Library of Congress Cataloging in Publication Data

On the Nature of Representation

A Case Study of James Gibson's Theory of Perception

Mark H. Bickhard and D. Michael Richie
with the assistance of
Robert Hughes and James Dannemiller

PRAEGER

PRAEGER SPECIAL STUDIES · PRAEGER SCIENTIFIC

New York · Philadelphia · Eastbourne. UK Toronto · Hong Kong · Tokyo · Sydney

Bickhard, Mark H.
On the nature of representation.
Bibliography: p.
Includes indexes.
1. Perception. 2. Gibson, James Jerome, 1904-1979. I. Richie, D. Michael. II. Title.
BF311.B4981983 153.7'092'4 83-13365
ISBN 0-03-069526-0 (alk. paper)
.ft

Published in 1983 by Praeger Publishers CBS Educational and Professional Publishing a Division of CBS Inc. 521 Fifth Avenue, New York, NY 10175 USA

3456789 052 987654321
Printed in the United States of America on acid-free paper

Library of Congress Cataloging in Publication Data

Acknowledgments

Grateful acknowledgment is made for permission to quote selections from the following works.

Fodor, J., & Pylyshyn, Z. How direct is visual perception? Some reflections on Gibson's ecological approach. *Cognition,* 1981, 9, 139-196. Gibson, James J. *The perception of the visual world.* Boston: Houghton Mifflin Company. Copyright © 1950, renewed 1977 by Eleanor J. Gibson and Pearl K. Carmichael.

Gibson, James J. *The senses considered as perceptual systems.* Boston: Houghton Mifflin Company. Copyright © 1966.

Gibson, James J. *The ecological approach to visual perception.* Boston: Houghton Mifflin Company. Copyright © 1979.

Ullman, S. Against direct perception. *The Behavioral and Brain Sciences*, 1980, **3,** 373-381.

Other debts of this book extend most especially to Robert Campbell, Robert G. Cooper, and Ernst von Glasersfeld for helpful and stimulating comments on earlier drafts of the manuscript, and to Tom Hoeffner for the design and creation of the indexes.

Preface

My awareness of the centrality of representation and of the untenability of standard (encoding) assumptions regarding the nature of representation has been developing for some time. Actually, an alternative, nonencoding, approach to representation was developed first, and only later did the depth of the incoherence in the encoding approach begin to become apparent. Initially, this awareness of the untenability of the encoding approach was stimulated by reflections on the impossible difficulties encountered in trying to account for language within the alternative—nonencoding—model: language too is standardly conceived of as an encoding phenomenon, and it simply did not fit the alternative model, despite years of trying. The result was a distinctly nonstandard conceptualization of language.

In conversations a few years ago with Richie, Hughes, and Dannemiller, the relationship of this issue of the nature of representation to perception was explored. In particular, Gibson's theory seemed to constitute an implicit internal struggle between the encoding and nonencoding perspectives. As such, his theory seemed interestingly illuminated by an explicit understanding of the two approaches and their relationships to each other, and, in turn, it served to illustrate the more general issues in a particular domain and historical context. Out of these discussions came a paper focused on Gibson, written primarily by Richie and myself. It soon became clear, however, that the points about Gibson per se could not be adequately understood without a more thorough understanding of the general analytic perspective on representation that was being brought to bear. Consequently, I set out to write a more thoroughly contextualized discussion of Gibson. The result is this book, in which Gibson's theory and the encoding-interactive perspective on representation serve respectively as illustration and illumination of each other.

Mark H. Bickhard

Introduction

THE CONTEXT

Representation is fundamental to the study of the mind. It is one of the essential differentiators of the mental from the nonmental. The corresponding importance of issues regarding representation is manifested both historically and in contemporary research: Representation has occupied the attentions of students of the mind for millenia, and representation permeates psychology. It is central to perception, cognition, and language, and deeply involved in all other areas as well. The broadest focus of this book is on the nature of representation.

In particular, two alternative approaches to the nature of representation are examined. The first, the encoding approach, assumes that representation is some form of encoding. In its strong form, it assumes that encoding is the essence of representation. In its weak form, it assumes that encodings constitute one logically independent and irreducible form of representation, though there might be others. The encoding approach to representation has been dominant throughout history, and it is still so today. It is not often recognized that an alternative even exists.

The alternative, the interactive approach, eschews encodings in any foundational role in favor of a conception of representation as an aspect of successful interaction. Its historical origins are relatively recent, and its recognition as a distinct alternative is incomplete and still evolving. A number of people (e.g., Piaget) propose a kind of hybrid between the encoding and the interactive approaches; in effect, these constitute weak versions of the encoding approach. A central theme of our examination is a critique of the encoding approach, in both of its versions, and a presentation and espousal of the interactive alternative.

Issues concerning the nature of representation permeate all of psychology, but they can become especially acute with regard to perception. It is relatively easy in many areas of psychology to avoid basic representational issues by pushing them off to some other area where the representations of concern are presumed to originate. Perception, being the presumed ultimate source of most representations, finds it more difficult to sidestep or postpone the basic issues involved. Correspondingly, perception is a central focus within the broader concern with the nature of representation.

More specifically, the two approaches to representation, the tensions and confusions between them, and the critiques involved are interestingly and deeply exemplified in James Gibson's theory of perception and in the controversy surrounding it. A dominant part of the analysis to follow is a case study of the Gibsonian theory and controversy from the perspectives of the encoding and interactive approaches to representation. The goals of this case study include an elucidation both of the issues involved between the two approaches and of the central themes and insights of Gibson's theory.

A failure to recognize the distinction between the two approaches to representation has infected both sides of the Gibsonian controversy, primarily to the detriment of Gibson's deepest insights. The confusion between the two approaches plays itself out, for example, in Gibson's blanket criticisms of mental-processing theories of perception, not recognizing that his own theory is a version of a mental-processing theory, and in Gibson's critics not recognizing that Gibson's own criticisms devastate the very model that they take him to be espousing—as well as their alternatives. The common error here is the presupposition that the encoding approach is exhaustive, the failure to explicitly recognize the interactive alternative.

After the analysis of the Gibsonian theory and controversy, the focus shifts back to the broader implications of the critique of the encoding perspective and of the interactive alternative. These implications pursue the phenomena of representation throughout all of psychology. The discussions, consequently, are of a more general and illustrative nature.

STRUCTURE OF THE BOOK

The book is organized in four major sections. The first introduces the foundations of the ensuing analyses. This includes both the initial presentation of the encoding and interactive approaches and some of the relationships between them, and the presentation of the basic themes and historical development of Gibson's theory. A central aspect of that historical development is the progressive deepening of Gibson's interactive insights.

The second major section focuses on Gibson's metatheory and metatheoretical critiques. It is here that Gibson's failure to recognize the interactive approach as a distinct approach shows up most strongly. In general, Gibson's arguments are telling, but the conclusions are overdrawn: They are stated as applying against all mental-processing models, but they apply only to encoding models, not to interactive models. In particular, they do not apply to Gibson's own interactive mental-processing theory.

The third section is devoted to two major published criticisms of Gibson. The central themes are that the criticisms are valid as applied to a common and quite plausible interpretation of Gibson, an encoding interpretation, but that they are invalid as applied to a deeper interactive interpretation and that the alternatives to Gibson offered within these criticisms—different versions of the encoding approach—are themselves untenable. The interactive interpretation of Gibson, then, is both more perspicacious as an interpretation and more defensible as a model of perception.

The fourth, last, and largest of the major sections closes out the discussion of Gibson per se and reintroduces the issues of representation in general. Further potential criticisms of the interactive approach are considered; some of the deeper conceptual foundations of and relationships between the encoding and interactive approaches are discussed; and some of the consequences of the interactive perspective are illustrated. The brief discussions of consequences include the role of encodings *within* an interactive approach and an analysis of functionalism as an approach to mental phenomena.

Contents

On the Nature of Representation

Chapter One

Foundations

The central theme of this book is the rivalry between the encoding and the interactive perspectives on the nature of representation. The first subsection on foundations provides the initial, introductory presentations of these two positions and some of the criticisms and relationships between them. A major claim is that the encoding approach is impoverished relative to the interactive approach—that the interactive approach provides much richer resources for understanding and explaining representation. In later sections, it is also argued that the encoding approach is an asymptotic limiting case of the interactive approach, derivable by taking certain interactive characteristics to their asymptotically unreachable limits.

The encoding and interactive perspectives are involved somewhat differently in Gibson's theory and in his metatheory. Gibson's theory shows a progressive evolution away from encoding conceptions, toward interactive insights. Gibson's metatheory shows a consistent critical stance toward mental processing theories of perception in all of their guises; Gibson's metatheoretical *arguments,* however, are valid only against encoding-based mental-processing theories, not against the interactive approach to mental processing. A metatheory that contained telling criticisms against encodings may have stimulated Gibson's theoretical evolution toward interactionism, while a metatheory that failed to explicitly recognize interactivism as an alternative to the encoding perspective may have simultaneously inhibited the full flowering of that evolution. A historical summary of Gibson's theory is presented in the second foundational subsection. Gibson's metatheory is addressed in a later major discussion of its own.

The central theme in the historical summary is Gibson's movement away from an initial, simplified encoding conceptualization of perception in accordance with his progressively deepening interactive insights. It is essential to understand this transcendence of his early model in order to understand Gibson's mature theory, yet both the fact and the nature of that transcendence are commonly overlooked. In this and later sections, it is suggested that Gibson's early position was very similar to the impoverished, direct encoding model with which he is still charged. The incomplete development of Gibson's final position, together with his continued overstatement of his criticisms of alternatives, obscured the fact and nature of Gibson's transcendence of that early position, thus allowing Gibson's critics to continue to attack a position that he had largely abandoned and generating an air of mutual incomprehensibility in the Gibsonian-antiGibsonian debate.

The presentation of foundations ends with a pair of conceptual tools to be used in later analyses. The first of these is a simple structure of conceptual distinctions, which relate closely to the encoding-interactive contrast. These distinctions underlie many of Gibson's criticisms of alternative approaches to perception. The second such conceptual tool is a brief outline of an interactive approach to perception. The primary purpose of this outline is to serve as a counterexample to some of Gibson's metatheoretical criticisms, thus demonstrating that their proper scope is not as broad as is usually stated and providing a framework and an anchor for the interactive interpretation of Gibson's own theory.

ENCODINGS AND INTERACTIVE PERSPECTIVES ON REPRESENTATION

In this section, the distinction and rivalry between the encoding and interactive perspectives on representation are introduced. The discussion is greatly inspissated, being primarily an initial framework to be added to and elaborated upon throughout the rest of the book. Other relevant discussions are contained in Bickhard (1980a, 1982). A prominent theme in this initial discussion is the sense in which the encoding approach is impoverished relative to the interactive approach. This is relevant at this point in that Gibson is most commonly construed as proposing an impoverished version of standard encoding models, just the opposite of what is argued in our analysis.

The intuition of encoded representation is that of a resemblance between the representation and that which it represents. The paradigm cases are paintings and statues. The essential idea in such resemblance is that of correspondence: The structure of the representation corresponds to the structure of the represented. In the paradigm cases, the correspondences are relatively direct and simple, but the basic idea can be elaborated into highly complex, sophisticated, and abstract correspondences without changing the critical characteristics of the approach: Simple encoding models generally rest upon point-to-point correspondences, as with pictures, but more sophisticated versions can involve features, properties, concepts, propositions, and so on. At whatever level of sophistication, however, the essential idea of encoding is that of representational correspondence between representing elements and represented elements.

Any structural correspondence involves two parts: a correspondence of basic elements and a correspondence of relationships among those elements. A structural correspondence is composed of such an element-by- element correspondence together with a correspondence of the relationships among the elements. Encoding correspondences, therefore, involve both element and relationship representations, with the elements in a foundational building-block role. The distinction between elements and relationships is critical to any actual encoding model, but the issues with which we are primarily concerned here involve either the elements per se or the broader idea of an encoding correspondence, independent of whether it is with respect to elements or to relationships.

There are two general versions of the encoding perspective. The strong version holds that encodings are the essence of all representation. The weak version acknowledges the existence of other independent forms of representation, but it holds that encodings themselves constitute a necessary and logically independent form of representation. The focus in the following discussions is primarily on the weak version and its deficiencies rather than on the strong version: If encodings are logically dependent on other forms of representation, then they are not necessary, and they certainly cannot be the essence of representation. That is, if the weak version is untenable, then so also is the strong version.

Arguments against the sufficiency of encodings to account for all representation are telling against the scope of the strong version. Such arguments include the claim that there is no ultimate atomic representational level at which the basic encoding elements can be defined, that no list of encoding elements can be adequate, and every attempt to make one so yields an ad hoc proliferation of elements, that there is no possible origin for new encoding elements and thus none for new kinds of knowledge, and that there are some kinds of knowledge, if not representation, such as skills, that cannot be completely captured as encodings. These arguments, like the strong version toward which they are directed, do not receive much attention. There is one argument against sufficiency, however, that is elaborated somewhat more fully, because it leads directly into one of the major arguments against the weak version. That is, it leads directly into an argument against the necessity of encodings.

Representation is a functional concept. An encoding represents something only insofar as it can *function* as such a representation for some agent. An encoding is presumed to represent to an agent what is and is not the case, and what is and is not possible, thereby influencing the processes and interactions of that agent. The point is that encodings are never representationally sufficient unto themselves; they always require an agent as an interpreter. Encodings are not only representations *with respect* to such interpreters, but also *relative* to interpreters and their interpretations: It is quite possible for a particular encoding to receive differing interpretations.

Thus, encodings are necessarily contextualized within interactive agents as interpreters. Encodings function as representations only insofar as they influence the flow of processing inside such an interactive system. But any such influences on system processing could in principle be constituted directly in the organization of the system. Encoding influences on system processing are composed of selections at various points among alternative paths of processing. If, at those points, the system were already differentiated in accordance with the alternatives available, and if the system were already *in* the appropriate differentiated condition to flow directly into the alternative' that would otherwise be selected by 'interpreting' the encoding, then the encoding would be superfluous. Such a condition would be obtained if, instead of developing (or being constructed with) the ability to set up and interpret such encodings in the first place, the system differentiated its organization in accordance with the possible selections of such encodings, and, instead of setting up one of those encodings in any particular instance, it simply entered the appropriate differentiated condition. The representational influences of encodings, thus, can in principle be incorporated into the organization of the system that would otherwise be the interpreter.

Such *state splitting* in lieu of encodings can quickly become combinatorially very complex, but it is always possible in principle. Thus, the differentiation of explicit encodings in a system may be desirable, perhaps for reasons of efficiency, but it is never logically necessary. Encodings, then, though perhaps desirable, are always logically eliminatable and, therefore, they cannot be logically independent forms of representation. Encodings are always subordinate to, and in principle eliminatable within, an appropriate interpreter.

An encoding represents only insofar as it represents something in particular; thus, only insofar as it is taken to represent that something by an interpretive agent. The second, and perhaps deepest, argument against logically independent encodings is a challenge to the possibility of a logically independent encoding representing *anything* in particular and, thus, encoding anything at all. That is, it is a challenge to the coherence of the concept of a logically independent encoding.

What a logically independent encoding represents can only be specified via that encoding itself. To do so in any other way is to specify it in terms of something else, some other representation, which is to render it logically dependent and, in principle, eliminatable in favor of that something else. But this yields the result that what a logically independent encoding represents can be specified only circularly, in the form: 'It represents whatever it represents'. This is incoherent as a specification of an encoding. The point is not that encodings cannot exist; they clearly do and are easily definable *in terms of some other representation* in the general form: 'This represents the same as what that represents.' The point, rather, is that *logically independent* encodings—encodings independent of any other representation—cannot exist. Encodings, therefore, cannot be a basic or irreducible form of representation.

The general nature of the alternative interactive approach has already been indicated in the earlier reference to an interactive agent. Knowledge is constituted as goal-oriented interactive competence, and representation is a functional aspect of such competence: Interactions and interactive systems that are not appropriate to an environment, that are not sensitive to that environment and to its potentialities, will not be competent in that environment. The interactive claim is that such interactive sensitivity, such ability to take into account an environment, its potentialities, and its changes, is the fundamental form of all representation. All other representation is constituted as differentiations and specializations of this aspect of interactive competence.

It is not logically necessary that explicit representations be differentiated and specialized in such an interactively competent system but, if they are, the first step will not be symbolic encodings, because of the necessary logical dependence of such encodings. Instead, the most fundamental form of explicit representations will be a system of internal bookkeeping for the interactive processes. Such bookkeeping will be in terms of the internal course and outcomes of some interactions, which may in turn be useful in determining the course of other interactions. Note that such bookkeeping does not itself constitute a system of symbolic encodings, but rather a system of *indicators.* An element in such a system at most indicates that such-and-such an internal outcome of an interaction has been reached, an indication that might be useful for some other interaction. It does not directly encode or symbolize anything, although it does indicate that an environment sufficient to that interaction outcome obtained. Clearly, however, derivative encodings can be defined in terms of such indicators.

We have argued that encodings are logically unnecessary and that the concept of logically independent encodings is incoherent, and we have briefly introduced a sense of an interactive alternative approach to representation. The two approaches are not simple alternatives however: they have deep relationships with each other. We first examine some ways in which the encoding approach is impoverished relative to the interactive approach and then some ways in which it appears to be the other way around from within the encoding perspective. Other relationships are examined later in the book.

First, note that although an interaction outcome indicator represents an environment sufficient to that outcome, it does not encode any particular characteristics of such an environment. The system, in fact, may have no knowledge whatsoever of any property of that environment other than its sufficiency to that outcome. However, as other parts of the system come to make appropriate usage of such outcome indicators in their own interactions, those potential usages, those interpretations, constitute knowledge of characteristics of such an environment. An encoding, on the other hand, represents within itself all of its possible interpretations; it *encodes* them—either explicitly, in its substructure of encoding elements, or implicitly, in its position within a general coding scheme. An encoding, therefore, is essentially an indicator together with all of its interpreters—and all of its possible interpreters—or, equivalently, it is a representation that, impossibly, needs and has no interpreters at all. The encoding approach, then, is necessarily impoverished in its treatment of the process of interpretation.

Second, an interaction outcome indicator can be the product of an interaction of indefinite complexity, involving procedures with unbounded organizations and hierarchies of decisions, subprocedures, and goals. An indicator may meaningfully depend on the whole of such an interaction, including its organization, without being determined by any, or even an exhaustive set, of its 'parts.' An encoding, on the other hand, is some sort of a direct record or copy or transformation of whatever it encodes. As the processes for detecting that which is to be encoded become more and more complex, the concept of encoding as direct transformation or transduction becomes more and more obviously untenable. Convergently, as the detection of that which is to be encoded becomes increasingly interactively complex, it becomes increasingly obvious that encoding is not fundamental, but rather is derivative from interactive detection (a version of interactive representation). Correspondingly, encoding models are inevitably grounded upon simple transformation or transduction steps that, with a little hand waving, appear to pass as direct, one-to-one, and certain—as nonderivative encodings. Essentially, the detection and identification process is functionally simplified to the point where it seems to epistemologically disappear. Correspondingly, the encoding approach is necessarily impoverished in its treatment of the processes of detection and identification. Complexity can be added to such a model by the addition of further stages of the encoding processing, so long as those stages do not critically involve further *interaction,* for to do so would be to undercut the entire conception of transformations-of-encodings. Further stages of processing *must* generally be added in order to somehow enhance the encodings with information that is not easily conceived of as being 'directly transduced.' Such multistage processing models based on some level of presumed basic encodings are the dominant approach in perceptual theorizing. The presupposition in such models, of course, is that such interactively simple conceptualizations are both possible and sufficient to the facts.

Third, new interaction outcome indicators, thus new representations, thus new potential knowledge, can originate with the construction of new procedures for interaction. The potentialities of such new constructions are ultimately bounded only by the intrinsic constraints on the nature of processes (e.g. Davis, 1958, Rogers, 1967), and this point holds both phylogenetically and ontogenetically. Encoding models, on the other hand, are restricted in their representational capacity to combinations of some basic level of encoding elements. The origin of new encoding elements is impossible because there is nothing for them to originate out of except already existing encoding elements, in which case they are not new elements at all, but simply new combinations of old elements. Thus, there must be a basic level of encoding elements, and all other representations must be combinations of these basic elements. The presuppositions, of course, are that such a basic encoding level is possible, that it exists, and that such a combinatorial constraint is adequate to the facts. It should also be noted that, as a consequence of the impossibility of the origin of new encoding elements, theorists within an encoding approach are ultimately driven to espouse an innatism of a basic, but combinatorially adequate, level of knowledge and representation (e.g., Chomsky, 1965; Fodor, 1975). Such an innatism, however, only pushes the logical problem of the origin of new encoding elements back into phylogeny, and it is no more solvable there than it is in ontogeny. The encoding approach, then, is necessarily impoverished in its treatment of representational adequacy and of representational innovation.

In these three senses, as well as others, encoding approaches are impoverished and limiting cases of interactive approaches, rather than the other way around. From a presupposed and unexamined encoding perspective, however, things appear differently. From an implicit and implicitly exhaustive encoding perspective, objections to the idea of transformations- of-transduced encodings can only lead to the conclusion that the objector is claiming that everything, no matter how complex, is somehow *directly* encoded, with no intermediate processing; such objections must be to the mediations, for the fact of encodings is unexaminedly obvious. Such an impoverished direct encoding model is obviously absurd and impossible, and one wonders why anyone in his right mind would seriously espouse it. The wonder derives from a lack of understanding that there is an alternative to the encoding approach, from a lack of understanding that the deepest objections are to the encoding approach per se, not to the mediations within that approach.

Similarly, from the encoding perspective, to break a problem of representation down into subproblems is to break an encoding down into subencodings or into prior encodings. Any objection seems to amount to the absurd and impossible claim that representation *has* no components, no differentiable subproblems. From an interactive perspective, however, to break a problem of representation down into subproblems is to break an interactive procedure down into subprocedures or into constructively prior procedures, none of which will necessarily involve intermediate or consequent encodings. The entire Gibsonian controversy has been bedeviled by the failure to recognize the existence and relationships of these two alternative approaches to the nature of representation, with the consequence that the encoding perspective is presupposed to be exhaustive.

The discussion to this point of the encoding and interactive perspectives has been greatly condensed and abbreviated. It is sufficient, however, to introduce the analysis of Gibson's theory, and it is elaborated upon throughout the remainder of the book.

