nestmate and kin recognition in interspecific mixed colonies of ants, Science 222 1027-1029.

Churchland, P. M. (1984) Matter and Consciousness. Cambridge, MA: Bradford/MIT.

Churchland, P. S. (1980a) A perspective on mind-brain research, Journal of Philosophy 77 185-207.

Churchland, P. S. (1980b) Language, thought, and information processing, Nous 14 147-170.

Churchland, P. S. (1983) Consciousness: The transmutation of a concept, Pacific Philosophical Quarterly 64 80-95.

Cummins, R. (1975) Functional analysis, Journal of Philosophy 72 741-760.

Cummins, R. (1977) Programs in the explanation of behavior, Philosophy of Science 44 269-287.

Cummins, R. (1978) Explanation and subsumption, Proceedings of the Philosophy of Science Association 1 163-175.

Cummins, R. (1982a) What can be learned from Brainstorms? In J. Biro and R. Shahan (eds) Mind, Brain, and Function. Norman, OK: Oklahoma University Press.

Cummins, R. (1982b) The internal manual model of psychological explanation, Cognition and Brain Theory 5 257-268.

Cummins, R. (1983) The Nature of Psychological Explanation. Cambridge, MA: Bradford/MIT.

Dennett, D. (1978a) A cure for the common code? In Dennett (1978c).

Dennett, D. (1978b) Artificial intelligence as philosophy and psychology, In M. Ringle (ed) Philosophical Perspectives on Artificial Intelligence. New York: Humanities Press. Repreinted in Dennett (1978c).

Dennett, D. (1978c) Brainstorms. Cambridge, MA: Bradford/MIT.

Dennett, D. (1983) Styles of mental representation, Proceedings of the Aristotelian Society 83 213-226.

Dennett, D. (1984) The logical geography of computational approaches, presented at the MIT/Sloan Conference on Cognitive Science, Cambridge, MA, May, 1984.

Dretske, F. (1969) Seeing and Knowing. Chicago: UNiversity of Chicago Press.

Dretske, F. (1977) Laws of nature, Philosophy of Science 44 148-268.

Dretske, F. (1981) Knowledge and the Flow of Information. Cambridge, MA: Bradford/MIT.

Dreyfus, H. (1979) What Computers Can't Do. New York: Harper and Row.

Dreyfus, H. (1981) From micro-worlds to knowledge representation, In Haugeland (1981b).

Eccles, J. (1980) The Human Psyche. Berlin: Springer.

Feldman, J. A. and Ballard, D. H. (1982) Connectionist models and their properties, Cognitive Science 6 205-254.

Fields, C. A. (1983a) Compartmental analysis as a formal language for ecological realist psychology, Nature and System 5 195-209.

Fields, C. A. (1983b) Cognitive penetration pathways in natural language understanders, Cognition and Brain Theory 6 449-461.

Fisher, J. (1974) Knowledge of rules, Review of Metaphysics 28 237-260.

Fodor, J. (1968) The appeal to tacit knowledge in psychological explanation, Journal of Philosophy 65 627-640. Reprinted in Fodor (1981).

Fodor, J. (1974) Special sciences, or The disunity of science as a working hypothesis, Synthese 28 97-115. Reprinted in Block (1980b), Fodor (1981).

Fodor, J. (1975) The Language of Thought. New York: Crowell.

Fodor, J. (1978) Computation and reduction, Minnesota Studies in the Philosophy of Science 9 229-260. Reprinted in Fodor

(1981)

Fodor, J. (1980) Methodological solipcism considered as a research strategy in cognitive psychology, Behavioral and Brain Sciences 3 63-109. Reprinted in Fodor (1981).

Fodor, J. (1981) Representations. Cambridge, MA: Bradford/MIT.

Fodor, J. (1983) The Modularity of Mind. Cambridge, MA: Bradford/MIT.

Fodor, J. (1984) Why paramecia don't have mental representations, Midwest Studies in Philosophy (in press).

Fodor, J. and Pylyshyn, Z. (1981) How direct is visual perception? Cognition 9 139-196.

Gibbs, J. C. (1979) The meaning of ecologically oriented inquiry in contemporary psychology, American Psychologist 34 127-140.

Gibson, J. J. (1950) The Perception of the Visual World. Boston: Houghton-Miffin.

Gibson, J. J. (1966) The Senses Considered as Perceptual Systems. Boston: Houghton-Miffin.

Gibson, J. J. (1979) The Ecological Approach to Visual Perception. Boston: Houghton-Miffin.