A HISTORICAL SUMMARY OF GIBSON S THEORY

Gibson's theory and his stance toward contemporary approaches to perception cannot be understood without attention to their historical context of origin and subsequent development. We review that context and development in order to highlight Gibson's early reactions against then current encoding-type models and his later movement toward an interactive model. A major part of our contention is that those early reactions have never been properly disentangled from and differentiated with respect to his later development.

Gibson (1950) points out that the study of perception had long been dominated by the problem of how the mind can generate our full experienced perceptual knowledge from the inadequate data provided by the senses, with vision and the eyes always the primary focus. The major approaches to this problem were based on the works of Berkeley (1709/1922), Muller (1838/1948), and Helmholtz (1896/1952), who proposed that the eyes directly receive and encode certain basic sensations, such as patches of color, lines, points, and so on, and that full visual perceptions are then constructed on the basis of such sensations through various processes of comparisons with memory, inferences based on cues within the sensations, and, ultimately, judgments concerning the nature of the external stimulus. Differences among theories generally were concerned with the nature and identity of the basic sensations and of the subsequent processing required to generate perceptions of the visual world. But all such models, including a slight variant in which the retinal image served in the role of sensations, assume that perceptions must be generated out of primitive sensations or retinal images. They assume that the senses receive fragmented or incomplete information about the world that must be enriched by mental processing (Gibson & Gibson, 1955).

The Gestaltists objected to this approach, arguing that the sensory elements seemed impossible to specify and that, in any case, such an approach “could never really explain how we *see* the world. . . but only how we can make *judgments* about the world" (Gibson, 1950, p. 22). The Gestaltists argued that "experience is not reducible to elements or additive units" and proposed instead that the process of perception "was one of a relatively spontaneous *sensory organization”* (Gibson, 1950, p. 22). Unfortunately, the concept of sensory organization was much less applicable to the perception of space than it was to the perception of form, and it proved difficult to specify in either case. Gibson suggests that the major contribution of the Gestalt theorists was that they formulated genuinely relevant problems for space perception, problems concerning the characteristics of the actual experienced visual world rather than the flat geometric visual field (Gibson, 1950, p. 23).

In the context of sensation or retinal-image-based theories on the one hand and of Gestalt theories on the other, Gibson began his own questioning of theories of perception while conducting experiments in depth perception during World War II (Gibson, 1950, 1966, 1979). His basic conclusion at this point was that depth perception was more accurate than could be explained by any model based on depth cues. Such an observation of the inadequacy of cues or sensations leads rather directly to a general questioning of sensation-based approaches, but Gibson found the available alternative of Gestalt sensory organization to be inadequate as well.

By the time he wrote *The perception of the visual world* (1950), Gibson had gone beyond both alternatives. From the Gestaltists, he accepted and adapted the idea that the most basic problems of visual perception were those regarding the experienced three-dimensional visual world, not the flat geometric visual field, but he rejected the proposed process of sensory organization. From the sensation-based approaches, he accepted very little, neither their basic problems nor their basic solutions.

Gibson argued that people and animals "appear to react to the spatial environment with an accuracy and precision too great for any known theory of space perception to be able to explain. ... If the solid visual world is a contribution of the mind, if the mind constructs the world for itself, where do the data for this construction come from, and why does it agree so well with the environment in which we actually move and get about?" (p. 14). This basic rejection of mental constructivism, of mental processing, was one of the most fundamental moves in the development of Gibson's own theory. Consistent with this rejection, Gibson also rejected the premise that made such processing necessary and the particular distinctions and processes by which it was presumed to occur.

In particular, and most fundamentally, Gibson rejected the basic premise that the data available to the senses were inadequate to perception: "Even complex perceptual qualities must have stimuli" (p. 8); "If the total stimulation contains all that is needed to account for visual perception, the hypothesis of sensory organization is unnecessary" (p. 25). Clearly, if the total stimulation contains all that is necessary to account for visual perception, then the (unconscious) inferences, comparisons with memory, and judgments—the mental processing—of the sensation-based models are also unnecessary. If we ask the right question, Gibson suggests, if we ask about the experienced visual world based on surfaces and edges, rather than about the flat geometric visual field, then we find that the information available to the visual senses is sufficient to perception, and information enhancement via mental processing is a superfluous and flawed postulate.

Rejection of mental processing in this enhancement sense entails a rejection of the classical distinction between sensations and perceptions; that distinction is based on the assumptions that sensations are informationally impoverished and that mental processing enriches them into perceptions. "Obviously these terms will have to be either discarded or redefined" (Gibson, 1950, p. 11).

Gibson also rejected the passivity of the perceiver, which was generally assumed in sensation-based models: The perceiver needed merely to passively and statically receive sensations in order to perceive. Gibson argued, however, that "the normal human being is active. His head never remains in a fixed position for any length of time except in artificial situations. If he is not walking or driving a car or looking from a train or airplane, his ordinary adjustments of posture will produce some change in the position of his eyes in space" (1950, p. 117). Such a stance was strongly consistent with Gibson's emphasis on perception as a process that can only be understood in terms of its natural ecology. In further support of that stance was Gibson's argument that "such changes [of motion of the perceiver] will modify the retinal images in a quite specific way" (p. 117). In particular, Gibson found that the activity of the perceiver provided powerful information for depth perception in the form of motion parallax.

An additional powerful argument that Gibson formulated against sensation-based models is the homunculus problem (1950, 1966, 1979), perhaps best illustrated by his criticism of retinal-image theories, which assume that a person must process the stimulation on the retina. Although Gibson had already formulated the argument by 1950 (p. 54), we quote a later, more developed statement of it.

> [It is] tempting to believe that the image on the retina falls on a kind of screen and is itself something intended to be looked at, that is, a picture. It leads to one of the most seductive fallacies in the history of psychology—that the retinal image is something to be seen. I call this the "little man in the brain" theory of the retinal image (1966, p. 226), which conceives the eye as a camera at the end of a nerve cable that transmits the image to the brain. Then there has to be a little man, a homunculus, seated in the brain who looks at the physiological image. The little man would have to have an eye to see it with, of course, a *little* eye with a *little* retinal image connected to a *little* brain, and so we have explained nothing by this theory. We are in fact worse off than before, since we are confronted with the paradox of an infinite series of little men, each within the other and each looking at the brain of the next bigger man (1979, p. 60).

There are a number of variants of this argument, corresponding to variants in sensation-based models: The commonality among them is that something, or someone, must ultimately do the perceiving, and that is what was to be accounted for in the first place. Such an argument applies to any form of inputs-followed-by-processing-followed-by-perception model. Although Gibson seems to have been most strongly persuaded by his argument that mental constructivism is too weak, too prone to variation and error, to account for observed accuracy, the homunculus problem appears to be his logically strongest and philosophically deepest criticism of sensation-based models.

Gibson's attack on mental constructivism as inadequate to observed accuracy and as subject to the homunculus problem was a rejection of the most fundamental assumption of sensation-based models concerning the *process* of perception. His assertion that the total stimulation is informationally adequate to perception was a rejection of the most fundamental assumption of sensation-based models concerning the *problem* of perception: the problem of how full perceptions are derived from impoverished sense data. Gibson was rejecting both the form of, and the need for, sensation-based models as forms of transformations-of- transduced-encoding models, and his arguments were powerful enough to span the entire domain of such models, including later input-processing models that were not strictly sensation or retinal-image models. Gibson continued to develop his arguments against sensation-processing and other input-processing models; we later focus more explicitly on the proper scope and validity of some of these arguments.

The model that Gibson (1950) proposed to replace sensation-based and Gestalt approaches to perception might be described as an ecological direct-encoding model. Gibson rejected the sensation-based conception of the perceiver as a passive individual confronting a flat visual field in favor of an active perceiver confronting an ecologically structured visual world—thus, an ecological model. He also rejected both the mental constructivism of the sensation-based models and the sensory organization of the Gestaltists in favor of a direct correspondence between stimulation and perception—thus, a direct encoding model.

The direct encoding aspect of Gibson's 1950 model was both a methodological move and a theoretical move. It was methodological in that he proposed a sort of ecological psychophysics, "establishing an empirical correspondence between the stimulus and its conscious resultant" (p. 52) as the basic problem of perception. It was theoretical in that he rejected any intermediate processing of encoded sensations between the stimulation and the perception and in his corresponding rejection of the sensation- perception distinction.

It is not entirely clear that Gibson would have agreed with the "encoding" part of our designation of his model as an "ecological direct- encoding," especially in his later career. In the light of his criticisms of sensation-based processing, however, and in the absence *of any explicit* alternative conceptualization of perception, his 1950 model at least seems committed to some form of a direct encoding model. Certainly it is as a direct encoding model that Gibson's theory is most commonly and most resoundingly criticized (as we argue later), though it is the subassumption of directness that is attacked, rather than the primary presupposition of encoding.

A direct encoding model poses the obvious question of how such encodings could occur. Gibson's conceptualization of an ecologically active perceiver contains the germ of his later answers to that question and, we argue, the germ of interactive insights that allowed him to largely transcend the encoding approach altogether. In pursuit of this point, we turn now to Gibson's later development of his model.

There is a powerful consistency in that development: The major features of Gibson's later model are all developments of the internal implications of Gibson's 1950 position. His theoretical evolution manifests a deep faith in and commitment to the unfolding of the internal logic of his original insights.

Gibson's 1950 statement that the total stimulation was sufficient to perception was still consistent with a retinal-image-based model of perception, just so long as the retinal image was not considered to require informational enhancement or enrichment. Correspondingly, Gibson wrote *The perception of the visual world* (1950) largely in terms of a retinal-image-based conception of perception. He described his psychophysics program as involving a "jump from the retinal image directly to the perceptual experience" (p. 51).

Such a retinal-image focus, however, was not tenable in conjunction with Gibson's ecological emphasis on the importance of the active perceiver.

The retinal image of an active perceiver changes too much, too fast, and too continuously, in contrast with relatively stable perceptions, to be the primary locus of perception. "The neural input of the mobile eyes in the mobile head of a mobile animal . . . cannot be thought of as the anatomical pattern of the nerve cells that are excited in the fiber bundle. This anatomical pattern changes from moment to moment" (1966, p. 4). "The active observer [however] gets invariant perceptions despite varying sensations" (1966, p. 3).

A different locus of perceptual information was required, one that maintained a stability comparable to that of perceptions and one that was adequate to those perceptions. A new perceptual locus was required by Gibson's recognition of the importance of the active perceiver; such a locus was suggested by that same recognition. Gibson's original emphasis on the active perceiver stemmed in part from the motion parallax information concerning depth that was thereby derived. Motion parallax is a phenomenon of the structure of the ambient light through which the eye moves. The clear suggestion is that the broader spatial and temporal patterns in the ambient light might well be the actual locus of visual perception.[1] Certainly, on the one hand, there is no information available in the retinal image that is not available in the ambient light, and, on the other hand, it is difficult to conceive what alternative external locus for visual perception might be possible. Furthermore, very encouraging success was obtained in investigating the information that was in fact available in the ambient light. Correspondingly, "In my book, *The Perception of the Visual World* (1950), I took the retinal image to be the stimulus for an eye. In this book I will assume that it is only the stimulus for a *retina* and that ambient light is the stimulus for the visual system" (1966, p. 155).

Thus, consideration of the fact and necessity of the active perceiver forced a shift in the postulated locus of visual perception from the retinal image to the ambient light. Consideration of the ambient light as the locus of perception forced, in its turn, a reciprocal revision of the conception of the perceiver. The logic of the second revision derives from the fact that such broader spatial and temporal patterns in the ambient light cannot simply be sought by the visual system, then, when found, statically, retinally perceived. They are, by definition, too big for that. They must be scanned, sampled, or otherwise interacted with in such a way as to detect and identify—to *pick up*—an encounter with a discriminable pattern.

The detection and differentiation of such a broader pattern, a variant or invariant in the ambient light—the pickup of such information—is intrinsically interactive. The *active* perceiver of 1950 had to become a truly *interactive* perceiver:

> There is a loop from response to stimulus to response again (1966, p. 31).
>
> An explanation of constant perception . . . should be sought in the neural loops of an active perceptual system that includes the adjustments of the perceptual organ. Instead of supposing that the brain constructs or computes the objective information from a kaleidoscopic inflow of sensations, we may suppose that the orienting of the organs of perception is governed by the brain so that the whole system of input and output resonates to the external information, (1966, p. 5).
>
> The process of pick up is postulated to depend on the input-output loop of a perceptual system (1979, p. 250).
>
> The process is circular, not a one way transmission (1979, p. 61).

The course of the whole *interaction* can be critical. It is the course of the interaction by the visual system, for example, the scanning, both input and output and the relationships between them, that differentiates the pattern interacted with; it is not the 'final·, static, retinal image that 'completes' the interaction that picks up such a pattern, nor even the 'succession of images' or, better, the flow of retinal stimulation that accompanies the interaction. Retinal stimulation is relegated to the input side of an overall interactive visual system that engages in such interactions and discriminates such patterns. It is the *pattern* of the interaction that differentiates and, thus, identifies the pattern interacted with; it is not any piece or component of the interaction.

Gibson referred to the ambient light patterns as the *information* in the ambient array (of light). Thus, perception was a process of information pickup via information-extracting interactions. He was well aware of the encoding connotation difficulties involved in this terminological choice: "The term *information* cannot have its familiar dictionary meaning of *knowledge communicated to a receiver.* This is unfortunate, and I would use another term if I could" (1979, p. 242). Certainly, however, ambient light patterns are information in the sense of constituting knowledge or providing knowledge; the issue is whether it has to be encoded, communicated, in order to be accessed, or whether there is an alternative—not whether it is information per se. Gibson's choice, then, would seem to be fully appropriate.

Gibson was also well aware that retinal stimulation does occur, that it plays a central role in visual perception, and that it is involved in (interactive) processes. The issue is the nature of that involvement: "The inputs of the receptors have to be processed, of course, because they in themselves do not specify anything more than the anatomical units that are triggered" (1979, p. 251). Information, however, "is not something that has to be processed" (1979, p. 251). "Information is conceived as available in the ambient energy flux, not as signals in a bundle of nerve fibers" (1979, p. 263). Information is extracted by the interactions of sensory systems, not encoded and transmitted by sensory organs. The eye and its stimulations participate in information-extracting patterns of interactions; they do not encode that information.

Gibson's evolution from the active perceiver to interactive information pickup was clearly interactively flavored, not just in the generic sense but in the encoding vs. interactive sense introduced earlier. His criticisms also developed from their original sensation-based targets to encompass the general encoding approach. The homunculus criticism, for example, at times became a general statement of the necessity of an interpreter for any encoding: "Signals must be in code and therefore have to be decoded; signals are messages, and messages have to be interpreted" (1979, p. 61).

Gibson's strongest research emphasis, however, was on the ambient light information available to be picked up, not on the process of pickup. His strongest suggestion concerning the process of pickup was in his use of the metaphor of resonance: The perceiver interactively resonates with the available information (for example, 1966, p. 5; 1979, p. 246). Consistent with this suggestive metaphor, he also referred to the process of becoming able to extract information, of learning to resonate to available information, with a metaphor of "tuning."

This structure of metaphors was unfortunate in several senses. First, resonance is only one of the ways in which energy patterns can be picked up without intermediate enhancement of encoded information. Second, resonance requires periodicities in patterns to resonate to, and those are not necessarily available in information to be perceived. Third, even if such periodicities were available, it is neither at all clear what it is about the interactive loop that would resonate to them nor how it would do so. Fourth, last, and most important, that which resonates generally resonates at the same (or a directly related) frequency as that which is resonated to. The resonant frequency is a copy, a duplicate, of the original frequency. Such vestiges of picture, of image, of encoding conceptualizations are regretfully distortive of Gibson's basic interactive insight in his concept of information extraction. The pattern of an interaction need not have any particular structural correspondence whatsoever with the pattern of ambient light that it differentiates.

Nevertheless, the basic direction of the evolution of Gibson's theory seems clear. It was not complete, however, with the advent of interactive information extraction. A still further step was required in order for Gibson to avoid his own homunculus criticism. This step involved the problem of meaningfull perception.

In *The perception of the visual world* (1950), Gibson made a distinction between "the perception of the substantial or spatial world and . . . the perception of the world of useful and significant things to which we ordinarily attend" (p. 10, italics omitted). The first kind of perception he called *literal;* the second he called *schematic.* Schematic perception was presumed to be based on literal perception because literal perception "provides the fundamental repertory of impressions for all experience" (p. 10), and the two forms of perception were presumed to have importantly different properties. Meanings were presumed to be attached to, and detachable from, the spatial impressions of literal perception.

Such a separation of literal from meaningful perception requires a homunculus to receive the literal spatial impressions and to attach to them, to interpret them as having, appropriate meanings. Literal spatial impressions must be enhanced, presumably via some kind of processing, with meanings; suddenly all of Gibson's general criticisms apply to his own model. Clearly, this would not do.

Again, the germ of Gibson's solution to this problem was already present in 1950. He recognized, for example, that "squeezableness is something which seems to be located in the object, not in the hand. . . . Visual objects appear to have soaked up such qualities and to be fairly saturated with them, the use of the object and the shape of the object being almost indistinguishable" (pp. 203, 204). The critical distinction here is between the spatial nature of the object and its functional, or useful, nature. The pernicious assumption is that the perception of the functional nature is dependent on the perception of the spatial nature.

Suppose instead that the most direct focus of perception is the functional nature of that which is perceived. Suppose that what are most directly perceived are functional potentialities, potential usefulnesses. The patterns of interactions that detect ambient light patterns, after all, are not in any sense copies of those light patterns nor of the physical surfaces and edges that yield them; they, rather, are simply interaction outcomes that may indicate potentialities for further actions and interactions. They are simply functional indicators.[2]

Such an imbuing of perception with direct, functional, ecological meaning, already hinted at in his 1950s discussion of squeezability, yields Gibson's concept of *affordance.* "The affordance of anything is a specific combination of the properties of its substance and its surfaces taken with reference to an animal" (1977, p. 67, italics omitted). Affordances are those things the environment "*offers* the animal, what it *provides* or *furnishes* either for good or ill" (1979, p. 127).[3] And such affordances are intrinsic to perception:[4]

> The composition and layout of surfaces constitute what they afford ... to perceive them is to perceive what they [surfaces] afford ... it implies that the "values" and "meanings" of things in the environment can be directly perceived (1979, p. 127).
> The perceiving of an affordance is not a process of perceiving a value- free physical object to which meaning is somehow added in a way that no

> one has yet been able to agree upon; it is a process of perceiving a value-rich ecological object (1979, p. 140).

Interactive information extraction and affordances were the culminations of Gibson's major moves away from his early ecological direct encoding. Although we later argue that those moves were nontrivially incomplete, nevertheless they transcended that early encoding model by constructing an intrinsically interactive model of perception. Essentially, Gibson started with ecological direct encoding, then filled in the detection-differentiation-identification process, the process of 'transducing' the encodings, with so much interactive activity—extraction, resonance, pickup, affordance—so as to make it clear that whatever ultimate perceptual encoding, if any, occurred it was not primary nor necessary nor independent, but, rather, subsidiary to interactive extraction. Gibson's basic insight was that it is possible to derive information about an environment from interactions with that environment without encoding anything from that environment.

SOME CONCEPTUAL DISTINCTIONS

We turn now to some conceptual distinctions that will be useful in later analyses of Gibson's position. The first is a distinction between preparatory processes and constitutive processes (Shaw & Bransford, 1977). Essentially, some processes are functionally preparatory to particular other processes, and some are constitutive of particular other processes. For example, buying ingredients is preparatory to making a cake, while blending those ingredients is constitutive of making that cake. Conversely, erecting the walls on a foundation is preparatory to putting up the roof, but it is constitutive of building the house. The distinction is useful because any process or element that is preparatory to any representational phenomenon, such as perception, is subject to the homunculus criticism: it requires a subsequent interpreter. Preparatory representations and representational phenomena necessarily involve encodings because interactive representational processes are intrinsically interactive and in process, and there is, consequently, no time during which a preparatory process could occur. The preparatory-constitutive distinction, thus, picks out an aspect of the necessity of an interpreter for any encoding.

The second distinction is among three levels of analysis: the phenomenological, the functional, and the material. The phenomenological level of analysis is concerned with the organism's experiencing of its world, the functional level with those abstract processes and mechanisms that constitute the phenomenological experience, and the material level with those physical (physiological) structures and processes in which the functional processes are instantiated. In other words, certain material processes constitute functional processes, certain functional processes constitute phenomenological processes, and certain of such material- functional-phenomenological processes constitute perception. The importance of the distinction is that it is only with respect to the full (interactive) phenomenological level of experience that the homunculus-argument demand for nonpreparatory models holds validity. Clearly, one functional process can be preparatory to another functional process, and one material process to another material process. The point is that no such process involved in a preparatory relationship can itself be intrinsically and independently representational in nature, for to claim such is to subject that process (or its output) to the need for a subsequent interpreter and, thus, either to an infinite regress of interpretive homunculi (if the interpretation is in terms of further encodings that then require further interpretation) or to the destruction of its presumed independent representational (encoding) nature (if the interpretation is in terms of, thus subsidiary to, an interactive interpreter). Such processes in preparatory relationships can, of course, be collectively *constitutive* of interactive representational phenomena, just so long as the representational nature is an aspect of the entire interactive organization and is not presumed to be resident in any component process per se.

Gibson did not accept the formal position of philosophical phenomenology, but his reliance on the phenomenological level in the previous sense is frequent and clear. These range, for example, from his seeking "to establish an empirical correspondence between the stimulus and its conscious resultant" *(1950,* pp. 51-52) to his claim that a theory of perception should account for "the eventful world and the perceiver's awareness of being in the world" (1979, p. 239) and to "what an object affords us is what we normally pay attention to" (1979, p. 134). More fundamentally, Gibson relies implicitly (and sometimes explicitly) on the phenomenological level of analysis every time he claims that perception is direct in the sense of not being mediated by preparatory processes: "The term *awareness* is used to imply a direct pickup of the information, not necessarily to imply consciousness" (1979, p. 250).

AN INTERACTIVE MODEL OF PERCEPTION

At this point, we briefly outline an interactive model of perception.[5] We present the model for purposes of comparison and contrast with Gibson's, with which, as might be expected, it shares several features, and as a counterexample to some of Gibson's general criticisms of mental-processing theories. Our primary concern at this time is neither with the ultimate correctness of the model nor with its details, but, rather, that it is an example of a mental-processing model that avoids the general Gibsonian critique and that it provides an interesting perspective on Gibson's own model.

Any goal-directed system is going to have to be engaging in interactions with the environment that are dependent upon, and modifiable by, internal indications about that environment—indicators in the interactive sense discussed earlier. We call the structure of such indications *the situation image* (clearly, a nonencoded image). It must be updated and kept current, both in terms of the passage of time and of the outcomes of interactions. This updating process is called *apperception.* Within this view, *perception* is the process of interacting with the world insofar as such interacting participates in the apperceptive updating of the situation image. Some interactions will be relatively specialized for the sake of their apperceptive consequences rather than for their consequences on the world, that is, some are more perceptually specialized than others. However, the two basic aspects of serving as a ground for apperception and of potentially transforming the world are present in all interactions.

A fundamental point to be noted is that the apperceptive consequences of an interaction always exceed the immediate outcomes of the interaction. In other words, the apperceptive updating of the situation image as the basis for potential future actions always exceeds the immediate indications of the interaction outcomes upon which those apperceptions are based. Thus, the outcome of a visual scan of a glass of water most immediately indicates that appropriate optical conditions for such a scan were present at the time of the scan, but it apperceptively indicates the potentialities for multiple future interactions, ranging from additional scans to taking a drink.

It should also be noted that there are no encoded inputs coming in from the environment to be processed in this model. There are instead interactions with the environment, which interactions yield internal outcomes, and which outcomes yield internal indications concerning possible future interactions. Inputs to the system are generated by the sense organs, of course, but the significance of those inputs concerning the environment resides only in their participation in an overall (perceptual) interaction.[6]

Chapter Two

Gibson's Metatheoretical Critique

In our review of Gibson's theory, we have already presented some of Gibson's basic metatheoretical arguments as they apply to sensation-based models in particular and to encoding models in general. In this section, we examine some important elements of Gibson's more general metatheoretical critique. This examination is both in terms of the internal logic of these metatheoretical stances and as they apply to mental-processing theories. We find the stances to contain valid and telling arguments, but we find the conclusions to be invariably overstated concerning the scope of the arguments. These errors, we suggest, have unfortunately obscured Gibson's basic insights and, we argue later, may well have inhibited critical developments in Gibson's own model. The section is intended both as an analysis and as a clarification of these issues per se and as a prelude to later discussions.

MEMORY IN PERCEPTION

In sensation-based models of perception, sensations are presumed to undergo informationally enhancing processing to yield perceptions. A major component of that processing has been assumed to be some sort of comparisons of sensations or retinal images with memories or the activation of associations to memories. Stemming from this connection, Gibson's criticisms of alternative approaches to perception have included criticisms of the idea that memory is involved in perception.

These criticisms are at times clearly specific to informational-enhancement models:

> All kinds of metaphors have been suggested to describe the ways in which sensory inputs are processed to yield perceptions. It is supposed that

> sensation occurs first, perception occurs next, and knowledge occurs last, a progression from the lower to the higher mental processes. . . . All [such] theorists seem to agree that past experience is brought to bear on the sensory inputs, which means that *memories* are somehow applied to them (1979, p. 251, emphasis added).

At other times they amount to a general derision of the concept of memory:

> We assume that memories accumulate and are stored somewhere; that they are images, or pictures, or representations of the past; or that memory is actually physiological, not mental, consisting of engrams or traces; or that it actually consists of neural connections, not engrams; that memory is the basis of all learning; that memory is the basis of habit; that memories live on in the unconscious; that heredity is a form of memory; that cultural heredity is another form of memory; that any effect of the past on the present is memory, including hysteresis. If we cannot do any better than this, we should stop using the word (1979, p. 254).

There are two ways in which memory might conceivably be involved in Gibson's own theory: in the process of information pickup itself, the resonating, and in the process of *learning* to pick up particular information, the tuning. Gibson has at times acknowledged the role of some form of memory in tuning:

> A kind of memory in a new sense of the term is definitely required if we are to explain not apprehension over time but *repeated* apprehension over time. For the fact is that an observer learns with practice (1966, p. 265);

and at other times denied it:

> [The theory of information pickup] needs to explain learning, that is, the improvement of perceiving with practice and the education of attention, but not by an appeal to the catch-all of past experience or to the muddle of memory (1979, p. 254).

At all times, however, Gibson has denied a role for memory in the actual extraction of information:

> An individual who explores a strange place by locomotion [does not have to] remember a series of forms. . . . What went out of sight as he moved one way comes into view as he returns. ... He does not have to remember it . . . but only to apprehend its place (1966, p. 264).
>
> The theory of information pick up does not need memory. It does not have to have as a basic postulate the effect of past experience on present experience by way of memory (1979, p. 254).

A clarification of these issues requires first that a distinction be made between the phenomenon of memory and the various possible models of memory. The phenomenon of memory is essentially that of informational past dependence: Current process is in some way informationally dependent on past process, that is, current interactive pattern is dependent upon past interactive pattern. In this sense, memory is an aspect of (interactive) process, an aspect consisting of such informational dependence holding within (or between) the process(es). There are many models of how such dependence could possibly occur and of the various ways in which it does occur, but of particular relevance to the current discussion are those models in which memory is presumed to occur via discrete stored memories and, still more specifically, models in which such discrete memories are presumed to play an essential role in perception. In this sense, memories are discrete entities (or processes) that are specialized for their role in memory in the broader aspectual sense.

Interaction outcome indicators, in the sense discussed earlier, might seem to be discrete memories. But such indicators do not represent anything in particular at all: They are simply internal switching signals from one process to others. Indicators may participate in interactive representational phenomena, but they do not in themselves constitute self-sufficient representations. This is not the manner in which discrete memories function; thus, indicators are not memories.

Discrete memories are encoded representations of past situations or events. They are commonly thought to arise from the storage of discrete frames of perceptual experience. Any role for such encodings in a representational phenomenon, such as perception, requires that such encodings first be retrieved and interpreted, because there is no other way for their informational or representational content to be accessed.[7] Thus, any such role must be preparatory to that representational phenomenon.

Furthermore, the only reason to retrieve and interpret such a specialized memory is exactly in order to access its informational content, that is, to enhance the process that is about to occur (or to resume occurring) with the information in the memory. Certainly such enhancing interpretations of memories do occur, as in reminiscing about the past; the issue concerns their role in perception.

Thus, any involvement of discrete memories, arising from discrete perceptual experiences, requires preparatory retrieval and interpretation and is for the purpose of informational enhancement; we find the whole package of "homunculus-requiring preparatory enhancement prior to the phenomenological experience of perception" arising intrinsically from the postulation of the essential involvement of discrete memories in perception.

Correspondingly, if we examine the particulars of Gibson's arguments against the involvement of *memory* in perception, we find that they are actually directed against various aspects of this package of the involvement of discrete *memories* in perception. Gibson does not deny that such remembering and related imaging phenomena occur; he only denies their role in perception:

> This is not to deny that reminiscence, expectation, imagination, fantasy, and dreaming actually occur. It is only to deny that they have an essential role to play in perceiving (1979, p. 254).

He does criticize their role in perception, however, at multiple levels, ranging from their presumed basis in discrete perceptual experiences:

> The trouble with the classical theory of memory as applied to apprehension over time is that it begins with passive sensations in a supposedly discrete series. It presupposed that the observer gets only a series of *stimuli* (1966, p. 264).
>
> [If perceptual experiencing is discrete] it should be possible to find out when perceiving stops and remembering begins. But it has not been possible. ... A special sense impression clearly ceases when the sensory excitation ends, but a perception does not. It does not become a memory after a certain length of time. A perception, in fact, does not *have* an end. Perceiving goes on (1979, p. 253).
>
> The act of picking up information ... is a continuous act, an activity that is ceaseless and unbroken. . . . perceiving is a stream, and William James' description of the stream of consciousness . . . applies to it. Discrete percepts . . . are "as mythical as the Jack of Spades" (1979, p. 240).

to their presumed role in perception:

> Memory in the traditional sense of stored engrams is not required (1966, p. 265).
>
> [Tuning] need not be thought of as depending on a memory, an image, an engram, or a trace (1979, p. 254).

Gibson's arguments, then, hold only against the role of memories in perception in the specialized encoding sense, not against a possible role for memory in perception in the more general aspectual sense of informational past dependency. Gibson's arguments are powerful, but his conclusion is overstated and invalid as stated. In *The senses considered as perceptual systems* (1966), Gibson suggested that a revised version of memory might be appropriate in a model of perception, at least for tuning (p. 265), but he apparently and unfortunately abandoned that position.

Certainly Gibson's arguments do not apply against the interactive model outlined earlier. Memory in the informational past dependency sense is involved in the apperceptive updating of the situation image in that the updating depends on the already extant indicators as well as on the current flow of interaction. But this involvement of memory is continuous, not discrete; it is constitutive of perceptual experiencing, not preparatory to it; it is an aspect of perceiving, not a component of it. Memory in a different version of informational past dependency is also involved in the development of such a system's abilities to engage in perceptual interactions and apperceptive updates, that is, in learning. Memory in this sense also involves neither discrete memories nor preparatory processing.

More fundamentally, Gibson's arguments do not apply against memory as it is involved in his own model, for informational past dependence is certainly involved, to begin with, in the learning, or tuning, aspect of his model (as he basically acknowledged in 1966, p. 265). Moreover, it is involved in interactive information extraction. Gibson, as we have seen, argues that perception is not instantaneous, but that it is a process occurring through time. It is the temporally structured pattern of interaction that is perceptually critical, and this is so in two senses: Some information that is picked up from the ambient light is intrinsically temporally structured, and even the temporal invariants must be interacted with through time to be detected. In all cases, then, the temporally structured pattern of perceptual interaction is critical. But the occurrence of a temporally structured pattern of interaction is informationally dependent on the occurrence of past aspects of that pattern—that is, perceiving involves memory.

Gibson's arguments, then, apply against memory as it would be involved in an encoding model of perception, not in an interactive model. His overstatement of his conclusion is understandable, in that he did not have a general interactive modeling approach available as an alternative to be contrasted with the encoding approach, but it is unfortunate in that, among other things, it obscures features of his own model and likely inhibited further development of those features.

INFERENCE IN PERCEPTION

The basic issues regarding the involvement of inference in perception are closely parallel to those regarding memory in perception: the historical origins of Gibson's criticisms, the nature of the criticisms, and the nature of the required revisions. This parallelism is a manifestation of a deeper point regarding the relationships between interactive and encoding perspectives.

Inference, along with memory, was one of the components of the generation of perceptions from sensations in sensation-based models of perception:

> The visual sense must be supplemented in some way by the mind. There must exist a special mental process over and above the visual sensations: a process which in some way constructs the world out of the "raw data" presented to the mind. Such a process might be one of association and *inference* (1950, p. 13, emphasis added).

Such a role for inference was no more acceptable to Gibson than was the comparable role for memory, and a shift to modern guise does not constitute an improvement:

> Not even the current theory that the inputs of the sensory channels are subject to "cognitive processing" will do. The inputs are described in terms of information theory, but the processes are described in terms of old- fashioned mental acts: recognition, interpretation, inference, concepts, ideas, and storage and retrieval of ideas. These are still the operations of the mind upon the deliverances of the senses. ... It will not do, and the approach should be abandoned (1979, p. 238).

As with memory, Gibson's argument against the involvement of inference in perception is the homunculus argument. The homunculus argument applies to preparatory, encoding, versions of a phenomenon, and, thus, again as before, the required revision of the argument depends upon a distinction between encoding versions of inference and nonencoding, interactive, versions.

Inferences that require subsequent interpretation are inferences that yield encodings. Within an encoding perspective, these are discrete processes that generate new encodings out of old ones. Old encodings, of course, are in general memories or sensations, and new ones may well be perceptions. Inference from this perspective, therefore, is a discrete differentiated phenomenon that is functionally complementary to the discrete differentiated phenomena of encodings—memories, sensations, and so on.

The nonencoding version of inference is, as with memory, an aspectual version. Inference is an aspect of choice or selection or decision, and such selection among alternatives is an aspect of any goal-directed activity, of any interaction, when more than one possibility exists. Selection and concomitant inference may well be *implicit* in the organization of an interactive system, but they will be *present* precisely in such an implicit aspectual sense.

Inference in the aspectual sense is clearly present in the interactive model of perception outlined earlier: The apperceptive construction of indicators involves selections, thus inferences, among many possible such indicators. It is also present in exactly the same sense (a related sense) as it is in Gibson's model: An active perceiver, seeking information, must implicitly (aspectually) decide what information to seek, that is, must implicitly infer which possible new information is worth seeking. Furthermore, those implicit inferences might be wrong, in the sense of not yielding the kind of information sought, and thus in need of revision and retrial; here we have an interactive version of (nonencoded) hypothesis testing. Gibson's arguments, then, again hold against encoding versions and miss the possibility of nonencoding versions, even in his own model.

Examination of the aspectual presence of memory and inference in interactive systems yields a deeper connection between interactive and encoding perspectives. We briefly explore the connection, both for its own interest and in the interest of later reference. Essentially, encoding phenomena turn out to be limiting differentiations and specializations of representational functions that originate as aspects of interactive systems.

Such specializations begin with one of the most fundamental aspects of interactive systems: control. Control is the influence of one (sub)process on the course or pattern of another (sub)process. The first specialization of control is in terms of interactive indicators, as discussed earlier. One aspect of control is representation; it is derived from the environmental detection, transformation, and creation aspects—the implicit definitional aspects—of interactions. As control indicators acquire further differentiation and specialization with respect to this representational aspect, they develop into encodings, as signals, for example, for purposes of transmission, or as memories for purposes of storage.

Such encodings, in turn, require appropriate auxiliary and complementary differentiations and specializations, such as for interpretation, decision making, inference, and so on. That is, when representational functional aspects of interactive systems become differentiated into specialized components, they must do so in complementary packages. They must become differentiated with respect to each other in such a way that complementary functions remain available at the differentiated level as well as at the implicit aspectual level. Selection of new encodings on the basis of old encodings, for example, *requires* differentiated inferences.

Such differentiation within an interactive system occurs, both phylogenetically and ontogenetically, precisely because of the specializations of function that are thereby obtainable. Specialization allows first of all for greater flexibility of use, as with a subroutine. It also allows the evolution and development of such differentiated subsystems specifically with respect to their specialized functions, and, thus, the development of much greater power than might otherwise be obtainable. However, such differentiations and specializations of aspects of interactive systems always occur within the context of interactive systems, always in the service of interactive systems, and always in principle eliminatable in those systems.

Encoding versions of cognitive phenomena, thus, are not only impoverished limiting cases of interactive versions when encodings are taken as independent forms of representation and are not only subsidiary to interactive systems, both as argued earlier, but furthermore they are differentiated specializations of phenomena that are intrinsically interactive to begin with. Correspondingly, as we will have occasion to point out later, appropriate specialized encoding phenomena are not only not contradictory to the interactive approach, but they are predicted by it.[8]

ENHANCEMENT IN PERCEPTION

Memory and inference are simply two of the most prominent features of sensation-based models' enhancement of sensations into perceptions. Gibson's criticisms of their involvement in perception is derivative from his rejection of such enhancement in general.[9] His basic criticism of preparatory enhancement is the homunculus argument, and, as has been noted, this carries over to the special features of memory and inference.

Gibson has an additional general argument against enhancement, however. The homunculus argument derives from the necessity for encodings to have interpreters, while this second argument derives from the impossibility of encodings being given independent definitions. Consider the sensation- based role of memories in the construction of perceptions out of sensations. Memories provide information required for perceptions that is not to be found in the sensations. The memories, in turn, are derived from prior perceptions, which are similarly dependent on prior memories, of still earlier perceptions, and so on. We are forced either into an infinite regress or into the typical encoding innatism. Thus, we find:

> The error lies, it seems to me, in assuming that either innate ideas or acquired ideas must be applied to bare sensory inputs for perceiving to occur. The fallacy is to assume that because inputs convey no knowledge they can somehow be made to yield knowledge by "processing" them. Knowledge of the world must come from somewhere; the debate is over whether it comes from stored knowledge, from innate knowledge, or from reason. But all three doctrines beg the question. Knowledge of the world cannot be explained by supposing that knowledge of the world already exists. All forms of cognitive processing imply cognition so as to account for cognition (Gibson, 1979, p. 253).

As with the homunculus argument, however, this circularity argument applies only to encoding versions of enhancement. Stated in its general form, this argument charges a circularity in that knowledge of the world is required to supply and explain knowledge of the world. What is actually required, however, is knowledge of the potential world to make possible perceptual knowledge of the current actual world. In this version, the circularity is potentially broken, though of course we must explain the origin of the knowledge of the potential world. If that is assumed to be intrinsically encoded knowledge, then we encounter the regress of seeking an independent encoding definition. If knowledge of the potential world is assumed to be interactive in nature, on the other hand, then that origin is no longer problematic in principle—the origin is via learning constructions of interactive systems—and the regress is avoided.

In one sense, then, in the encoding sense, enhancement is neither required nor possible. If the information in the ambient light is not sufficient to specify the available affordances, then certainly nothing else is either. A model of informationally deficient encoded sensations enhanced by encodings of prior perceptions, or by any prior encodings, is untenable. In an interactive aspectual sense, however, enhancement *is* required. The information may be sufficient to specify the available affordances, but those specifications must be learned and explicated. Enhancement is an intrinsic aspect of such explication. Information pick up *is* an enhancement beyond the immediate patterns of interaction that specify that information. If, however, the phenomenological-level experience of perception is *constituted* by such enhancing pickup, by apperceptive updating, then neither the homunculus nor the circularity arguments apply.

DIRECT PERCEPTION

Gibson's metatheory is basically an elaboration of his rejection of sensation-based enhancement theories. This includes his rejection of such particular features as memory and inference, as well as of the general phenomenon of enhancement. Instead of the sensation-to-perception model, Gibson claims that his theory is, and that all theories of perception should be, of *direct* perception.

> When I assert that perception of the environment is direct, I mean that it is not mediated by *retinal* pictures, *neural* pictures, or *mental* pictures. Direct perception is the activity of getting information from the ambient array of light. I call this a process of information *pickup* (1979, p. 147).

Direct perception, then, is Gibson's primary term for his metatheory. It is partly positively defined by his own theory, but largely negatively defined by his criticisms of enhancement theories.

We have found, however, that Gibson's metatheoretical arguments actually hold against *encoding* models of perception, not against all possible mental-processing models—in particular, not against interactive models. This is fully consistent with the interactive insights of Gibson's own model, and delineating this consistency, in fact, strengthens Gibson's position. But Gibson did state his criticisms as applying generally, not just against encoding models. The term *direct perception* is appropriately suggestive of Gibson's rejection of mediating encodings in perception, but it is also unfortunately suggestive of his early, ecological, direct encoding model, and his overstatement of his metatheoretical criticisms is supportive of that interpretation. A perceptual model without memory, without inference, without enhancement in any sense, *is* a direct encoding model. It cannot be interactive, for, if so, it would involve memory, inference, and so on in the aspectual senses discussed earlier. Therefore, it must be an encoding model. It cannot be a mediated encoding model; such encoding mediations are the original targets of Gibson's criticisms. Therefore, it must be a direct encoding model. Thus, Gibson has evolved his way into an interactive theory and an antithetical direct encoding metatheory.

The evolution of Gibson's theory has already been traced. The evolution of his metatheory is already clear: without a general alternative to contrast with encoding approaches, Gibson's criticisms would appear to apply with full generality, and that is exactly how they have been stated. Most of Gibson's available targets, in fact, have been exactly the kind of enhancement of encoding models to which his criticisms *do* apply.

> Adherents to the traditional theories of perception have recently been making the claim that what they assume is the processing of information [is] in a modem sense of the term, not sensations, and that therefore they are not bound by the traditional theories of perception. But it seems to me that all they are doing is climbing on the latest bandwagon, the computer bandwagon, without reappraising the traditional assumption that perceiving is the processing of inputs (Gibson, 1979, p. 251).

With such a conflict between theory and metatheory, however, Gibson's position is obviously subject to difficulties of interpretation. Critics of Gibson usually do not focus on the specifics of what he has learned about perception; the importance of his theoretical contributions is generally accepted. Nor do they usually focus on the particular logics of his metatheoretical arguments, which would be difficult to defeat within a presupposed encoding perspective. Instead, they focus on the implications of Gibson's overstated metatheory, charge that the alternative he offers is impossible, and, therefore, conclude that his metatheory need not be taken seriously. What is impossible, of course, is a direct encoding theory, but with the criticisms misplaced on the directness rather than on the encoding.

The conflict between Gibson's theory and his overstated metatheory tends to obscure his basic interactive insights. More unfortunate, however, is the probable effect on the development of those insights. Gibson could not possibly begin to be specific about the processes of information pickup without it becoming obvious that past dependencies, selection, inference, and so on were at least implicitly involved. Without any caveats on his own criticisms of such phenomena in perception, this path of development was closed to Gibson. Since information pickup is taken as an impossible direct encoding by many, if not most, non-Gibsonians, examination of its nature is closed to them as well. Exploration of information pickup as a process, not just as a function, thus has yet to begin.

Direct perception, then, in the sense of not involving mediating encodings, is a valid metatheoretical position consistent with Gibson's theory. Direct perception in the sense of direct encoding is an invalid metatheoretical position inconsistent with Gibson's theory. The ease with which the former is interpreted as the latter, both within and without the Gibsonian camp, has obscured and retarded the basic task of understanding perception.

Chapter Three

Critics of Gibson

In this section we examine two recent criticisms of Gibson's position (Ullman, 1980; Fodor & Pylyshyn, 1981) to illustrate the distortion and confusion that result from taking Gibson as espousing a direct encoding. In both cases, the authors, as is almost universal, proceed from an encoding perspective themselves, interpret Gibson's position as being that of a direct encoding, and attack the presumed directness of that position. Their critical points, we suggest, are generally telling against a direct encoding position, though not against an interactive position, and they presuppose a mediated encoding as the alternative to direct encoding. They are, therefore, subject to all of the difficulties inherent in that "alternative."

It must be emphasized that to take Gibson as espousing a direct encoding model is *not* inaccurate: To do so is simply to take Gibson's own metatheoretical statements literally. It is, however, not perspicacious in that it obscures and ignores what is most novel and, we suggest, most valid in Gibson's theory. It is, unfortunately, to accept Gibson's own metatheoretical strawman in place of his genuine theoretical and metatheoretical insights.