Gierer, A. (1981) Some physical, mathematical, and evolutionary aspects of biological pattern formation, Philosophical Transactions of the Royal Society of London B295 420-440.

Gilbert, C. D. (1983) Microcircuitry of the visual cortex, Annual Review of Neuroscience 6 217-247.

Gregory, R. L. (1970) The Intelligent Eye. New York: McGraw-Hill.

Grossberg, S. (1973) Contour enhancement, short term memory, and constancies in reverberating neural networks, Studies in Applied Mathematics 52 217-257. Reprinted in Grossberg (1982).

Grossberg, S. (1980) How does the brain build a cognitive code? Psychological Review 87 1-51. Reprinted in Grossberg (1982).

Grossberg, S. (1982) Studies of Mind and Brain. Dordrecht, Netherlands: D. Reidel.

Haugeland, J. (1978) The nature and plausibility of cognitivism, Behavioral and Brain Sciences 1 215-226. Reprinted in Haugeland (1981b).

Haugeland, J. (1979) Understanding natural language, Journal of Philosophy 76 619-632.

Haugeland, J. (1981a) Semantic engines: An introduction to mind design, In Haugeland (1981b).

Haugeland, J. (1981b) Mind Design. Cambridge, MA: Bradford/MIT.

Hawkins, R. D. and Kandel, E. R. (1984) Is there a cell-biological alphabet for simple forms of learning? Psychological Review 91 375-391.

Hebb, D. O. (1949) The Organization of Behavior: A Neurophysiological Theory. New York: Wiley.

Heil, J. (1981) Does cognitive psychology rest on a mistake? Mind 90 321-342.

Heil, J. (1983) Perception and Cognition. New York: Wiley.

Hubel, D. H. and Weisel, T. N. (1977) Functional architecture of macaque monkey visual cortex, Proceedings of the Royal Society of London B198 1-59.

Jantsch, E. (1980) The Self-Organizing Universe. New York: Permagon.

Kalmijn, A. J. (1971) The electric sense of sharks and rays, Journal of Experimental Biology 55 371-383.

Kalmijn, A. J. (1974) The detection of electric fields from inanimate and animate sources other than electric organs, In A. Fessard (ed) Handbook of Sensory Physiology. Berlin: Springer.

Kandel, E. R. (1976) The Cellular Basis of Behavior. San Fransisco: Freeman.

Kandel, E. R. (1979) Behavioral Biology of Aplysia. San Fransisco: Freeman.

Lara, R., and Arbib, M. A. (1982) A neural model of interaction between tectum and pretectum in prey selection, Cognition and Brain Theory 5 149-171.

Lehrman, D. S. (1970) Semantic and conceptual issues in the nature-nurture problem, In L. Aronson, E. Tobach, D. Lehrman, and J. Rosenblatt (eds) Development and Evolution of Behavior. San Fransisco: Freeman.

Levine, D. S. (1983) Neural population modeling and psychology: a review, Mathematical Biosciences 66 1-86.

Lewis, D. (1966) An argument for the identity theory, Journal of Philosophy 63 17-25.

Lindsay, P. H. and Norman, D. A. (1977) Human Information Processing. New York: Academic.

Lund, R. D. (1978) Development and Plasticity of the Brain. New York: Oxford University Press.

Marr, D. (1981) Vision. San Fransisco: Freeman.

Marr, D. (1982) Visual information processing: The structure and creation of visual representations, In D. Albrecht (ed) Recognition of Pattern and Form. Berlin: Springer.

Marr, D. and Poggio, T. (1979) A computational theory of human stereo vision, Proceedings of the Royal Society of London B204 301-328.

Marr, D. and Ullman, S. (1981) Directional selectivity and its use in early visual processing, Proceedings of the Royal Society of London B211 151-180.

Marshall, N. B. (1966) The Life of Fishes. Cleveland: World.

McClelland, J. L. and Rumelhart, D. E. (1981) An interactive activation model of context effects in letter perception, Psychological Review 88 375-407.

McDermott, D. (1976) Artificial intelligence meets natural stupidity, SIGART Newsletter 57. Reprinted in Haugeland (1981b).

Merleau-Ponty, M. (1962) The Phenomenology of Perception. London: Routledge and Kegan Paul.

Michaels, C. and Carello, C. (1981) Direct Perception. Englewood Cliffs, NJ: Prentice-Hall.

Morse, G. (1984) Mating signal mediates strong bonds in microbes, Science News 126 68.

Nagel, E. (1961) The Structure of Science. New York: Harcourt, Brace, and World.