ULLMAN

In "Against direct perception," Ullman (1980) contrasts the "direct approach" to visual perception and cognition with the notion, "which is central to contemporary cognitive science . . . that mental processes involve computations defined over internal representations" (p. 373). The heart of Ullman's analysis is his conceptualization of what it would mean for perception to be direct: "In the direct theory, perception does not involve computations of any sort; it is the result of the direct pickup of available information" (p. 373). The stimulus-percept relation would be direct, or immediate, in Ullman's sense, "if the relation had no meaningful decompositions into more elementary constituents" (p. 374, emphasis omitted), that is, "if a further elaboration of [the information-pickup] operation would only be possible in physiological, but not in psychological terms" (p. 374). Ullman illustrates the distinction that he intends between direct and indirect with the example of the addition of two integers via a direct "look up" in a table (corresponding to direct pickup) versus addition via computational manipulations of integer representations (p. 374).

Ullman's own commitment to an encoding perspective is implicit in his rendering of computations as "defined over" representations and is evident in his choice of symbolic addition to exemplify process: The process of addition is direct if it is a nondecomposable table look up of the symbolic code for the sum, and it is indirect if it is decomposable into operations on such encodings. Ullman's encoding perspective is also explicitly manifest in such phrases as "the method by which the stereo information is encoded" (p. 380).

Ullman's rendering of Gibson as a direct encoding theorist is the essence of his interpretation of *direct* as akin to a table look up that is not decomposable in psychological terms. Furthermore, we find, "A formulation in terms of invariances would be advantageous for the theory of direct perception if invariances could be discovered in the changing visual array that would be (a) informative enough to specify the structure of the moving objects, and (b) simple enough so that it would be reasonable to suggest that they are picked up directly" (p. 378); in his suggestion that things must be "simple" in order to be "picked up directly," it is clear that Ullman considers direct perception to be an impoverished version of the mediated encoding approach that he advocates.

Given Ullman's position and interpretation of Gibson, his criticisms are understandable and appropriate. He argues on logical ground, for example, that "if we consider all distinguishable perceptions (such as the perception of all different shapes) as distinct percepts, the number of possible stimuli and percepts becomes too large to be amenable to direct pairing" (p. 376), that is, direct encoding, and, therefore, suggests that perception must "employ processes or rules of formation" (p. 376). The argument, of course, has a powerful impact on an impoverished direct encoding model, but none whatsoever on an interactive model. Such comparative impoverishment is, in fact, a point in favor of interactive models over encoding models (with their combinatorial restrictions), whether mediated or not.

Ullman also presents two empirical examples as contraindications to direct encoding. The first is an illusion from which Ullman concludes that "the perception in Mach's illusion evidently depends on the internal state of the observer" (p. 379). The conclusion is not unreasonable, and we do not examine it in any detail; the point is that while the conclusion, if true, may be telling against a direct encoding model (which will not accommodate such a dependence on internal state), it is essentially a theorem of the interactive apperceptive model. Ullman's second example, involving apparent mental rotations of mental representations in a problem-solving task, is still less conclusive. The conclusion that mental representations are involved is again not unreasonable, though perhaps contestable, but, in any case, 1) it is not clear that such representations need be encodings, 2) it is not clear that they would trouble Gibson's position, whether encodings or not, because problem solving rather than perception is involved, and 3) they would be completely consistent, even as encodings, with the interactive position, which explicitly postulates the differentiation of such representations in higher order processes.

Ullman's criticisms, then, whether logical or empirical, apply with varying strengths to a direct encoding interpretation of Gibson, but not at all to an interactive model and, therefore, not to an interactive interpretation of Gibson.

Ullman not only criticizes the directness of Gibson's direct encodings, but he also suggests a criticism of the logical form of the support that Gibson offers for his position. He suggests, in particular, that

> The argument that a Gibsonian theory of direct perception is required simply because . . . sensation-based theories are considered untenable suffers the fallacy of "argument by selective refutation." That is, only one of the alternatives to "direct perception", not all of them, is refuted. Association of sensations is not the only conceivable form of a mediating perceptual process (p. 375).

We have seen earlier, however, that Gibson's criticisms have a much broader scope than just sensation-based models. They apply to any perceptual model whatsoever that involves mediating encodings. Ullman does not consider the content of Gibson's major arguments at all, not to contest their validity, not to restrict their scope, and not to show that his own position is immune to them.[10] We argue, of course, that, as a mediated encoding position, Gibson's arguments not only apply to Ullman's position, but they are validly telling against it.

There is, in fact, an important sense in which Ullman himself has inadvertently committed the fallacy of "argument by selective refutation." He has assumed that if perception is not a process of direct encoding, then it must be a process of mediated encoding. This ignores the third possibility of it being an interactive process.

Ullman ends his discussion of Gibson with a defense of, and argument for, the general "computations defined over internal representations" approach to psychological phenomena. Critical to this discussion is the claim that:

> The calculator example examined above illustrates in what sense processes and representations are amenable to empirical investigation: certain events and components within the calculator can consistently be interpreted as having their meaning in the domain of numbers and operations on numbers. There is nothing mysterious or mentalistic, then, in accepting and studying those intermediate representations and processes (p. 380).

From the interactive perspective, of course, the critical questions here are "Having meaning for whom (or what)?" and "What constitutes the having-meaning-for relationship?" Semantic relationships, such as "having meaning for," cannot be indefinitely explicated in terms of semantic relationships, that is, encodings explicated in terms of encodings, upon pain of either circularity or infinite regress. What is the ground of such explications? Encodings, whether innate or otherwise, will not do.

FODOR AND PYLYSHYN

The general logic of Fodor and Pylyshyn "How direct is visual perception? Some reflections on Gibson's ecological approach" (1981) is identical to Ullman's: The two possibilities are direct encoding and mediated encoding; therefore, if direct encoding is proven untenable, then mediated encoding must be the case. Thus, both analyses miss the possibility of the interactive approach.

Fodor and Pylyshyn arguments, however, proceed somewhat differently from Ullman's. The fundamental issue for Fodor and Pylyshyn is whether perception involves mediating inferences: Direct perception assumes that mediating inferences are not involved, while the "Establishment theory (sometimes referred to as the 'information processing' view) is that perception depends . . . upon *inferences*' (p. 139). They point out that some form of noninferential direct detection is common to both types of theories:

> Even theories that hold that the perception of many properties is inferentially mediated must assume that the detection of *some* properties is direct (in the sense of *not* inferentially mediated). Fundamentally, this is because inferences are processes in which one belief causes another. Unless some beliefs are fixed in some way other than by inference, it is hard to see how the inferential processes could get started. Inferences need premises (p. 155).

Thus, since direct detection is common to both approaches, the basic issue is whether direct detection is sufficient or whether inferences are necessary in addition.

They approach this issue via “the trivialization problem" (p. 142). Essentially, if anything whatsoever can be directly picked up, and if any process counts as a direct pickup, then the theory is trivial: It excludes nothing, and therefore tells us nothing. The problem, then, is to find appropriate constraints that render the notions nontrivial.

> Gibson's account of perception is empty *unless* the notions of 'direct pick up' and of 'invariant' are suitably constrained. For, patently, if *any* property can count as an invariant, and if any psychological process can count as the pick up of an invariant, then the identification of perception with the pick up of invariants excludes nothing. We will show, however, that Gibson has no workable way of imposing the required constraints consonant with his assumption that perception is direct. To put the same point the other way around, our argument will be that the notion[s] of 'invariant' and 'pick up' can be appropriately constrained only on the assumption that perception is inferentially mediated. This is hardly surprising: Gibson and the Establishment agree that pick up and inference exhaust the psychological processes that could produce perceptual knowledge; hence, the more pick up is constrained, the more there is left for inference to do (p. 141).

Fodor and Pylyshyn proceed to consider and demolish several conceivable constraints that might appear to save Gibson's direct encoding from triviality, to examine some proper nontrivial constraints and show how these constraints require ensuing inferences, and to suggest that a deep source of Gibson's difficulties is that he overlooked the necessity of accounting for intentionality in human cognition, and, therefore, he correspondingly overlooked the necessary role of mental representations and concomitant inferences in accounting for such intentionality. There are also a number of subordinate and side points and arguments along the way.

We do not contest Fodor and Pylyshyn's general arguments that a direct encoding interpretation of Gibson is untenable. The assumption that anything and everything perceived is directly encoded *is* trivial and empty: It yields an uninformative and unbounded ad hoc proliferation of basic encodings.[11] Our attention is much more focused, however, on two strongly interrelated issues: The failure of their critical arguments to apply to the interactive approach, and the failure of their arguments concerning what *is* the case to withstand a critical analysis of their encoding assumptions.

First, it is clear that Fodor and Pylyshyn are implicitly assuming an encoding perspective in their arguments: Their ultimate proposal is for a mediated encoding perspective in which certain properties are directly encoded—transduced is their term—and perceptions are generated from those encodings via inference.

> *All* theories of perception must acknowledge the direct pick up of *some* properties. In Establishment theories, what is directly picked up is often taken to be the properties to which *transducers* respond (p. 152). We believe that there *are* ways of constraining the notion of a perceptual mechanism—via an independent characterization of transduction—but the price you pay is that many perceptual processes turn out to be nontransductive, hence presumably inferential. This is Gibson's characteristic dilemma, and we claim he has no way out of it (pp. 152-153).
>
> The more pick up is constrained, the more there is left for inference to do (p. 141).
>
> If, as we have argued, perception is an inferential process, then what goes on in perception is the construction of certain kinds of "arguments"—*viz.* from the premises that transducers make available to conclusions which express perceptual beliefs (p. 183, footnote 15).

The argument that Fodor and Pylyshyn propose, then, is that if direct detection is appropriately constrained so as to not be trivial, enhancement via inference is necessary.

Second, it is clear that Fodor and Pylyshyn interpret Gibson as espousing an impoverished direct encoding version of an encoding model. This follows from their general treatment of Gibson as proposing a model of direct detection with no inference and of the Establishment approach as involving direct detection *plus* inference. It also follows more specifically from their equation of direct pickup with transduction:

> We will see . . . that Establishment theories do propose mechanisms for the direct pick up of certain sorts of stimulus properties. ... In Establishment theories what is directly picked up is . . . the properties to which *transducers* respond (p. 152).
>
> For us, "direct" means only "noninferential" (p. 156). . . . noninferential processes like transduction (p. 183, footnote 15).

and most clearly in:

> [In Gibson's model] the objects of direct detection (transduction) must be so specified that no perceptual judgements turn out to be inferentially mediated (p. 157).
>
> We will assume, in what follows, the identification of what is "picked up" with those properties that transducers respond to (p. 158).

We also find overt manifestations of the impoverishment view of Gibson's model: In a discussion of what kinds of properties could be directly detected and, thus, what kinds of things could be transducers, they refer to them as "being primitive in the required sense" (p. 175).

Third, as mentioned previously, Fodor and Pylyshyn regard the general encoding perspective as exhaustive; the issue is whether there is mediation. If they did not, they could not claim that the untenability of direct encoding entails mediated encoding (p. 166).

These three points are contextually important in recognizing the limited scope of Fodor and Pylyshyn's trivialization argument—in particular, in recognizing its failure to apply to interactive models and, thus, its failure to apply to an interactive interpretation of Gibson. The argument, simply, is that if *anything* can be directly encoded, then we have explained nothing. Notice, however, that if we argue "If *anything* can be *mediatedly* encoded, then we have explained nothing," then the argument does not seem to go through. The missing premises in the argument are that direct encodings are nondecomposable, while mediated encodings *are* decomposable and that explanations require decompositions. Decomposition of encoding problems is in terms of subencodings, and prior encodings, and their associated inferences—hence, the mediated encoding model. Interactive models, on the other hand, are also decomposable, though not necessarily in any way involving encodings or inferences, and, therefore, they are not subject to the trivialization argument. To assert that anything whatsoever can be interactively detected is not a trivialization of the interactive approach.

In Fodor and Pylyshyn's discussion of what constraints and consequent conclusions they view as necessary in order to avoid the trivialization problem, they propose two major considerations. The first concerns an argument that on pain of trivialization, the physical layout of the environment cannot be considered to be directly picked up—only certain properties of the ambient light can be—and, therefore, that perception of the layout must involve inferences based on the pickup of ambient light properties. The second involves an analysis of what properties of the light could possibly be directly picked up, that is, transduced, to serve as foundations for those inferences. We consider these two analyses to be the core components of Fodor and Pylyshyn's position. At the level of generality at which they are stated, in fact, they form the essential foundations of *any* version of the strongest form of the encoding approach to perception; the mediated encoding approach. We argue, respectively, (1) that the case for mediating inferences based on the pickup of properties of light is invalid when the interactive possibility is taken into account and (2) that the attempt to define transduction, of *any* properties, is internally incoherent.

We address first the argument for the necessity of mediating inferences. Fodor and Pylyshyn begin with an examination of the related concepts of "containing information about" and "specifying." Essentially, the idea is that things contain information about each other if they are correlated, and one thing specifies another for a particular organism if, assuming those things in fact contain information about each other, that organism can find out about the second given an occurrence of the first:

> The state of affairs Si contains information about the state of affairs S2 if and only if Si and S2 have correlated properties. ... As Gibson repeatedly remarks, this is an entirely "objective", nonpsychological notion of information. Information in this sense is something "in the world", just as correlation is (p. 158).
>
> But, whereas information is an ontological category, *specification* is an epistemological one. The idea is basically that when two states of affairs are correlated, the organism can use the occurrence of one to find out about the other. Under such circumstances, the first state of affairs is said to *specify* the second (for that organism). Correlation (hence information) is presumably a necessary condition for specification: When Si specifies S2,
> Si and S2 are correlated, and Si contains information about S2. Gibson's favorite example is the relation of specification that holds between features of the ambient light and features of the distal environmental layout (p. 159).

These conceptualizations, Fodor and Pylyshyn point out, give rise to a problem:

> Now, the relation of *containing information about* is symmetrical, but, in the general case, the relation of *specifying* is not. Suppose that the state of the layout is correlated in a certain way with the state of the light. While it is then true that the properties of the light contain information about the properties of the layout, it is equally true that the properties of the layout contain information about the properties of the light. However, for no organism that we know of . . . does the structure of the layout specify the light. Organisms just do not use the properties of the layout to find out how the light is arranged. Notice that that is not because the information is not there. . . . This raises a problem, though not one that Gibson discusses in these terms: *Viz.* what *determines the direction of specification!* (p. 159).

The solution to the problem, they claim, is forthcoming:

> As soon as the problem is put this way, the principle at issue seems clear enough. What determines the direction of specification is the nature of the detectors (transducers) available to the organism. Light specifies layout and not *vice versa* precisely because we have transducers for light and no transducers for layout (p. 159).

To this point, we have no basic objections to Fodor and Pylyshyn's reasoning. The point that has been reached, however, is a critical one in that, if transducers are assumed to be direct encoders, one set of conclusions follows, while if they are assumed to be interactive detectors, a quite different set of conclusions follows. We have already argued, and will see still more evidence in a moment, that Fodor and Pylyshyn interpret transducers as encoders.

The next critical step (for our current purposes) in Fodor and Pylyshyn's discussion is concerned with the question of how the specification relationship works for the organism.

> Gibson seems to agree that picking up features of the light is causally necessary for "directly perceiving" features of the layout. . . . Where Gibson speaks of directly perceiving features of the layout in consequence of picking up features of the light, the Establishment theory speaks of perceiving features of the layout in consequence of transducing features of the light. . . . There is, however, this difference: The classical theory has a story about *how you get from detected properties of the light to perceived properties of the layout* (p. 165).

It should be noted at this point that the critical implicit claim in this last sentence is not so much that Gibson does not have such a story, but that no such story is possible for Gibson. To not have such a story could be taken simply as a focus for further work; for such a story to be impossible is fatal. We agree that no nontrivial story is possible within a direct encoding interpretation of Gibson, but will show that such a story is possible, and is in large respects already contained in Gibson's writings, for an interactive interpretation.

The critical point for the moment, however, is Fodor and Pylyshyn's version of the story.

> The story is that you infer the latter [the layout] from the former [the light] on the basis of (usually implicit) knowledge of the correlations that connect them. Gibson clearly does not like this story, but it is quite unclear how he is going to get along without it. . . . The basic problem for Gibson is that picking up the fact that the light is so-and-so is *ipso facto* a *different* state of mind from perceiving that the layout is so-and-so (p. 165).

The essence of the matter is in this last sentence. The sentence is unobjectionable per se, and it certainly leads to the conclusion that Fodor and Pylyshyn want: "Some process *must* be postulated to account for the transition from one of these states of mind to the other, and it certainly looks as though the appropriate mechanism is inference" (p. 166). The difficulty is that the sentence and its consequences are non sequiturs from the discussion preceding them.

The original formulation of the issue was in terms of moving from picking up features of the light to perceiving features of the layout. Fodor and Pylyshyn have shifted from this to a formulation in terms of moving from picking up that "the light is so-and-so" to perceiving that "the layout is so-and-so." This critical non sequitur is based on the assumption that picking up features of the light is equivalent to picking up that the light is so-and-so. If this equivalence holds, Fodor and Pylyshyn's conclusion would seem to hold, and if picking up features of the light is the same as encoding those features, then the equivalence holds.

Picking up features of light, however, need not be the same as encoding those features, or as encoding anything else, for that matter. If an interaction, engaged with certain pattern(s) of light, arrives at a particular internal outcome that is correlated with those pattern(s), then it has detected those pattern(s) in an interactive sense. That detection, however, does not at all necessarily constitute or yield a representation "that the light is so-and-so." That is, the detection of a pattern is not necessarily the same thing as the structural representation of that pattern. So the critical equivalence in Fodor and Pylyshyn's argument fails.

Furthermore, it fails in a way that provides an answer to their problem of specification. If the internal interaction outcome is correlated with the pattern of light, then it will also be correlated with anything that is correlated with that pattern of light—in particular, with features of the layout. That is, if the light contains information about the layout, then so will the interaction outcomes. And that internal information will have been 'directly picked up' by the interactions, without any mediating encodings or inferences.

Of course, if the internal outcomes contain information about anything correlated with the light patterns, then they contain information about a vast number of states of affairs: intermolecular forces underlying the substance of the layout, electron energy-level transitions in the surfaces of the layout, interference interactions in the light, and so on. The problem of specification, then, is not so much how the organism gets from the light to the layout as it is how the organism gets from internal interaction outcomes to the layout *in particular,* among all the other states of affairs about which those outcomes contain information. The answer, of course, is that it learns to do so. Other interaction systems are constructed that make use of those internal outcomes in selecting which kinds and courses of interactions are likely to be successful, that is, are likely to accomplish their (internal) goals. The organism, then, learns to pick up the information in the light and to make use of the information in the resultant internal outcomes, concerning what the environment will permit, deny, and inflict on the organism with respect to its own interactions. That is, the organism learns to use its internal outcomes to specify affordances. The organism learns to pick up affordances, directly.[12]

There are no encodings and no inferences in this story. The pickup is direct, yet it is not trivial in the sense that Fodor and Pylyshyn charge: Both the learning of and the actual processing of such interactive detections, pickups, and specifications are, in principle, highly decomposable, analyzable, empirically investigatable, and explanatory.

In further support of their mediating inference story, however, Fodor and Pylyshyn point out that "the subject's epistemic relation to the structure of the light is different from his epistemic relation to the layout of the environment, and . . . the [latter] relation is causally dependent upon the [former]" (p. 165). This is true, but Fodor and Pylyshyn conclude that, since the epistemic relation to the layout is *causally* mediated by the epistemic relation to the light, the epistemic relation to the layout must therefore be *epistemically* mediated by the epistemic relation to the light: "What we detect is not the information *in* Si but rather the informative properties *of* SI. Then what we learn about S2 in consequence of having detected these informative properties depends upon which inferences we draw from their presence" (p. 166). From an interactive perspective, this is exactly wrong. The inference from a causal mediation of an epistemic relation to a corresponding epistemic mediation of that epistemic relation holds *only* if the original epistemic relation to the light is an encoding. Then, by definition, the generation of new encodings for the layout from the original encodings for the light would involve inferences. If the epistemic relation to the light is one of interactive detection, however, it is equivalently an epistemic relation to any state of affairs that is causally mediated by, hence correlated with, that light—without any necessary epistemic mediation whatsoever.[13] It is exactly the information *in* the light, the correlation, that is picked up, *not* the informative properties *of* the light.

Fodor and Pylyshyn, however, take the case for their view to be so persuasive that they wish to explain how Gibson could have missed it:

> Something has clearly gone wrong, and it is not hard to see what it is. Having introduced the (purely relational) notion of states of affairs *containing information about* one another (i.e., being correlated) Gibson then slips over into talking of *the information in* a state of affairs. And, having once allowed himself to reify information in this way (to treat it as a thing, rather than a relation), it is a short step to thinking of detecting the information in the light on the model of, for example, detecting the frequency *of* the light; *viz.* as some sort of causal interaction between the information and the states of a perceptual mechanism (the information makes the perceptual mechanisms "resonate") (pp. 166, 167).

That is, having reified information as a thing in the light, rather than as a relation between the light and the layout, Gibson can talk of picking up that information directly and can avoid the issue of how the relation to the layout is epistemically captured:

> How (by what mental processes) does the organism get from the detection of an informative property of the medium to the perception of a correlated property of the environment? How does the fact that certain properties of

> the medium are *de facto* informative manage to have any epistemic consequences? ... It is perfectly possible that an organism should pick up a *de facto* informative property of the light but not *take it to be* informative, e.g., because the organism does not know about the correlation (p. 167).

From an interactive perspective, it is Fodor and Pylyshyn who have reified information. By assuming that information can only be picked up as "things," encodings, their argument follows: If the information in the light can be picked up only by encoding the informative features of the light, then the layout must be inferred from those encodings. Certainly Fodor and Pylyshyn are correct in that you cannot detect a correlation per se (p. 166), though you might learn about one. But to conclude that "you cannot pick up the property of *being correlated with*" (p. 168) and, therefore, that you must pick up the "informative properties" and infer the layout holds *only* when "pick up" is read as "encode." From an interactive perspective, picking up the property of "being correlated with" is exactly what is going on: The light is correlated with the layout—that is the information in the light. By interactively generating internal states that differentiate among relevant states of affairs in the light, those internal states are correlated with the light and, therefore, with the layout; the information (correlation with the layout) in the light has been directly picked up so that that same information (state of being correlated with the layout) is now present inside the organism.[14] There is no encoding that the light is so-and-so, no encoding of informative properties, and no subsequent mediating inferences in this story. Only when "pick up" is read as "encode," only when "information picked up" is reified as "information encoded," does the Fodor and Pylyshyn story follow.

Fodor and Pylyshyn's argument for the necessity of mediating inferences has three major steps: 1) only features of the light can be (causally) directly detected; 2) to detect or pick up features of the light is to encode that the light is so-and-so; and 3) to get from encodings that the light is so- and-so to encodings of the layout requires mediating inferences. We have no objections to steps one and three, but have argued that step two is invalid and, therefore, that the overall argument is invalid.

We have seen how two of Fodor and Pylyshyn's major arguments hold only if an encoding approach is presupposed. Their trivialization argument holds only against a direct encoding interpretation of Gibson, not against an interactive interpretation, and their first argument concerning what is necessary in order to avoid trivialization—that the layout must be inferred from the light—holds only if it is assumed that the properties of the light are encoded, not if they are interactively detected.[15]

We turn now to their second application of the trivialization consideration—to, in fact, the foundation of Fodor and Pylyshyn's position: their model of direct encoding, or transduction. Their crucial objection to Gibson's direct encoding is its triviality and consequent emptiness. Correspondingly, their primary concern in their own position is to define transduction in such a way that it does what they need it to do (directly encode light properties), but it is nontrivial and, consequently, nonempty. We undercut this issue entirely. We claim that the issue of triviality is irrelevant in this circumstance since the conceptualization of a direct encoding, whether trivial or not, is incoherent, and, therefore, the existence of a direct encoder, a transducer, is impossible.

Fodor and Pylyshyn's definition of a transducer is focused on the problem of excluding inference from transduction: Transducers must generate the noninferential basic ground of premises for subsequent inferences.

> We say that the system S is a detector (transducer) for a property P only if (a) there is a state S_i of the system that is correlated with P (i.e., such that if P occurs, then S_i occurs); and (b) the generalization *if P then* S_i is counterfactual supporting—i.e., would hold across relevant employments of the method of differences (p. 161).

From Fodor and Pylyshyn's perspective, "It is, of course, condition (b) that does the work" (p. 162). This is because the mere fact of an *if P then* S_i correlation (condition a) does not preclude mediating inferences: Such a correlation holds with, among other things, the layout, which is ultimately to be perceived. By requiring that the *if P then* S_i correlation be such that it is supportive of counterfactuals, that is, it holds even when other conditions are changed, Fodor and Pylyshyn intend to exclude mediating inferences.

The basic reasoning here derives from the fact that mediating inferences are required when one state of affairs is inferred from the encoded premise of some other state of affairs: The point of the counterfactual-supporting condition is to exclude such epistemically mediating states of affairs and, thus, to exclude the inferences that connect the mediating- states-of-affairs premises to the ultimately encoded conclusions. The principle of this exclusion is that, when such epistemically mediating states of affairs are involved, the *if P then* S_i connection can be broken by eliminating the mediating correlate. For example, the layout cannot be perceived if the mediating light patterns are missing—if the lights are off—thus the layout-to-perception correlation is *not* counterfactual supporting. But such errors

> . . . are possible only where the perception of P is mediated by the detection of one of its correlates, the [error] occurring when the correlation fails. Since, however, transduction is, by assumption, direct—i.e., *not* dependent on specification—failure-of-correlation illusions cannot, by definition, arise in the case of transduced properties (p. 162).

Thus, the requirement that the connection be counterfactual supporting eliminates connections for which such failure-of-correlation illusions are possible, which eliminates connections mediated by a correlated state of affairs, which eliminates the inferences from that correlated state of affairs to P itself. A transduced property P is not detected via correlational specification; it is detected directly.

At this point we encounter a lacuna in Fodor and Pylyshyn's argument. Their constraint that the correlation underlying a transducer must be associated with a counterfactual-supporting *if P then* S_i generalization is intended to eliminate mediating inferences, an issue we address later. For now, however, we wish to point out that they need not only a P to S_i connection, but also an S_i to P connection: S_i as a state in the transducer is supposed to count as a detection of P. They have the simple S_i to P connection in that the P with S_i correlation is symmetric, but if the reverse generalization, *if* S_i *then P,* is not also counterfactual supporting, then the transducer must (attempt to) distinguish between those conditions in which the process potentially yielding S_i is a detection of P, and those cases in which it is not—that is, the detection (transduction) yielding S_i will involve inference, inference involved in determining whether S_i should in fact be generated.

Fodor and Pylyshyn have in fact assumed this reverse counterfactual-supporting condition without stating it. They clearly want the status of S_i as a detection of P to be counterfactual supporting, for the counterfactual-supporting *if P then* S_i assures that P is *sufficient* to S_i, thus eliminating *mediating* inferences based on mediating correlations, but it requires a counterfactual-supporting *if* S_i *then P* to assure that P is *necessary* to S_i, thus eliminating *discriminative* inferences in determining whether P really exists and, thus, whether S_i should be generated. They have assumed this reverse condition in a number of ways: in construing a transducer as a detector (pp. 160, 161, 163), in their claim that detectors (transducers) for P are "illusion-free with respect to P" (p. 162), in their claim that transducers are noninferential and thus cognitively impenetrable (that is, not subject to influence by inner premises and biases) (pp. 182-184), and so on. Henceforth, then, we assume both counterfactual-supporting constraints on transducers.[16]

Assuming such constraints, however, does not entail accepting them. Both constraints are needed (if not intended) in the definition of a transducer, and we assume them as such. We argue against them, however, both in terms of their separate possibilities and in terms of their joint adequacy. We also argue against the adequacy of even stronger conditions. If our arguments are valid, if these constraints are not possible, and not adequate, and if they cannot be strengthened in a way sufficient to be adequate, then the definition of a transducer fails to yield the properties that

Fodor and Pylyshyn need, and no definition can yield those properties. What a transducer is supposed to do cannot be done.

First, consider the possibility of a counterfactual-supporting *if P then S_i*, that is, that P is sufficient to S_i. P sufficient to S_i, in turn, implies that any time the state of affairs P occurs, S_i will occur. But such a certain relationship will obtain only if it is impossible to break the micromediating causal chain that connects P to S_i, for to break that chain would yield a case of P without S_i, thus violating the condition. The breaking of such a causal chain, however, is always a possibility. In the case of the eye, such a chain might be broken, for example, by a distorting or translucent material in front of the eye, or by an opaque lens, or by an opaque vitreous humor, or by an absence of some necessary chemical in the receptors, or by an interference with some necessary reaction, and so on. All that is required is that some connection in the causal sequence be temporarily disrupted.

There are several apparent counterarguments to this point. The first is the claim that such disruptions simply render the system nonfunctional—it is not a transducer so long as the connection is broken. This has an intuitive appeal, especially when considering permanent disruptions, but when considering temporary and reversible disruptions it breaks down. The whole point of the counterfactual-supporting condition is to render the detection of P a nonmediated detection by virtue of making the connection from P to S_i sufficient and certain, thus obviating the need for any mediating considerations of any mediating states of affairs. If those mediating states of affairs can potentially be blocked, however, then they potentially must be taken into account, and the detection of P cannot be considered to be 'direct.'

Another way to see this is to note that if such disruptions can be ruled out as invalid, if the transducer only exists or only functions under 'normal' conditions, then the *if P then S_i* generalization is no longer counterfactual supporting, but is only valid across employments of the method of differences (the altering of normal correlational relationships) *that do not alter the normal ecological conditions and relationships.* But, of course, under normal ecological conditions, the *layout* is correlated with internal states, and the layout-to-internal state relationship remains unchanged so long as such normal conditions remain unchanged. Thus, the layout (or anything else, for that matter) can be transduced, and we have returned to the trivialization problem. Fodor and Pylyshyn are aware of this sort of problem, and they point out the invalidity of such appeals to ecological normality (p. 161), but it is not clear how their own definition is to escape the same problem.

Considering the second counterfactual-supporting condition, *if S_i then P*, that is, that P is necessary to S_i, we find similar problems. For P *sufficient* to S_i, we have the potentiality of *breaking* the causal chain between P and

S_i; for P *necessary* to S_i, we have the possibility of *initiating* the causal chain somewhere between P and S_i, thus generating S_i without P: a P-illusion. We could employ physical stimulation of the retina, for example, or radiation, or chemical stimulation, just so long as the process yielding S_i is initiated at some point between P and S_i. As before, attempts to rule such possibilities out of bounds by appeal to definition or to ecological validity simply eviscerate the definition and allow the layout and anything else to be transduced.[17]

In effect, the constraints of necessity and sufficiency between P and S_i are attempts to put constraints on the web or network of potential causal processes in which they are both embedded, in particular, to require a single necessary and sufficient connection between them. Fodor and Pylyshyn believe that the relationship between P and S_i will have the required properties if that relationship is based on a physically lawful generalization (pp. 163, 164), but, so long as the chain of connections has any temporal extent, there remains the possibility of disruption and illusion within that chain, unless it is assumed to be isolated and protected within a closed, ecologically normal' system. Thus, the kind of relationship, the kind of lawfulness, that Fodor and Pylyshyn need between P and S_i can be obtained in only two ways: by requiring an assumption of ecological validity, which makes their proposal no different from Gibson's direct encoding, or by having P and S_i be simultaneous, thus allowing no possibility of disruption or interference. Such simultaneous lawful relationships certainly exist in physics, for example, in Maxwell's equations, but recourse to this point would leave only the most elementary subatomic quantum interactions as possible grounds for transduction and, thus, only the corresponding quantum states of affairs as properties that could possibly be transduced. But this would have the layout being inferred from transduced premises concerning the quantum states of affairs (inside the rods and cones?), not a likely proposition and certainly not the one that Fodor and Pylyshyn propose.

The necessity and sufficiency conditions between P and S_i are intended to eliminate inference in the detection of P, so as to provide a ground of premises for subsequent inferences, by eliminating specification and differentiation relationships that would require inferential computation. The same arguments that apply to Gibson with regard to this quest also apply to Fodor and Pylyshyn: The only way to arrive at such necessity and sufficiency is to postulate the transduction (direct encoding) of only the most elementary and unlikely states of affairs. There is a deeper argument against this concept of transduction, however, that does not depend on any difficulties in finding such necessary and sufficient relationships. To present this argument, therefore, we assume that a transducer requires a counter- factual-supporting relationship of *P if and only if* S_i and that reasonable such relationships can be found.

Essentially, our argument is that a transducer will still require inference. The necessity and sufficiency requirements are intended to eliminate the need for inferences of differentiation and mediation, respectively, and, for current purposes, we grant that. But it is still the case that S_i must be taken as a representation of, or a premise about the existence of, P in order for it to ground the ensuing inferences—S_i must be or contain an encoding of P. The point that still requires inference concerns that encoding: How is it determined that P is what is to be encoded? No matter how necessary, sufficient, or certain, the *correlation* between P and Si cannot *itself* be detected;[18] the epistemic link from S_i to P requires a differentiating inference, not differentiating among the many *causally possible* antecedents to S_i, for by current assumption there is only one, but differentiating among, specifying the particular one among, the unbounded number of *epistemically possible* antecedents. That is, it must still be inferred that it is P that has been detected. The 'learning' required for this inference might be pushed off into phylogeny, but the inference must take place nevertheless. Thus, the postulation of a transducer not involving inference yields the necessity of an inferential step in the transduction: a contradiction.

A possible countermove in response to this point might appear to be to take explicit account of the necessity of such an initial inference and simply to modify the definition of a transducer accordingly, either so that that one inference is the last step in the transduction or so that it is the first step after the transduction. In other words, if such an inference to P is required, then make it. But neither move will suffice, for, no matter where it is located, that step is impossible.

The first point to be noted is that making that inference on the basis of the internal state S_i rather than on the basis of P itself offers no particular improvement: The underlying epistemic problem is essentially equivalent between the two cases. The equivalence is in terms of detecting P *as P* and detecting S_i *as a necessary and sufficient correlate of P:* In both cases, the inference to P must be made, and the necessary and sufficient causal relationship between P and S_i does not change the epistemic issue at all. The point here is that P is its own best necessary and sufficient correlate, and, if the inference to P cannot occur in terms of P itself, then the interposition of a distinct necessary and sufficient correlate, such as S_i, will not improve matters.

The issue focuses, then, on whether the inference to P can be made at all, not on whether it is to be made from P or from S_i. Such an inference to P must be based on some identifying properties of P or, perhaps, on the basis of some single, essential, identifying property (such as some property that underlies P being a necessary and sufficient correlate of S_i—note that the correlation itself cannot be detected), and this raises the question of how these properties are to be detected. They might, of course, be detected interactively (as might P itself), but then we would no longer be considering logically independent encodings—the encoding perspective would have been abandoned.

The inference to P, then, must be based on subsidiary encodings of the identifying properties of P (and, perhaps, subsidiary inferences as well). These subsidiary encodings, of course, must themselves be generated without inferences—they must be transduced—so as to provide a ground for the inference(s) to P. But P was itself to have been transduced in the first place, so as to itself provide the ground upon which perceptual inferences can be based. Thus, the basic problem of mediated encoding models of perception, the problem that transduction is supposed to solve, has simply been iterated, and we begin an infinite regress. (We also encounter the regress of having to postulate prior encodings of P and of the relationship between P and its identifiers, in order to account for the inferences to P, and then of trying to account for those prior encodings.) Fodor and Pylyshyn's move, of basing transduction on necessary and sufficient causal relationships, epistemically buys nothing.

As Fodor and Pylyshyn point out (though in a different context), "The existence of a correlation [or an identity] between two states of affairs does not, in and of itself, explain how the detection [or generation] of one of them could eventuate in . . . knowledge of the other" (p. 167). The problem encountered in attempting such an explanation within an encoding perspective is that we are attempting to generate something that is meaningful, an encoding, out of something that is not, a physical event or state of affairs, but encodings can be defined only in terms of *already existing* meanings. The regress that is generated by that attempt is, interestingly, a converse of Gibson's homunculus problem: The homunculus problem focuses on the necessity of an interpreter for *already existing* encodings; the regress above stems from the necessity for an interpreter, a provider of meanings, to generate any *initial* encodings.

There is still a further difficulty, however, even if the necessity and sufficiency difficulties and the initial interpretive regress were all discounted. If an encoding of P is to be generated, then something, perhaps S_i, is to be taken as such an encoding of P. But a major difficulty ensues: What is it for S_i to be an encoding of P? It is for there to be an epistemic correspondence between S_i and P, for there to be knowledge of what S_i encodes. But how can there be knowledge of the P that S_i is to be taken as encoding? Only in terms of S_i, the encoding for P. Thus we have: S_i encodes whatever S_i encodes. We have encountered the incoherence problem.

One last possible move might seem to be to drop the claim that S_i is a premise or involves an encoding, and simply posit that initial internal states such as S_i yield and are transformed into other internal states by various internal processes, without any encoding claims for those states. But, on the one hand, if those states are presumed to yield encodings of the layout or of the environment, if they are presumed to yield encoded perceptions at any point during those ensuing processes, then we simply encounter all the transduction, regress, and incoherence problems at that point. On the other hand, if those states are not presumed to be or to yield encodings, but are presumed simply to indicate to other processes what courses of processing might be most successful, and if the great power for such purposes is recognized to allow the S_i to be set by interactive processes, not just by passive one-directional processes, then we have an interactive model.

We claim, then, that the concept of a transducer is incoherent because neither the necessity nor the sufficiency constraints can be meaningfully met and because neither the initial interpretive regress nor the encoding incoherence can be avoided. If transduction is incoherent, then so also is any model of mediating inferences grounded on transduced premises. We conclude that not only is the direct encoding approach to perception untenable, but so also is the mediated encoding approach.[19]

CONCLUSIONS

In both cases examined, Ullman (1980) and Fodor and Pylyshyn (1981), the arguments are directed against a direct encoding interpretation of Gibson from a mediated encoding perspective on perception. In both cases, the general arguments are valid against a direct encoding interpretation, though not against an interactive interpretation. Furthermore, in both cases, the mediated encoding alternatives proposed are fatally vulnerable to criticisms directed at their encoding foundations. In the Fodor and Pylyshyn article, that mediated encoding alternative position is particularly well elaborated; its vulnerabilities are thus exposed in corresponding detail. In spite of Gibson's own metatheory, then, a direct encoding interpretation of Gibson is untenable. So also is the mediated encoding alternative to Gibson. This leaves the interactive interpretation of Gibson and the interactive approach in general.

Chapter Four

Perception and Beyond

DIRECT ENCODING MODELS OF PERCEPTION

Gibson came closest to making an explicit commitment to a direct encoding model in his 1950 book. The direct retinal-stimulus-to-perception relationship of the model is tantamount to a direct encoding. Intrinsic tensions in that model, however, stimulated and guided the evolution of an increasingly interactive model of perception, including a shift in the causal locus of perception to the ambient light patterns, the development of an interactive conceptualization of the perceiver, and the realization that functional affordances were the primary objects of perception, rather than meaningless physical arrangements.

Gibson's metatheory, on the other hand, retained and elaborated a strong implicit commitment to direct encodings. Gibson's basic stance had originated as a negative reaction against the inadequacies of mediated encoding models, and, while his theory went on to develop a positive interactive content, his metatheory retained a basically negative critical character. A metatheory consisting largely of criticisms of mediated encoding models continued to back Gibson into a direct encoding comer because he had no explicit metatheoretical sense of any other direction to move. If perception were not a mediated encoding process, then it apparently had to be a direct encoding process, because he had no explicit sense of any other alternative. Gibson, thus, had no explicit interactive metatheory to match his interactive *theory* and to give positive guidance, not just negative critical guidance, to the development of that theory. Accordingly, he continued to overstate his metatheoretical criticisms in such a way as, if taken literally, to implicitly commit him to a direct encoding model—the overstatements left no other possibility—and he correspondingly constrained and inhibited the understanding of the interactive processes posited by his theory—his overstated critical metatheory prohibited some of the most fruitful directions of exploration.

The intrinsic tensions and strains in Gibson's position were still operative, however. To posit, even implicitly, a direct encoding model of perception is to make perception preparatory to cognition and action in the same sense in which sensations are preparatory to perceptions in mediated encoding models and is, thus, to make Gibson's model subject to a version of the same homunculus argument that he applies so tellingly against those mediated encoding models.[20] It seems likely that it is this tension that yielded Gibson's realization that his general position had powerful implications for cognition and action, not just for perception (Gibson, 1979). Such an extension of his basic interactive insights beyond perception has deep possibilities, and it is tragic that he was not to pursue it further.

A direct encoding model of perception is untenable. It is, first of all, not meaningfully decomposable (Ullman, 1980) and is thus explanatorily trivial (Fodor and Pylyshyn, 1981). These criticisms, however, are stated from a mediated encoding perspective and have a correspondingly limited impact. Gibson's homunculus argument, in contrast, applies both to mediated encoding models and to direct encoding models. So also do the more general interactive-perspective arguments of logically independent encodings being unnecessary and incoherent. A direct encoding model, then, is more deeply untenable as an *encoding* model than it is as a *direct* encoding model.

For all these reasons, a direct encoding interpretation of Gibson's theory, however appropriate and understandable it may be given his metatheory, is an unfortunate and unperspicacious interpretation. In focusing on the untenability of the direct encoding commitments of his metatheory, the validity of those metatheoretical criticisms against *mediated* encoding models is overlooked, and the nature and value of the interactively interpretable insights in his theory is obscured.

MEDIATED ENCODING MODELS OF PERCEPTION

The deficiencies of mediated encoding models of perception formed the point of origin for Gibson's theory and have formed the carapace of his metatheory ever since. From Gibson's perspective, those deficiencies included the assumption of a passive perceiver, the infinite regress involved in trying to explain encodings (perceptions) in terms of enhancing prior encodings (memories based on prior perceptions), and the infinite regress of homunculi needed to interpret the presumed resultant encodings.

Internal to an encoding perspective, however, the issues appear differently. From an encoding perspective, all representation, including perception, *must* be some kind of encoding, for representation *is* encoding. From such a perspective, Gibson's criticisms seem somewhat to miss the point: They may raise interesting issues, but to charge any particular, current, mediated encoding model with not having a solution to the problem of the ultimate origin of knowledge (the regress of encodings) or of the nature of awareness (the regress of homunculi) seems rather premature. Of course, Gibson's criticisms are with regard to deficiencies in principle, not just with regard to current fact, but they must be wrong nevertheless because they yield such untenable conclusions: They (appear to) leave a direct encoding model as the only alternative, but that is so obviously both logically and empirically untenable. The fruitful task, then, is to proceed with determining the specifics of the mediated encoding perceptual process and to leave such issues as the origins of knowledge or the nature of awareness to the philosophers, or at least until much later in the development of the science.

The basic components of that task are to determine the nature of the basic encodings and encoders, of the transducers, and of the subsequent, presumably inferential, enhancements that yield perceptions. Transducers cannot be too powerful, for to assume unlimited transducing power is to render the concept trivial and explanatorily empty. In addition, it is simply not the case that everything that needs to be perceived permits the lawlike certainty of relationship between stimulus and encoding that is necessary for a noninferential transducer to function—that is, not all encoding relationships are capable of being physically lawful. Perception, then, requires inferences based on sufficiently simple, materially lawful, direct encodings (transductions). Gibson's direct encoding model, obviously, is (or seems to be) inadequate to these considerations.

Such a stance is probably the most powerful position possible, given the constraints of an encoding perspective. Given the premises of that view, the reasoning regarding transducers and inferences does hold (as far as it goes), and Gibson's (or anyone else's) arguments that the reasoning can not go far enough, that the approach is fatally flawed in principle, simply must be wrong because they yield such impossible consequences.

To dismiss Gibson's arguments this easily, however, is simply to ignore that they *do* apply to mediated encoding models (especially the two infinite- regress arguments) and, correspondingly, to ignore that there is a serious problem of some sort involved. More subtly, it is to ignore the deeper fact that Gibson's most powerful arguments, the regress arguments, apply not only to mediated encoding models, but to direct encoding models as well. They apply, in fact, to any assumption of encodings as logically independent forms of representation. Thus, despite the fact that Gibson's overstatements regarding memory, inference, and enhancement have implicitly committed him to a direct encoding position, his own deepest arguments in those regards apply to *both* forms of encoding models.