Nagel, T. (1974) What is it like to be a bat? Philosophical Review 83 435-450. Reprinted in Block (1980b).

Neisser, U. (1976) Cognition and Reality. San Fransisco: Freeman.

Nelson, R. J. (1969) Behaviorism is false, Journal of Philosophy 66 417-432.

Newell, A. (1980) Physical symbol systems, Cognitive Science 4 135-183.

Newell, A. (1982) The knowledge level, Artificial Intelligence 18 87-127.

Newell, A. and Simon, H. (1972) Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall.

Nicolis, G. and Prigogine, I. (1979) Self-Organization in Non-Equilibrium Systems. New York: Wiley-Interscience.

Ottoson, D. (1983) Physiology of the Nervous System. New York: Oxford University Press.

Poggio, T. (1982) Trigger features or Fourier analysis in early vision: A new view, In D. Albrecht (ed) Recognition of Pattern and Form. Berlin: Springer.

Poggio, T. and Fuhle, M. (1983) Spatiotemporal interpolation in vision, In J. Beck, B. Hope, and A. Rosenfeld (eds) Human and Machine Vision. New York: Academic.

Putnam, H. (1967) The mental life of some machines, In H. Castenada (ed) Intentionality, Minds, and Perception. Detroit: Wayne State University Press.

Pylyshyn, Z. (1980) Computation and cognition, Behavioral and Brain Sciences 3 111-169.

Pylyshyn, Z. (1984) Computation and Cognition. Cambridge, MA: Bradford/MIT.

Raff, R. A. and Kaufman, T. C. (1983) Genes, Embryos, and Evolution. New York: Macmillian.

Reed, E. S. (1983) Two theories of the intentionality of perceiving, Synthese 54 84-94.

Reed, E. S. and Jones, R. (1982) Reasons for Realism. Hillsdale, NJ: Erlbaum.

Rock, I. (1975) Introduction to Perception. New York: Macmillian.

Rock, I. (1983) The Logic of Perception. Cambridge, MA: Bradford/MIT.

Rorty, R. (1979) Philosophy and the Mirror of Nature. Prineton: Princeton University Press.

Rosenthal, M. P. (1965) Fundamentals of Radio. New York: Rider.

Runeson, S. (1977) On the possibility of 'smart' perceptual mechanisms. Scandanavian Journal of Psychology 18 172-179.

Rus, T. (1979) Data Structures and Operating Systems. New York: Wiley.

Salm, W. G. (1966) Abc's of Modern Radio. Indianapolis: Sams.

Searle, J. (1979) What is an intentional state? Mind 88 74-92.

Searle, J. (1980) Minds, brains, and programs, Behavioral and Brain Sciences 3 417-457.

Searle, J. (1981) Analytic philosophy and mental phenomena, Midwest Studies in Philosophy 6 405-424.

Searle J. (1983) Intentionality. Cambridge. Cambridge University Press.

Shannon, C. (1948) The Mathematical Theory of Comminication. Urbana, IL: University of Illinois Press.

Shaw, A. C. (1974) The Logical Design of Operating Systems. Englewood Cliffs, NJ: Prentice-Hall.

Shaw, R. E. and McIntyre, M. (1974) Algoristic foundations to cognitive psychology, In W. Weimer and D. Palermo (eds) Cognition and the Symbolic Processes. Hillsdale, NJ: Erlbaum.

Shepard, R. N. (1984) Ecological constraints on internal representation, Psychological Review 91 417-447.

Skinner, B. (1953) Science and Human Behavior. New York: Free Press.

Sterling, P. (1983) Microcircuitry of the cat retina, Annual Review of Neurosciences 6 149-185.

Stich, S. (1983) From Folk Psychology to Cognitive Science. Cambridge, MA: Bradford/MIT.

Turvey, M. T., Shaw, R. E., Reed, E. S., and Mace, M. W. (1981) Ecological laws of perceiving and acting, Cognition 9 237-304.

Ullman, S. (1980) Against direct perception, Behavioral and Brain Sciences 3 373-415.

Ullman, S. (1983) Recent computational studies of the interpretation of structure from motion, In J. Beck, B. Hope, and A. Rosenfeld (eds) Human and Machine Vision. New York: Academic.

Villwock, W. (1973) Order: Anglerfish, In Grzimek's Animal Life Encyclopedia, V. 4. New York: van Nostrand Reinhold.

von Uexkull, J. (1957) A stroll through the worlds of animals and men, In C. Schiller (ed) Instinctive Behavior. New York: International Universities Press.