This later fact leaves only a few possible conclusions regarding the place of Gibson's criticisms within his general position. First, it might be concluded that there is a contradiction in Gibson's model that has not generally been appreciated. This might seem a fairly obvious conclusion, once pointed out, but note that, among other possible implications, it might indicate that Gibson is not as logically committed to a direct encoding model as has usually been assumed. More deeply, however, note that this conclusion, although easy to accept from the interactive perspective (we have argued it for many pages now), is not so easy to draw for the mediated encoding theorist. The purported contradiction is between a direct encoding model and criticisms that apply to direct encoding models. The mediated encoding theorist might be basically unsurprised and pleased by additional arguments against direct encoding models, except that these particular arguments, if countenanced, apply just as strongly to mediated encoding models. The mediated encoding theorist, then, needs both the direct encoding position and the criticisms against them to be wrong.

This leads to the second possible conclusion: that Gibson's criticisms are wrong in a way or ways not generally considered. Gibson's metatheory (and theory) is commonly considered to be in error by virtue of yielding (or constituting) an untenable direct encoding position, but this purported reductio argument cannot apply in any simple way to the infinite regress criticisms because they both *apply* to a direct encoding position. Again, it is relatively easy to make sense of this position from an interactive perspective: The regress arguments are valid and do apply to all logically independent encoding models, but Gibson has overstated his conclusions from those arguments (regarding memory, inference, and so on—for understandable but regrettable historical reasons), resulting in the maintenance of an implicit commitment to direct encodings. This move of accepting the arguments but narrowing the conclusions, however, is not available to the mediated encoding theorist: The regress arguments cannot be accepted as valid within the mediated encoding view because they apply to *all* encoding models. The mediated encoding theorist, then, must find some alternative way to invalidate the regress arguments.

This leads to the third possible conclusion, an attempt at a deeper reductio of Gibson's criticisms: Gibson denies all forms of representation (after all, his criticisms apply to all logically independent forms of encodings, and what is there besides encodings?) and is thus committed to the ultra-untenable position of a strict extensionalism. This, of course, is contradicted by major components of Gibson's theory—the role of learning, acknowledgment of reminiscence, expectation, imagination, fantasy, and dreaming, and so on—but this might simply be dismissed as a basic contradiction in Gibson's theory. More fundamentally, however, the extensionalist reductio of Gibson's arguments depends on the premise that encodings are the only essential forms of representations.

This, of course, leads to the fourth (and last) possible conclusion: There is some nonencoding perspective regarding representation from which Gibson's theoretical content and metatheoretical criticisms might both begin to make sense. We, obviously, are arguing for this last possibility, the one least often considered, with that nonencoding perspective being the interactive perspective.

From the interactive perspective, the basic issues are with the general assumption of encodings as logically independent forms of representation, whether they be directly perceptual, transduced, or whatever their presumed origin. From this perspective, Gibson's regress arguments are special cases of more general arguments regarding strict encodings being impoverished, insufficient, unnecessary, and incoherent. From this perspective, the logical difficulties with Gibson's position are neither so much with his theory nor with his metatheoretical criticisms per se, but with his overstated conclusions from those criticisms. It is those overstatements that provide the primary tensions and contradictions in his position.

From an interactive perspective, there is a still deeper reason for the inadequacies of a strict encoding perspective, a reason that underlies and explains the impoverishment, insufficiency, lack of necessity, and incoherence arguments. Strict encoding views, whether direct or mediated, assume that encoded representations emerge directly from a physiological (material) level of analysis. They look to material level "lawfulness" or "resonance" for their critical explications. From an interactive perspective, however, there is at least one level of emergence *between* the material and the representational: the level of interactive control structures.[21] Representation, then, is an emergent functional property of certain forms of goal- directed interactive control structures, which, in turn, are emergent properties of certain patterns of material processes.[22] Encodings are specializations of this emergent representational function.

To define general representational (psychological, intentional) phenomena in terms of logically prior encoded representations is to reverse this sequence of emergences and, thus, to eliminate the grounds for explicating representation and encoding. This, however, is exactly what is done, for example, in present day information-processing or computational approaches, in which encodings are presumed to be somehow physiologically explicable and all other psychological phenomena are presumed to be explicable in terms of encodings and operations on symbolic (encoded) representations. In such approaches, the level of interactive control structures and processes that is properly *between* the material level and the representational level has instead been moved *above* the level of encoded representations, leaving that level of encodings hanging in midair with no grounds for explication. Encodings can only be explicated in terms of interactive interpreters, yet such intentional interactive systems have been presumed to be definable only in terms of (operations on) encodings. Strict encodings, then, are circularly incoherent: They presuppose themselves.

The incoherence of strict encodings, then, derives from their presuming to be the grounds for explicating that which must in fact be the grounds for their own explication. The impoverishment, insufficiency, and lack of necessity of encodings all derive from encodings being specializations of a broader representational function: the impoverishment and insufficiency because that specialization does not capture all of the broader interactive representational power, and the lack of necessity because such specialization, however efficient it might be, is itself not necessary.

The claim that a level of interactive systems exists in this manner between the material and the representational encoding levels actually has four parts: 1) that such a nonmaterial, not necessarily representational, level exists, 2) that representation can be explicated in terms of this level, 3) that representation must be explicated in terms of this level, and 4) that encodings can and must be explicated in terms of such interactive representations. That such a level exists is easily established: It is the domain of such areas as automata theory, control structure theory, cybernetics, and so on. All of these areas study processes, potentially interactive processes, in terms of their abstract pattern of organization—that is, without essential reference to the particulars of their material level instantiations—and without presupposing or (necessarily) instantiating any representational phenomena.

That the general phenomena of representation can be explicated in terms of the level of interactive systems, the second claim, is, of course, a programmatic claim. As such, it cannot be proven in detail except through the ultimate completion of the program. It can, however, be disproven by the failure in principle of some particular case of such an explication; correspondingly, a successful explication constitutes a corroboration of the program. For example, the sense in which the outcome state, or final state, of a process can be said to differentiate, or to 'recognize,' the string of inputs leading to that state has been well studied in passive automata theory (Eilenberg, 1974; Hopcroft & Ullman, 1969). This is an important explication and corroboration.

A recognizer automaton in an interactive mode, rather than a passive mode, is much more powerful in its ability to model processes of detection, transformation, and creation (Bickhard, 1980b). A special and important version of this is the representation-by-differentiation of a simple, goal-directed, interactive feedback system. The feedback in such a system need not be an encoding of the situation, nor an encoding of error (though such encodings might be indirectly generated). The feedback need only sufficiently differentiate the environment so that ecologically useful internal differentiations of next courses of action can be made. Such differentiations do not necessarily require the full information of (presumed) encodings: It can occur with internal outcome states serving as indicators.

An interactive recognizer is, in fact, equivalent in power to a Turing machine (Bickhard, 1980b): The claim that representational phenomena can be interactively explicated, thus, constitutes a special case of Turing's (or Church's) thesis that any effective procedure can be carried out by some Turing machine. Turing's thesis has been corroborated (it cannot be proven) by some 50 years of mathematical research (Brainerd & Land- weber, 1974; Davis, 1958; Minsky, 1967; Rogers, 1967).

The third claim, that representational phenomena *must* be explicated in terms of interactive systems, is equivalent to a claim that knowledge and representation have an essentially interactive or, roughly, pragmatic character. As such, it makes contact with the tradition of pragmatism (and related approaches) (Bernstein, 1971; Scheffler, 1974; Thayer, 1973) and with its critics. It is also equivalent to the combination of the claims that representation can be so explicated (the previous point) and that it cannot be explicated in any differing terms. Since the general encoding approach is the only evident alternative to the interactive approach, arguments against strict encodings constitute arguments in favor of the necessary interactive character of knowledge and representation.

Arguments against the sufficiency of encodings, for example, are arguments that there are forms of knowledge and representation that cannot be explicated in terms of encodings. Such arguments may concern limitations in principle or may exhibit particular forms of knowledge with arguments that those forms cannot be captured via encodings. Concerning limitations in principle, for example, it can be argued that there is no atomic level of representation at which basic encodings elements can be defined, that there is no possible origin of basic encoding elements and rules, that encodings cannot account for new kinds of knowledge, and that encodings require interactive interpreters (Bickhard, 1980b, 1982).

Concerning specific limitations, it can be argued, for example, that skills and values can be understood interactively, but not in terms of encodings. The representation by indicator differentiation of a simple feedback system also provides an example. An internal outcome indicator can be used to differentiate and select ensuing courses of interaction, but it contains essentially *no* information about what in the environment yielded that outcome; it contains or constitutes no information about what it might be taken to encode and, thus, no information that would allow it to *be* an encoding. We, as observers of the system, might claim that it encodes such- and-such a physical state of affairs that yielded that outcome, but the system itself has no necessary knowledge of that sort, and it certainly need not infer its next course of interaction on the basis of such encodings: Does a thermostat infer its actions on the basis of encodings of temperature, or does it simply discriminatively respond to internal indicators? We confuse our perspective as observers with the perspective of the system being observed when we take such indicators as encodings. In effect, encodings presuppose and require too much knowledge to be able to explicate such simple differentiating indicators.

In response to such arguments against sufficiency, however, it might be concluded that interactive and encoding approaches capture fundamentally different kinds of knowledge and that neither one is explicable in terms of the other. Piaget, for example, adopts that position (for example, Piaget, 1969, pp. 356-360, 1970, p. 14). The arguments against the necessity and the coherence of logically independent encodings, however (note Piaget, 1970, p. 15!), preclude that possibility. Therefore, we have strong arguments and indications that knowledge and representation can be explicated interactively—the already established instances and Turing's theses—and that it cannot be explicated in terms of encodings—the arguments that strict encodings are insufficient, unnecessary, and incoherent. The implication is that knowledge and representation *must* be explicated interactively.

The fourth claim, that such encodings as *do* exist can and must be explicated in terms of interactive representations, follows readily in principle from the equivalent broader claims for representation in general. The specifics of such explication are important, however, not only in terms of the details of specific models, but also more broadly in terms of the constraints that such an explication imposes on the nature, power, and rationale for encodings. In particular, such encodings are going to be derivative and subsidiary encodings only, not logically independent encodings, and the properties are not going to be identical to those assumed in a strict encoding perspective. The nature of those constraints is largely unexplored.

Gibson's metatheoretical criticisms against mediated encoding models, then, can be seen as deeply consistent with an interactive approach. They can, in fact, be seen as special cases of general interactive criticisms of strict encodings of any sort. To interpret his criticisms this way does some violence to the conclusions he draws concerning memory, inference, and so on, which, if taken literally, appear to commit Gibson to a direct encoding model. But Gibson's overly broad conclusions concerning these phenomena derive from relatively subtle and historically understandable errors. The violence to Gibson's position, therefore, is relatively minor. To interpret Gibson's position as a direct encoding position, however, not only contradicts the content and historical thrust of his theory, it also is contradicted by his own metatheoretical premises that yield the supposed direct encoding conclusions.

Gibson, then, exemplifies in his theoretical and metatheoretical development a progressive deepening and elaboration of interactive insights. This development was propelled and guided by an understanding of some fundamental deficiencies of mediated encoding models that, in fact, are deficiencies of *any* strict encoding approach. This latter fact is what seems to have guided Gibson into progressively more consistent and adequate interactive models. The deficiencies that Gibson focused on are special cases of deep flaws in the encoding approach that arise from a reversal of the sequence of emergence of representations and encodings from material-level phenomena. Gibson's positions, then, together with the auxiliary arguments presented, participate in, are supported by, and support, in turn, an interactive pragmatic understanding of knowledge, representation, and meaning.

INTERACTIVE MODELS OF PERCEPTION

The overstated conclusions of Gibson's metatheory indirectly commit him to an untenable direct encoding approach. The *premises* for those conclusions, however, apply *both* to direct encodings approaches and to mediated encoding approaches, thus indirectly committing him to an interactive approach. This more fundamental commitment to an interactive approach, however indirect, is consistent with the basic interactive character of his theory. That indirect commitment to the interactive approach, in fact, seems likely to have been a primary historical inducer of the interactive evolution of his theory.

The fundamentally interactive character of Gibson's theory is manifested in two ways: 1) the explicit interactive content, and 2) the implicit interactive content constituted by the ability of the general interactive approach to make sense of otherwise obscure parts of the theory. The explicit interactive content includes the intrinsically active and interactive character of the perceiver; the recognition of ambient light patterns, rather than retinal images, as the *causally* proximate locus of perception; the recognition of functional affordances, rather than meaningless surfaces and edges, as the *epistemically* proximate locus of perception; and the recognition that the theory has implications beyond perception.

The implicit interactive content of Gibson's theory is evidenced most strongly by the ability of the interactive approach to make sense of Gibson's concept of direct pickup. To directly pick up (the information concerning) the functional affordances of the layout is an inexplicable and impossible task if that pickup is interpreted as an encoding, but not if it is interpreted interactively. It is clear that within a material level model of a perceiver, the pickup of information concerning properties of the layout must be causally mediated by the patterns of ambient light; the question is how that causal mediation is epistemically realized. The *fact* that (properties of) the light patterns are correlated with the layout *cannot* be directly picked up, but the *state* of being correlated with the layout *can* be directly picked up. To interactively detect the informative properties of the light, those properties that are correlated with the layout, is not to encode them but to generate internal states that differentiate those properties, that are thus correlated with them, and that thus contain the same information (concerning the layout) that they contain. Furthermore, the informative properties of the light cannot be directly encoded—cannot be transduced—because, among other reasons, to do so requires knowledge of those properties that is independent of the transduced encodings of them—otherwise, the system cannot know what it is that it is encoding and, therefore, cannot encode them, and there is no possible source of that independent knowledge within a strict encoding perspective. Still further, nothing in this argument depends on the particulars of what is being transduced, only on the conception of transduction, so the failure is one of principle, not of detail. Information concerning the functional properties of the layout, then, cannot be directly encoded, nor can it be epistemically inferred on the basis of transductions of causal mediators; but it can be directly picked up in the internal states of an interactive system. The interactive approach, thus, fills in much of the content of Gibson's "direct pickup," a concept that is otherwise rather obscure. The appropriateness of the interactive interpretation of Gibson's position is correspondingly enhanced—enhanced by such an implicit presence of interactive properties in a major part of Gibson's theory that is not explicitly elaborated.

There is, however, at least one important sense in which an interactive interpretation of Gibson's theory imposes a change on that theory, not just on the metatheory (though the connection with the metatheory is strong). That change has to do with the possibility of state dependency in perception. The interactive process of apperception is a continuously updating transformation of structures of indicators, not simply a registering of interaction outcomes. As such, there is an intrinsic dependency on the prior indicator structure as well as on the interactive transformation. Gibson's theory does, in fact, involve state dependencies—such as in his emphasis on the importance of temporally structured perception—but his overly broad metatheoretical conclusions concerning enhancement in perception would seem to rule out such a possibility: State dependency sounds too much like enhancement via encoded memories, when, in fact, the latter is a special (unacceptable) case of the former. Gibson's theory, correspondingly, does not acknowledge a role for state-dependent perception.

In terms of its relationship to Gibson's model, an interactive interpretation seems highly appropriate. It systematizes much of his theoretical content, clarifies obscurities, corrects errors—both theoretical and metatheoretical—and makes sense out of his primary metatheoretical premises.

It is certainly possible, however, to challenge the interactive perspective per se, not just as an interpretation of Gibson. However much it may make sense out of Gibson's insights and provide content for such processes as pickup, and however much the strict encoding alternatives are impoverished and flawed, it might still be contended that the interactive approach is itself in some way(s) inadequate or untenable. There are many possible such challenges, and not all can be considered here,[23] but two illustrative objections are addressed.

The first of these has an empirical flavor and remains within the realm of perception, while the second is an example of a broader attempted objection in principle. The 'empirical' objection is to claim that, whatever the interactive approach may claim, perceptual encodings *do* occur as a matter of empirical fact, and the whole package of encoding critiques and interactive alternatives simply founders on that fact. Support for such "facts" can be derived by adducing the marvelously detailed and engineered "encodings" of physical sensory dimensions and characteristics into corresponding frequencies or axons of neural activity. Such frequency and line (axon) encoding of sensory properties is a generally unquestioned presupposition of neurophysiological explorations of perception (for example, Carlson, 1981; Geldard, 1972). Light pattern and color, for example, are line encoded, while intensity is frequency encoded. Sound frequency, on the other hand, seems to be encoded by some combination of line and frequency encoding. In any case, so the reasoning goes, line encoding, frequency encoding, or some combination thereof is the only physiologically possible way for the organism to get information about its environment; the interactive critique, therefore, must be wrong.

The error in this position is the conclusion that the facts support the existence of *encodings.* What is factually, materially demonstrable is a complex and detailed pattern of physically realized correspondences between lines and frequencies of neural activity, on the one hand, and various proximal physical stimulus properties, on the other. *Note that within a material-level perspective, some such correspondences between the internal system and the external environment both constitute and are necessary to the differentiations of the interactive model.* Correspondence is the converse, singular aspect of differentiation, and differentiation *is* the functional resultant of sets of correspondences. Such passive one-way correspondences and differentiations are less powerful than the more fully developed interactive correspondences and differentiations in which they might participate, but, without some such differentiations, no information about the environment could be derived at all.

Sensory correspondences, then, are essential for *both* the encoding and the interactive approaches. Evidence for such correspondences, therefore, does not differentiate between the two positions. What is the difference then? The difference is that the interactive approach asks *only* for such correspondences and differentiations, while the encoding approach requires something else in addition. The encoding approach requires that the system in some way also have knowledge of what the correspondences are with. A correspondence relation per se does not constitute an encoding relation: The internal end of the correspondence must somehow in addition represent, "encode," the external end of the correspondence. It is precisely this step, of course, against which the incoherence argument is directed.

In the interactive perspective, in contrast, the material correspondences exist in fact, but the system does not know (except perhaps via higher order models of itself) what the correspondences are with, nor even that such correspondences exist, and the resultant differentiations do occur, but the system does not know what the differentiations are among. What the system *does* know is that it is useful to differentiate its own activities in accordance with the internal manifestations of such correspondences and to so regulate its own activities in accordance with potentially quite complex such differentiation strategies. The neurophysiological sensory correspondences, thus, are 'simple,' passive, one-way versions of interactive implicit definition, and the resulting differentiations are similarly passive versions of interactive differentiation.

The broader, in-principle objection to the interactive approach goes outside the boundaries of perception. It may seem plausible that interactions can pick up perceptual information about the immediate environment, but, as a general alternative to strict encodings, the interactive perspectives lays claim much more comprehensively than that—to cognition in general, for example—and, if the interactive approach cannot plausibly make good on that broader claim, then it cannot in general supplant the encoding approach. In that case, either the arguments against strict encodings are invalid, and, therefore, strict encodings can exist, in which case their role in perception might be reestablished, or there is some third alternative to both encodings and interactive representations that must be conceived and considered.

One challenge to that broader claim of the interactive approach is to object that, however much interactive apperception might be able to explicate immediate perception, there is no way that the interactive approach can explicate cognitions of higher order abstractions: Light patterns as well as the surfaces and edges that they indicate can be interacted with, but there are no corresponding realms of interaction for the abstractions of, for example, logic and mathematics. The answer to this objection is that it is false: There are such realms. In particular, the properties of the primary interactions with the world, both perceptual and otherwise, as well as of the systems that engage in those interactions, are all more abstract than the proximate environmental objects of those interactions.[24] If a higher level interactive system can interactively know about (represent) those properties of the first-level system, in the same sense in which that first-level system can know about the environment, then that second-level system can address properties more abstract than those accessible to the first. That second-level system, in turn, will have properties that are accessible to a third level, and so on. The interactive approach, then, is not embarrassed by the problem of abstraction—it has an unboundedly rich approach to it.[25]

There are many other challenges, and corresponding responses, that are not pursued here. The general outcome, however, is that the interactive perspective seems not only perspicacious as a way of viewing and understanding Gibson, but it also seems viable and richly promising as a general program.

INTERACTIVELY DERIVATIVE ENCODINGS

The arguments presented against encodings have been against logically independent encodings, not against subsidiary encodings. The implications of those arguments are that approaches to representation that presuppose encodings to be the essence of representation, or at least to be an irreducible independent form of representation, are untenable. Those arguments do not yield the conclusion that every aspect and component of the ordinary intuition of an encoding is wrong, but rather that the conceptual program inherent in those intuitions can only be carried out in a subsidiary and dependent manner—dependent upon some other (form of) representation. It is not immediately apparent, however, what the function and functioning of such logically dependent encodings would be like; it is not immediately apparent how much, and with what modifications, the conceptual program in the encoding intuition can be carried out within a representationally dependent framework. At this point, we examine a few of the major features of such a dependent encoding program.

In the classical view, encodings may be defined in terms of other encodings, which may be defined in terms of still other encodings, and so on, until at some point a basic level of encodings is reached that is defined "directly," in terms of what the encodings represent. The encodings defined in terms of other encodings are defined in terms of a representational correspondence between the defined encoding and (some structure of) the defining encoding(s). There is nothing objectionable in this process in general, though the specifics of how it is carried .out may well be modified by the introduction of representational dependencies. It is the last (or first) step, the "direct" definition step, the foundational step, that is impossible: There is simply no way to specify directly what a logically independent encoding represents because, by definition, there is no other representation available with which to so specify. Encodings must always be defined as representing "the same as" some other representation—encodings are always stand-ins—and there must be some other form of representation upon which such a structure of stand-in relationships can be founded.

The only apparent such alternative form of representation, an alternative for which there is reason to believe that it does capture the essence of representation, is interactive representation. The task, then, is to examine what encodings are like or could possibly be like, within an interactive representational framework.

We briefly pursue two issues concerning interactive dependent encodings: some ramifications of the implicit and differentiating character of interactive representation and some ramifications of the relationships and differences between encodings and interactive indicators. These issues are interesting and important, but they (and their discussions) are far from exhaustive. Primarily, they illustrate the fact that interactive dependent encodings are not the same as classically conceived encodings.

The implicit and differentiating character of interactive representation visits itself on subsidiary encodings. This has immediate consequences for the presumed referential nature of encodings, but those consequences can perhaps be best illustrated from the perspective of an approach to representation that is essentially the equivalent of the encodings approach: representation as truth conditions. There is also an interesting and revealing exception to the implicitness of interactive representations: There is a special sense in which direct encodings *are* possible. The discussion begins with a consideration of derivative encodings.

First, note that the classical stand-in relationship is (perhaps complexly) transitive and, second, that the so-called direct definition of the classical model is actually a hidden stand-in relationship: "'X' (directly) represents Y" simply means " 'X' represents the same thing as (stands in for) 'Y'." Thus, within the classical view, even the most multistepped indirectness of encoding definition could in principle be eliminated in favor of some direct (referential) definition of what was being represented. Within the interactive perspective, however, that is not possible. The fundamental interactive representations are not themselves defined in terms of what they represent, and so they provide no ground for eliminating the encoding stand-in definitions in favor of such (complexes of) 'direct' definitions in any derivative encodings. The implicit sense in which interactive representations represent thus visits itself, via the transitivity of the stand-in relationship, on all subsidiary encodings.

Now consider the truth conditional approach to representation. Classical direct encodings are assumed to somehow make direct reference to, or to directly represent, that which is encoded. In that, they can be taken as specifying truth conditions: That which is referred to either exists or not, or obtains or not.[26] Derivative encodings, correspondingly, specify more complex truth conditions, which, again, either obtain or not, and much can apparently be made of representation in general in terms of such conditions (for example, Fodor, 1975).

This treatment of representation in terms of environmental truth conditions, however, is not only constructible on the basis of encodings, it is essentially equivalent to direct encodings—a direct encoding is simply an elemental or atomic truth condition representation—and it is, therefore, subject to essentially the same arguments. How is the truth condition specified, except circularly and incoherently in terms of the encoding itself? How is it interpreted; how is it known whether or not it obtains, except via interactive detection, which makes it no longer necessary or logically independent, and, therefore, certainly not the essence of representation? What are the basic encoding elements (truth conditions)? Are they adequate to all representation without ad hoc proliferation? How could they possibly have originated? How can they account for fundamentally new knowledge? And so on. Truth conditional and direct encoding approaches are interdefinable; they are subject to the same arguments; and they fall together.

In apparent rejoinder, it might be pointed out that an interaction-outcome indicator represents that some predicate P obtains in the environment, a predicate corresponding to the sufficient conditions for that particular outcome, and that such a representation is a truth conditional representation. The premises of this rejoinder are true, but the conclusion does not follow. The reason is that nothing is known or represented by that outcome indicator about what those sufficient environmental conditions are, and, therefore, no such predicate P is specified or represented, however much one may want to agree that such a predicate must exist (and hold true in this case). In the classical approach, not only are truth conditions "represented," but also it is known *which* truth conditions are so represented. That is not, and cannot be, the case in the interactive perspective. Without knowing which truth conditions are directly represented, there is no way for the classical approach to get off the ground.

It might seem that the relevant predicate P could be specified in terms of the internal procedure (and its outcome) that is being indicated and that that information, being internal to the system, is available. This is true, and it begins to capture the valid aspects of the truth conditional approach, but it must be noted that specifying P in this way represents nothing directly about the environment—it is an internal truth condition, not an environmental truth condition. Concerning the environment, it is apparent that P represents implicitly, but it is not known what; it is known in general that P_1 does not equal P_2 and does not equal P_3, and so on—that is, it cannot be directly known what that outcome represents, but it can be differentiated from others. Basing a "truth conditional" approach on such implicitly and differentiatingly representational "predicates," instead of on explicit and referential predicates (encodings), is what the interactive approach requires.[27] The implications of that requirement are largely unexplored.[28]

There is an interesting exception to this implicitness of representation, however. Interaction-outcome indicators implicitly represent encounters with their sufficient conditions in the environment. This implicitness with respect to the environment remains in any derivative encodings. The relationship between an encoding and the more fundamental encodings or indicators in terms of which it is defined, however, is not one of implicitness, but is, rather, direct. That is, if X is an encoding defined in terms of Y, if X is a stand-in for Y, then X does not encode whatever Y represents, for Y does not encode it either, but X *does* directly represent, directly encode, Y itself (whether Y is an encoding or an indicator). Thus, there is a special sense in which direct encodings can exist.

That special sense is possible precisely because both ends of the encoding relationship involved, that is, both X and Y, are directly accessible to the representational system involved—they are both part of that system. X cannot be defined directly as encoding Y if Y is in the environment because the system has no independent, and by definition cannot have any independent, way of specifying what Y is. But if Y is part of the system that is defining X, then that specification is in principle trivial: The stand-in relationship is explicitly constructed, and whenever X is to be representationally interpreted, it is interpreted via that stand-in relationship. As before, in order to define a direct encoding, what is to be encoded must be already known. With respect to the environment, that requirement begs the epistemological question at issue. With respect to inside the system, however, that epistemological question is void.

The second major issue pursued here concerns the relationships and the differences between encodings and indicators. It turns out that the difference is constituted by the functional manner in which they are interpreted, and that this difference has rather far-reaching consequences for the functional utility and developmental origins of encodings. In particular, they will tend to be used, and to develop, in conjunction with generalized strategy procedures in attempts to cope with novel and unfamiliar situations. The core discussion is with respect to *general* encodings and procedures, and it particularizes fully to the case of perception.

Consider first the basic question: What is the difference between derivative encodings and derivative indicators? The apperceptive process generates new indicators on the basis of old indicators as it updates and explicates the implications of prior indicators. Similarly, encodings are generated on the basis of prior encodings or, ultimately, of prior interactive indicators. Is there a difference in these two cases, and, if so, what is it?

There is a difference, but it is constituted not so much by how the derivative indicators and encodings are generated as by how they are interpreted. Indicators, including derivative indicators, are interpreted directly by being "consulted" by procedures in the course of their interactions; encodings, all of which are derivative, are interpreted by first "interpreting"—explicating, filling in, tracing back—the relevant stand-in relationships and *then* by interpreting the resulting indicators. Representation via stand-in relationships is the quiddity of encodings.

A critically important consequence of this difference is that if an encoding is used frequently, there will be selection pressures to increase efficiency by skipping, or short-circuiting, the stand-in relationship interpretation and learning to directly interpret the "encoding" instead. That is, there will be a tendency for encodings to develop into indicators. It would even be quite plausible that some procedures would still be treating particular elements as encodings, while other procedures would have learned to treat those identical elements as indicators. Among other things, this point illustrates that the essences of encodings, of indicators, and of the differences between them are strictly functional in kind.

This tendency of encodings to evolve into indicators has important implications concerning the functional usefulness of encodings: If the situation has recurred with some frequency, then the relevant representational elements will tend to be indicators; conversely, if the elements are encodings, then the situation must be novel or infrequent. Encodings stand in for indicators when the system does not (yet) know how to use the grounding indicators (or the resultant encodings) directly.

By their nature, then, encodings will tend to be used in novel and infrequent situations. Used for what? Used by what? It may be clear that encodings will with time and use tend to become indicators, but how and why would a system ever develop the procedures for constructing and using encodings in the first place? The question is in large part a developmental one and so, correspondingly, is the answer. The discussion considers some constraints on the development of procedures, which must ultimately create and make use of any encodings, before returning to encodings per se.

One of the primary functions of procedures is to compute interactive strategies toward assigned goals in terms of the available subgoals and actions provided by subordinate procedures.[29] From the perspective of computational efficiency, such strategy computations should be integrated and specialized as tightly as possible into the general interactive procedures in which they occur. From the perspective of learning and developmental procedures, however, the relevant considerations differ strongly.

The pressures of interactive computational efficiency select for integration and specialization; the pressures of developmental efficiency, on the other hand, select for generality and flexibility. Strategy computations involved in two different procedures may be formally identical, yet computational efficiency considerations considered alone would yield two tightly integrated and specialized *copies* of that single formal strategy, one in each procedure, rather than a single copy accessible to both procedures. But to independently discover and construct that same strategy all over again in the second procedure, once it had been developed for the first procedure, would be immensely inefficient from the perspective of the heuristics for development and learning and would become still more so if that particular strategy should prove useful in still other applications. Developmental heuristics are under strong selection pressure to develop generalized strategy procedures, and aspects and components and structures and constructors of such procedures, that can be flexibly tried out as single whole elements during the construction of new procedures—during the attempts to solve novel problems—rather than being rediscovered and reconstructed from scratch each time. If frequency of use warrants, then separate in-line copies of those strategy procedures can be integrated and specialized into particular procedures as a later step.

An important realm of development, then, is these generalized strategy procedures[30] together with the heuristics for constructing new ones, the differentiations into generalized aspects and components that makes such new construction of generalized procedures most efficient,[31] and the means for making use of and trying out such general procedures during problem-solving attempts. In fact, since much of higher level cognition is involved in such novel problem solving, these developments will yield much of higher level cognition.

Making use of and trying out general strategy procedures impose a number of important requirements of their own. First of all, the relevant indicator information upon which and with respect to which the procedure is to function may not be in appropriate form and may need to be summarized or transformed in some way. Second, the tryout may not work, so the information in the original indicators may need to be retained and protected from the operations of the general procedure by making a "copy" of the original indicators for that procedure to work with, leaving the original indicators alone. Third, the general procedure computation may take significant time to complete, and the information in the original indicators may need to be protected *for* that general procedure from the ongoing operations of interactive apperceptions, again by making some sort of copy of the original indicators. For such reasons, an important aspect of the development of generalized strategy procedures is the development of generalized copiers and transformers of indicator information for the use of such procedures and the development of corresponding interpreters of the products of computations of such procedures in terms of the relevant original or related indicators. That is, general strategy procedures generally operate on indicator-derived *encodings,* and an important aspect of the development of general strategy procedures is the supportive development of general encoding and decoding procedures.[32]

Note that a critical aspect of strategy computation is selection among alternatives, or decision making, and that a critical aspect of decision making is inference. As these aspects become differentiated and specialized as autonomous generalized procedures, then, together with the accompanying encoders and decoders, we have the familiar picture of inferences operating on encodings to yield new encodings—but in the unfamiliar form of both inferences and encodings being subordinate within an interactive knowing system.

Encodings, then, are used in novel situations in order to make use of the potentially immense power of a potentially vast array of strategy, decision-making, and inference procedures. In so doing, they tend to step out of ongoing interactive and apperceptive processes via the encoding step and then, if relevant, to step back in via the decoding step. With use, both the encodings and the general procedures tend to disappear (from those computations) in favor of directly interpreted indicators and specialized procedures whose computations participate directly in the interactive and apperceptive flow.

These considerations and conclusions apply directly to the case of perception; this is to be expected because, from the interactive perspective, perception is a specialized aspect of all interaction. One fairly common example of stepping out of the normal apperceptive flow for the sake of explicit inferences is provided by the cases in which some part of a visual scene is unclear or ambiguous—perhaps what might be an object is largely obscured by leaves—so we stop to try to figure out what it might be, and then look again. We might even alternate several times between the processes of inferential consideration and visual examination before finally figuring out what it is that we are looking at. Once we do figure out that we are looking at an X, then it simply looks like a leaf-obscured X, and it is no longer unclear or ambiguous. Clearly, the explicit inference processes can influence the apperceptive processes; also clearly, the inferring is not constitutive of the seeing—the explicit inference process is distinct from the apperceptive process.

The progression from encodings and explicit inferences to indicators and direct apperceptions is illustrated any time an individual learns to "see" and understand the environment in a fundamentally new way or via new means. Perhaps the clearest examples are provided by the progressive learning to use and see via new technological means of apperceptive interactions, such as radar, sonar, X-ray, CAT (Computerized Axial Tomography) scan, and so on. Similar processes of learning to perceive via new means occur with somewhat more temporally extended and strategically structured apperceptive skills (with perhaps more frequent novelty and consequent explicit problem solving, even for the skilled individual) such as analytic chemistry, medical diagnosis, mechanical or electronic troubleshooting, et cetera. Learning to apperceive in a new way, whether or not via new means, is a part of virtually any skill acquisition: The 'reading' of investment situations, understanding the data in quantum physics, learning to develop an understanding for a client in psychotherapy, et cetera.

Perceptual development, of course, is precisely a process of learning to perceive in new ways and via new means. Whatever the innate grounds may (or may not) be, the perceptual-development process itself can be expected to make use of and to contribute to the parallel development of general strategy, inference, and coding procedures, and to involve the progression-with-use from encodings to indicators. In early development, of course, there is not likely to be much available in the way of general strategies and associated procedures; this raises interesting questions about possible differences between early perceptual development and later perceptual development as well as about the influence of perception and perceptual development on general strategy development, not just the other way around. Perceptual development, then, and consequent perceptual process, are likely to involve an interaction between the development of the apperceptive processes per se and that of the relevant general procedures; as usual, questions concerning psychological nature are not really separable from questions concerning psychological origin.[33]

COGNITION, LANGUAGE, AND OTHER PHENOMENA

In the classical view, elemental direct encodings generate sensations (or retinal images, or 'information,' or some other equivalent) from which are generated encoded perceptions, which yield encoded cognitions, which can be transmitted via encoded sentences. Such an approach would like to have even actions encoded as "responses," though it is by now largely recognized that actions must somehow be computed relative to goals. If the interactive view is correct, however, then the foundational encodings in this sequence are impossible and unnecessary, and the sequence itself therefore collapses.[34] The interactive perspective, then, has important consequences far beyond perception; it affects all of psychology. At this point, we would like to illustrate some of those consequences.

The necessity of a nonencoding conception of representation in general as well as the more specific necessity of a nonencoding conception of perception impact most directly and obviously on cognition. Our primary discussion, correspondingly, focuses on cognition, particularly on the now- dominant approach of functionalism. The nonencoding impact is no less powerful, though slightly less direct, on language; this is somewhat more briefly mentioned. There are also distinct implications for other areas of psychology, such as for memory, learning, motivation, and emotions, that are only minimally indicated here. Just as representation permeates all of psychology, so also do the consequences of the encoding versus interactive issue; we illustrate only a very limited, though we hope illustrative, set of those consequences.

Before turning to functionalism per se, it should be noted that the critical sequence for current purposes in the classical progression from perception to cognition to language is a sequence of logical dependencies, not a sequence of processing stages. The representational content of the encodings at any particular level is logically dependent on the content of preceding levels; this is what causes the whole structure to collapse when the presumed foundation of transduction fails (in addition to the direct arguments against encodings at each level). Such a logical sequence may *motivate* a model of corresponding sequential stages of processing of inputs, but it is logically independent of such issues of processing organization (so long as sufficient epistemic grounds somehow exist at each point in the process). The implication of this point is that the recent moves away from strictly sequential information-processing models in cognitive psychology (for example, Glass, Holyoak, & Santa, 1979), which might at first appear to be addressing the problem of encodings in cognition, are in fact moves orthogonal to this basic issue: They constitute moves away from an extremely simple-minded version of the encoding approach, but the continued basic assumption of logically independent encodings is clear (for example, Glass et al., 1979, pp. 3-24).

In its broadest sense, functionalism is the tenet that the essences of mental phenomena are to be found in the abstract patterns of processes that manifest those phenomena, independent of the particular natures of the physical realizations of those processes. This is in strong contrast to certain versions of physicalism, which hold that mental phenomena are in some sense intrinsically tied to particular physical realizations. The thesis that mental phenomena are characteristics of process patterns, independent of physical realizations, is clearly compatible with interactivism: Nothing about interactivism makes any intrinsic reference to any particular physical process. In fact, as is argued later, interactivism is the only pure functionalism.

In the common versions of functionalism, however, mental phenomena are analyzed not just in terms of process patterns, but in terms of organizations of processes *operating on representations.* The paradigmatic case, though certainly not the only version, is that of computer programs operating on symbolic (encoded) data. This basic split between computation (process) and representation permeates artificial intelligence, information-processing approaches to cognition, classical approaches to memory, to learning, to motivation, and to emotions, and various philosophical versions of functionalism (Block, 1980a). It is also intrinsic to Piaget's distinction between operative and figurative knowledge (Bickhard, 1982; Boden, 1979) and to standard approaches to language.[35]

The presupposition of such an intrinsic distinction in kind between process and representation precludes any possibility of explicating representation in terms of process—the interactive move—and leaves only the possibility of explicating representation in terms of what is represented. But representations defined in terms of what they represent *are encodings.* In the reverse direction, it is clear that the presupposition of encodings as logically independent forms of representation requires an intrinsic differentiation between process and encoded representations. The process-representation distinction of standard functionalism, thus, is equivalent to the encoding approach, and the arguments against that approach apply directly.

Standard functionalism, with its process-representation split, is a kind of half step from pure physicalism toward pure interactive functionalism. There is an important consequence of this half step, which is that the half step is insufficient to escape physicalism and its problems: Standard functionalism, in spite of its nonphysicalistic insights, is still a version of physicalism. The argument is in three parts: 1) standard functionalism is a version of the encoding approach; 2) the encoding approach *must physically* specify the representational content of its basic encodings, its basic inputs and outputs; and 3) the necessity of the physical specification of basic representational inputs and outputs is equivalent to the necessity of the physical specification of the mental processes and states as well, which *is* physicalism.

That standard functionalism is a version of the encoding approach has already been established. That the encoding approach must physically specify its basic inputs and outputs is a consequence of the fact that without such specification of the physical-environmental end of the encoding relationship, there is no way to specify what the encodings are supposed to encode. Physical specification, such as with presumed transduction relationships, provides representations of both ends of the encoding relationship, thus allowing that relationship to be defined from an external perspective and providing the appearance that it can be defined or interpreted from within the perspective of the relevant system.

That the necessity of physically specifying inputs and outputs is equivalent to physicalism follows from the fact that the distinction between inputs and outputs, on the one hand, and system states and processes, on the other, is epistemically arbitrary and vacuous, so that a *necessary* physical specification of the first, necessary no matter where the distinction is drawn, is equivalent to a necessary physical specification of the second because that arbitrarily placed distinction can always be drawn inside the system itself. That is, it is always possible to collapse (or expand) the arbitrary boundary between system and inputs-outputs so that the inputs and outputs *are* the states and processes of the system. Then, if inputs and outputs must necessarily be physically specified, so must they be in this case in which the specification of one *is* the specification of the other.

Fundamentally, the distinction between inputs-outputs and system as normally drawn · is a physical distinction, a product of the physical analysis of epistemic systems. It is not an epistemic distinction at all: There is no direct way for an epistemic system to make or to make use of this distinction for itself (though it might do so indirectly within a model of itself). From an epistemic perspective, inputs are any states or changes in the system that are not determined internally by the organization of the system, and outputs are any internally determined states or changes in the system that have any functional (controlling) relationships with inputs.[36] That is, if there is to be any epistemic distinction between input-outputs and system at all, it is a strictly *functional* distinction made entirely *within* the system. This is clearly not the same as the usual physical distinction, cannot support the encoding interpretation normally built on that distinction, and shows very simply how standard functionalism, with its physicalistic definitions of inputs and outputs, fails to be a true functionalism.

This argument does not lead to the conclusion that the standard physical distinction is useless or invalid. It is quite useful in its own sphere of analysis, and the only invalidity is in the attempt to smuggle invalid epistemic content (encodings) into the distinction. In undertaking an external-perspective physical analysis of an epistemic system, one of the things we will want to know is the physical and causal realizations of the interaction paths and how those physical realizations serve and manifest the necessary epistemic functions. Among those functions will be included the pickup of interactively useful differentiations and the accomplishments of transformations. When we locate physicocausal points in the interaction paths that give us particularly strong explanatory perspectives on how these epistemic functions are served, such as receptors in the retina or inner ear, or motor unit effectors, we make useful physical input-output distinctions at those points.[37] The difficulties arise from confusions about what kinds of analyses are taking place when we do this.

The preceding argument—that standard functionalism does not succeed in escaping physicalism—is a derivative of the earlier argument concerning the sequence of emergence of mental phenomena: If representation emerges as a phenomenon of goal-directed interactive systems and such systems emerge as kinds of physical systems, then the reversal of the representation-from-system emergence, as when encoded representations are taken as independent, leaves encoded representations hanging with no ground for explication. Simply left hanging, they become purely formal 'symbols' with no representational content; if an attempt is made to give them content by grounding that content in their physical realizations (the standard course), then, as just outlined, the commitment to physicalism is inadvertantly reintroduced.

Block (1980b, which motivated much of the preceding discussion of functionalism) recognizes both of these dangers, but without any suggestions about how to avoid them. He exemplifies the fact that "inputs- outputs" and "system" can always be made physically identical by introducing a science fiction example in which brain EEGs (electroencephalograms) *constitute* the inputs and outputs, and he argues that the consequent physicalism is intrinsically "chauvinistic"—it always underselects within the class of possible minds, always fails to acknowledge legitimate possible minds, with a chauvinistic bias toward the particular physical models taken as paradigmatic (that is, humans).[38] He exemplifies the emptiness of systems operating on 'empty symbols' with an example of the Bolivian economy being constrained to function in accordance with all the "correct" state transition rules in order to make it sentient.[39] The absurdity of the Bolivian example illustrates the basic problem with empty-symbol explication of mental phenomena: It is inevitably guilty of the error of "liberalism"—overselecting with respect to the class of possible minds so as to include absurd and impossible instances. Block, however, presents no way out of the dilemma of either physicalism or of empty, formal symbols.

The dilemma between physicalism and empty symbols is Block's (1980b) observation, but the concern with the problem of empty, formal symbols per se has been widespread. Searle (1981), for instance, in a series of examples with impact similar to Block's Bolivian example, argues that a strictly formalistic approach can never account for intentionality, that simply following a set of formal rules can never be sufficient. Again, however, no alternative is offered.

Field (1981) argues that mental representations must have representational content, that is, that they cannot be strictly formal, in order to account for phenomenological experience within a framework of materialism. He further argues that this representational content must be understood in terms of truth conditions. In order to avoid a regress problem, he argues that mental representations must *have* truth conditions, but that we need not necessarily *know* those truth conditions.[40] It is not at all clear, however, how representations can *have* truth conditions *without* our knowing those truth conditions, nor how representations can represent for us if we do *not* know their truth conditions. Interactive implicit definition and differentiation provide answers to these questions, of course, but they are not answers that are available within Field's Tarskian encoding approach.

Fodor (1981) reaches similar conclusions about the consequences of formal symbols, but an opposite, sanguine judgment about the import of those consequences. He argues that a psychology based on strictly formal symbols is all that is to be expected, because the task of providing a semantics, a representational content, for these symbols is practically impossible. That he does not seem particularly disturbed by this conclusion appears to be a result of his considering the impossibility of a semantics to be one of methodological practicality, not of impossibility in principle.[41]

Dreyfus (1981), to the contrary, argues that the empty representations of Artificial Intelligence research in principle cannot account for human knowledge and suggests, in a position informed by Heidegger's analysis of the human condition, that the critical aspect missing is the intrinsic position of knowledge within performance, of human know-how. It is a fatal error to try to split knowledge from this context. In this case, an alternative is suggested, and it has a strongly interactive flavor, but Dreyfus provides no suggestions about how it might be approached in terms of precise models.[42]

From the interactive perspective, the dilemma of standard functionalism in trying to provide a semantics for its representations is real and insurmountable. In not recognizing the emergence of representation from competent interaction, in positing an intrinsic split between the two, standard functionalism has created a deep unsolvable problem. Attempts to provide a semantics with a physical (transducer) specification of inputs and outputs are logically impossible and reintroduce the physicalism that was to be escaped in the first place, and attempts to make do without an intrinsic semantics, with purely formal symbols, leave most of the basic issues of mental phenomena—representation, meaning, intentionality, et cetera—unexplained and unexplainable.[43] The deepest functionalist insights can only be maintained in an interactivism.

It might be wondered, however, whether those insights can be maintained even within a interactivism. It is clear that the interactive approach is not subject to the problems of the physical specification of inputs and outputs—inputs and outputs are defined strictly functionally—but it might appear that it is vulnerable to the problem of empty, formal representations and that such counterexamples as that of the Bolivian economy would apply. From the interactive perspective, however, representation is in terms of and by virtue of interaction, not encoding, and the "inputs and outputs" in the Bolivian counterexample do not interact with anything. The example, therefore, is not relevant to the interactive perspective.[44]

In the interactive perspective, an input represents whatever interactive potentialities it differentiates and indicates for the system. This may be causally dependent on particular prior conditions, but it is not epistemically so dependent. The input, thus, may be empty as a symbol, as it must be if the arguments against encodings are valid, but it is *not* empty as an interactive representational indicator—it participates in the course of the interactive flow. “Formal" (functional) inputs and outputs, therefore, are not problematic for interactivism.

Standard functionalism is dominant in contemporary cognitive psychology and with good reason: The nonphysicalist insights that give rise to functionalism are valid and important. The encoding approach to representation, however, stifles those insights and blocks the attempted escape from physicalism. Interactivism, thus, is the only true functionalism in which all psychological phenomena can be approached functionally.

What would an interactive cognitive psychology look like? Little can be said at this time: Logically independent encodings must be eschewed; representation must be based on interactive differentiation and implicit definition. An initial process is to explore what current issues and results look like within this altered perspective. Some preliminary, related questions are explored in Neisser (1976), and an initial, restricted model of a few phenomena is presented in Bickhard (1980a).

An interactive approach to psychology has as great, if not greater, an impact on language studies than it does on cognitive psychology. Less of that impact is explored here, however, because 1) the connection with the perceptual issues explored in this book is less direct; 2) the current approaches to language studies are more diverse than to cognition; and 3) more has been presented about an interactive approach to language elsewhere (Bickhard, 1980a). In general, the consequence of the interactive approach for the study of language and communication are deep and pervasive across the full range of current approaches to the subject. Most fundamentally, this is so because virtually every current approach considers language to be essentially representational and assigns representation in language to encoded propositions. This is precluded in an interactive approach, and the consequences are far-reaching.

As standardly conceived, an utterance is an encoding of mental contents, and the task of grammar is to capture the regularities among encodings. The representational content of those encodings is approached via a wide variety of devices, but there is always some version of representational encodings. These range from semantic features (Chomsky, 1965), to natural logic (Lakoff, 1972),[45] to case-structured propositions (Fillmore, 1968), to possible worlds models (Montague, 1974), and so on. Even functional grammars retain a role for encoded representational propositions (Dik, 1978; Silverstein, 1976).

Montague grammars exemplify a general approach to language and meaning via Tarskian model-theoretic semantics. Others making use of variants of a model-theoretic approach include Field (1980), Davidson (1975, 1980),[46] Donnellan (1977), Kripke (1972, 1977), Putnam (1977), and many others. The differences among these particular approaches are striking, but the incorporation of the interactive perspective would require fundamental changes in each, because model theory is at its root a formalization and extension of encoding conceptions.[47]

Recognition that language involves more than just encoded representation was initiated most strongly by Austin (1962) and the later Wittgenstein (1958) but, in both cases, propositional encodings were retained at the core of the representational function.[48] They have been similarly retained by Searle (1969) and by Grice (1967, 1969, 1971).

Psychologists and sociologists, despite being much closer to the interactive realities of actual language behavior, have had little choice but to follow the encoding lead of linguistics and the philosophy of language. Psycholinguistics, correspondingly, is a straightforward study of presumed processing of encodings (e.g., Foss & Hakes, 1978). Those investigators whose subject matter is most directly natural language interaction, students of language acquisition and of microsociological processes, have the strongest interactive intuitions, but no interactive framework within which to formulate them (for example, Bruner, 1975; Cicourel, 1974; Dore, 1975; Turner, 1974).

Most fundamentally, the interactive approach requires that the assumption of language as being intrinsically representational be given up. Language is most deeply a phenomenon of *operations* on representations, not of representation per se, and the consequences of that shift in perspective are wide ranging (Bickhard, 1980a). Standard approaches can again be viewed as asymptotically limiting cases of the interactive approach.

The ubiquitous involvement of encoding assumptions varies in explicitness across the differing areas of psychology. In memory studies, however, it is completely explicit. If memory is concerned with the processes and organizations of the storage and retrieval of representations and if representations have no way of being conceptualized other than as encodings, then memory studies are going to look for generations, transformations, organizations, and retrievals of encodings, usually of propositions (for example, Anderson & Bower, 1973; Melton & Martin, 1972). Rather straightforwardly, "Thinking occurs at the propositional level; language is the expression of thought" (Kintsch, 1974, p. 5).

The storage of control structures, however, cannot at its most basic level involve the storage of encodings. Control structures, of course, *can* be encoded, but then there must be an interpreter for those encodings. That interpreter could, in turn, be encoded, but there must at some level be a control structure (interpreter) that exerts competent control over relevant processes via *self-(*intrinsic) interpretation. This is clearly the case, for example, with well-learned motor skills. If representation is derivative from interactive competence, thus from interactive control structures, then such a nonencoding form of memory for control structures forms an alternative foundation and context for the study of memory in general. It also provides a much more powerful perspective than usual for the integration of otherwise diverse areas such as motor, episodic, semantic, and conceptual skill memory. The root question becomes: What is the nature of the interactive representation involved? Storage, organization, and retrieval should be very strongly constrained by that nature.[49]

Encoding assumptions show up in the study of learning not only in terms of what is assumed to be learned, but also in the presumed processes of learning. Encodings are easily presumed to be learnable via some sort of passive "stamping in," association, or induction process—a kind of temporally extended transduction. Interactive control structures, on the other hand, cannot be stamped in from the environment; they must be constructed from the inside and tried out against the environment. Learning, then, must be some sort of a variation and selection process, a hypothesis-testing process, but with procedural rather than encoded hypotheses (compare with Levine, 1975). As before, an initial step is to look at what gets learned, that is, the relevant representations, from an interactive perspective. Variation and selection paths, constraints, and heuristics are organized accordingly. Again as before, the interactive perspective offers significant integration: All learning must be variations, components, or aspects of control-structure construction.[50]

The classical view of motivation arises in the context of a model in which encoded information is interpreted and made use of as the occasion warrants by an agent interacting with an environment. The agent is active only as conditions demand, and the problem of motivation is to determine what initiates activity in the agent, what turns the agent on. The problem, in other words, is to explain why there is (some) activity rather than there being none. Some sort of energy metaphor, which "drives" the agent, is the usual answer.

In the interactive perspective, however, the interactive system is not merely a component of the overall organization that can be turned on and off from outside of itself as needed. The interactive system is *all* of the system, and it is always active, always doing something (even if it be sleeping) so long as it is alive. The issue, then, is not one of why there is any activity rather than none, but instead it is one of why there is this particular activity rather than some other. The problem of motivation is to explain selections of activity, the course of interacting, the modulations of control during interacting (compare with Atkinson & Birch, 1970). Organizations of potential such selections of activity, however, *are* control structures. Thus, motivation too is an aspect of interactive competence, and representation and motivation turn out to be intimately related *aspects* of the same interactive control structures, not separate and disparate *components* as in the classical view.

Representational aspects of emotions include the beliefs that are presumed to initiate and organize emotions and the objects toward which at least some of them are directed. In the usual view, these are some sort of encodings, and in between is some sort of motivational energy with various presumed effects on the manner and direction of activity in the agent. In the interactive perspective, emotions must be interactions or aspects of interactions: There are no other possibilities. One proposal is that emotions are special forms of interaction, differentiated by that with which they interact. In particular, they are explicated as a kind of metainteraction in which the system interacts with its own internal condition of well-definedness or, conversely, 'uncertainty' about the course of interaction selections to be made. The adaptive value of emotions, then, is the ability of the system to develop generic heuristics for coping with situations of ignorance, danger, and novelty.[51] It is not clear that the interactive perspective forces this explication of emotions, but it is one plausible explication consistent with that perspective.

Examples could continue to be presented, but it is by now clear that if the critique of logically independent encodings and the proposed interactive alternative are valid, then the impact is broad and deep: representation is everywhere. Recognizing the fact of an impact, however, does not necessarily provide a clear indication as to the nature of that impact. The nature of the interactive-perspective impact requires its own investigation.

The movement to an interactive perspective, however, and the extirpation of myriads of implicit and subtle encoding commitments from the contents and manners of thought are difficult and long-term tasks. An initial step can be to develop some conceptual bridges, or translations, from a familiar area to an interactive perspective on the same area. Part of the difficulty in doing even this, however, is that the subject matter(s) may not have the same cohesions and boundaries within the interactive perspective as within an encoding perspective, so that, for example, asking the interactive approach to account for some presumed unified phenomenon X in an integrated manner may inadvertently presuppose the encoding approach in the very posing of the question.[52] Nevertheless, such attempts at "saving of the phenomena" or "accounting for the facts" within the interactive perspective are absolutely necessary, and the difficulties in doing so, once overcome, are among the richest sources of new insights: The deepest difficulties are precisely the deepest, most taken-for-granted errors in the encoding approach.

With respect to representation per se, some general translation heuristics can be outlined. Encoded representations are supposed to represent by virtue of encoding correspondences—references, denotations, and so on—with that which is represented. To use or to evoke an encoding is to invoke that correspondence relationship and thereby to specify, to mention, whatever is on the other end of the correspondence. The incoherence problem, of course, arises from asking how the 'represented' end of the encoding correspondence can be specified for a logically independent encoding.

Interactive representation is most fundamentally one of patterns of interactive potentialities via interactive implicit definition. *Knowledge of what is represented* is constituted in terms of what *further* patterns of potentialities are indicated by the occurrence of a *given* (implicitly defined) pattern, rather than in terms of the epistemically circular "knowledge of what is encoded." That is, the encoding correspondence is refined into an interactive, implicit-definitional correspondence, and the knowledge of what is represented is removed from that representational (encoding) correspondence altogether and is placed within the organization of the system's indicative uses that can be made of an instance of the implicitly defined pattern.

The *specification* of a particular interactive representation is via context-dependent differentiations within the larger pattern of represented potentialities of which it is a part, not via the invocation of context- independent encoding correspondences. Representational specification, whether mental or linguistic, is a function, a use, to which differentiations can be put, not an intrinsic property of encoding correspondences, and unique specification is a differentiation goal that can be attempted, not the quiddity of an encoding. Presumed encoding specification—reference, denotation, and so on—is an assured successful, one-step, unique (interactive) differentiation. As such, it is impoverished and opaque with respect to its mechanism and is asymptotically unreachable with respect to its accomplishment.

Encoding "contacts with the world," thus are replaced by interactive implicit definitions; encoding knowledge of the world is replaced by organizations of indications of further interactive potentialities, and encoding specifications within the world are replaced by differentiations within organizations of interactive representations: If epistemic contact with the world is implicit in the interactive relationship between the system and the world, rather than explicit as presumed with encodings, then knowledge *must* be constituted in terms of the organization of the system, and specification *must* be accomplished via functional differentiation *within* that organization of the system.[53]

The interactive perspective rests on the critique of the encoding approach and grows most directly out of the recognition of the necessity for an interactive interpreter. Its impact is most direct and obvious with respect to issues of representation, but it provides an integration of representation, motivation, and competence at its base and a framework for approaching psychological phenomena in general. The shift to an interactive perspective is difficult, but it offers both its own deep insights and the possibility of rescuing valid knowledge in encoding models via appropriate models of interactive systems, organizations, and differentiations. Encodings are limiting case approximations of interactive realities.

EPILOGUE

James Gibson's theory of perception, both internally and with respect to the controversy surrounding it, exemplifies the conflict between two fundamental approaches to the nature of representation: the encoding approach and the interactive approach. We have argued that the encoding approach forms an impoverished, asymptotically limiting case of the interactive approach and that it is insufficient, unnecessary, and incoherent when taken as a logically independent approach to the nature of representation. The issues have been examined primarily with respect to perception and with a primary focus on Gibson's theory, but those issues are fundamental to all of psychology. It is hoped and expected that the exploration of those issues will prove interesting and productive throughout their domain.

Notes

1. Such a shift to patterns in the ambient light as the locus of perception is clearly prefigured by his 1950 point that patterns of stimulation could themselves be stimuli (p. 9), even though at that time he was referring to retinal patterns. The shift is also consistent with his general ecological emphasis, but neither of these points is sufficient to force that shift—the active observer is sufficient.

2. From this perspective, in fact, the spatial is subsidiary to the functional. Surfaces, objects, and the like are constructed as patterns of potential interactions, including further perceptual interactions, that may be indicated by particular perceptual interactions, that is that may be perceived. Such construction of the physical and spatial out of the functional is in the general spirit of Piaget.

3. Affordances, of course, are therefore "*relative to the animal.* They are unique for that animal. They are not just abstract physical properties" (Gibson, 1979, p. 127). "Knee-high [therefore affording the potentiality of sitting on] for a child is not the same as knee-high for an adult" (Gibson, 1979, p. 128). Horizontal support for water bugs is different than for heavy terrestrial animals (Gibson, 1979, p. 127).

4. Gibson's discussion, however, still suggests too much independence of the spatial from the functional; there is an incomplete recognition of the construction of physical and spatial representation out of functional representation. (Such construction would be a part of Gibson's tuning, not his information extraction.) Gibson still wants to go "from surfaces to affordances," he does so by having "the composition and layout of surfaces *constitute* what they afford" (1979, p. 127), but such constitution still leaves the question of what a representation of a surface is as logically prior, though no longer temporally prior, to a representation of an affordance. Yet infants can perceive affordances without necessarily perceiving the surfaces, edges, and full objects that provide, or constitute, those affordances.

5. This model is suggested by work in Bickhard (1980a).

6. Sense organs generate inputs that indicate encounters, but they do not indicate what the encounters are with. What the encounters are with is detected by patterns of perceptual interactions and is constituted by patterns of potential interactions. This is so even though we, as observers who already have our own representations of both the environment and the sensory inputs, can take those inputs as encoding that environment ("this representation represents the same as that") and then, forgetting that this move required our own independent representations of the environment, assume that they encode that environment for the organism under consideration as well.

7. Because indicators are also interpreted, there might appear to be little distinction here. But the accessing of an indicator *is* an interpretation of it, and such accesses collectively *constitute* its representational content. Neither of these characteristics is presumed to be true of encodings.

8. A number of interesting questions arise at this point but are not pursued. These include further exploration of the nature and structure of the lattice of cognitive aspects of interactive systems and of the constraints and possibilities created by such a lattice for evolution and development.

9. There is a related issue, really a pseudo-issue, in the literature that concerns whether Gibson's theory is an active theory (Gibson, 1976; Richards, *1976).* In the sensation-based model, the perceiver is essentially passive in the sense that activity is not an essential aspect of perceiving. As we have seen, Gibson's perceiver is essentially active and interactive. There would seem to be no room for controversy here. But the term *the active perceiver* is also used to refer not to the physical and interactional activity of the perceiver, but rather to the active (enhancing) contributions that the perceiver makes to perception. As we shall see, Gibson objects strenuously to the idea of such contributory perceptual activity—though, as is by now familiar, we have some revisions to make concerning the proper scope of Gibson's arguments. The issue of the activity of the perceiver in Gibson's model, then, is really a pseudo-issue based on an equivocation on the term *active.* The more substantive issue is that of the presence or absence of contributory enhancement in perception.

10. Ullman does criticize Gibson's tendency to appeal to phenomenological experience as evidence for direct perception. Arguing that "the mediating processes ... do not operate on subjective experiences, nor are they intended to account for their origin" (1980, p. 380) and that "the perceptual processes are not necessarily open to conscious introspection" (p. 380), Ullman concludes that "the introspective impression that the perception of objects is immediate and unanalyzable cannot be taken as evidence supporting the theory of immediate visual perception" (p. 380). Ullman's argument is valid with respect to his conclusion, but it does not address the homunculus problem that is implicit in Gibson's appeal to the phenomenological level: mediating encodings, whether accessible to consciousness or not, must ultimately be meaningfully interpreted, but such meaningful understanding of the environment is what was to be explained in the first place.

11. The problem of such ad hoc proliferations of basic encoding elements is intrinsic to all encoding approaches, whether mediated or not. In a direct approach, each different *thing* to be perceived requires a distinct new encoding element; in the general (mediated) case, each different *type* of thing to be represented requires a distinct kind of encoding element. In both cases, the proliferation of encoding elements is ad hoc, and in neither case is anything explained. See Bickhard (1980a).

12. Two subsequent questions at this point are "Why *don't* we directly pick up the light interactions, intermolecular forces, and other states of affairs about which those internal outcomes contain information?" and "Why does our knowledge of such things involve inferences?" The answer to the first question is that the specifics of such states of affairs have little direct relevance to our basic goals and available ways of reaching those goals; only the interactive affordances manifested by such states of affairs are proximately relevant. In other words, we learn those specifications that are most relevant to our goals, starting as infants. The answer to the second question is that, relative to our basic ecological interactions, such states of affairs, for example, molecules, are differentiated from other possible states of affairs only by complex and specialized patterns of basic interactions. (This is largely a function of our transducers, as Fodor and Pylyshyn point out, but of transducers not as encoders, but as interactive detectors of evolutionarily selected, ecologically relevant affordances.) Such complexities of interaction pattern relative to basic ecological patterns will generally involve the differentiation and specialization of representations, decision-making rules, inference procedures, and so on. If, however, the environment were adjusted in some way so that such states of affairs became directly ecologically relevant, then those ecological relevances, those affordances, would tend to become picked up. For example, a world in which diffraction gratings and prisms were common ecological objects would give us a direct pickup of interactive possibilities provided by diffraction patterns. A world in which optical focuses were

common would yield a pickup of the opportunities for magnification, for being burned, or for starting fires. The learning to pick up such affordances, the relevant 'tuning,' might or might not be aided by explicit inferences, though they would not in general be necessary, but the eventual pickup would *not* necessarily be inferentially mediated.

13. Another possible reading of the point at the beginning of this paragraph derives from reading "has an epistemic relation to" as meaning "having an explicit representation of" rather than as "having information about." This, of course, is equivalent to reading "epistemic relation" as "encoding," and Fodor and Pylyshyn's conclusions follow. There is substantial evidence later in their article that this is the reading Fodor and Pylyshyn presuppose. Under this reading, we would deny that the subject has any epistemic (encoding) relation to the light patterns at all. When the interactive perspective is taken into account, however, such a reading, of "epistemic relation" as "encoding," is seen to be too narrow.

14. As Fodor and Pylyshyn point out, "If there is a subsystem of the organism whose states are correlated with properties of the light, then the states of that subsystem will also be correlated with the properties of the layout that the light specifies" (p. 162). They miss, however, the deeper implications of that fact.

15. "Detected" in such contexts carries connotations of "encoded" and is to that extent unfortunate. Perhaps "differentiated" would be better. Certainly "differentiated" is accurately descriptive in the sense that the particular interaction outcome, thus the light patterns sufficient to that outcome, are differentiated from other possible outcomes and corresponding sets of sufficient light patterns. And certainly Gibson relies strongly on the concept of differentiation. But detection is not necessarily of a thing or state of affairs in all of its detail and uniqueness; it is in general of an instance within a class, and neither that instance nor that class are *necessarily* explicitly represented. For example, a frog might detect a fly, in the sense of responding appropriately to the situation, without being able to differentiate finely enough to distinguish flies from moving pencil points and without having an explicit representation at all, only a response to the act of detection. Similarly, what is detected in visual perception is not a light pattern per se, but rather an instance of an affordance-equivalence class of light patterns, and what is explicitly represented, if anything at all, is the affordance, not the light. In spite of its connotation, then, "detection" is not inappropriate. The nonencoding reading of it, however, must be kept in mind.

16. It is possible that our claim that the reverse counterfactual-supporting constraint is required might be contested. It might be claimed that Fodor and Pylyshyn either do not intend or do not need it or both. Even if accepted, however, such a contention would not harm our subsequent arguments, for, in adding such an additional constraint, we can only have strengthened the case against which we argue. Thus, if our general points are telling against the stronger position that we direct them against, then they would surely be telling against the position derived by eliminating the second constraint.

17. It might appear that an appeal to interactive detection at this point might salvage the definition: it would certainly eliminate the more simplistic ways of generating S; without P. But, aside from the fact that such advantages of interactions are clearly not being considered by Fodor and Pylyshyn, and that the point is contrary to the spirit and content of their discussion (e.g., the sense in which transducers are based on physical law, p. 163), still the introduction of interactions would simply make P-illusions more difficult to create, not impossible—the proper interactive properties would have to be simulated, not just the proper causal chain initiated.

18. Fodor and Pylyshyn make a similar point, but in a different context (p. 166).

19. Fodor and Pylyshyn end their discussion of Gibson with the claim that his approach to perception ultimately founders on the problem of intentionality. The starting point is the

fact that we do not merely *see* things; we see them *as* something or other. Thus, for example, we might see a particular spot of light as a firefly, as one of myriads of stars, or as the Pole Star (pp. 188, 189). The consequences of these different ways of seeing can be drastic, if, for example, we are lost and need to know the cardinal directions. Accounting for such facts is the problem of intentionality.

Establishment theorists, according to Fodor and Pylyshyn, approach such phenomena in terms of mental representations; mental representations allow us "to understand *seeing as* in terms of *seeing* and *mentally representing*" (p. 190). We can, in these terms, represent the light as a firefly or as a star or as the Pole Star. Gibson, on the other hand, is supposed to attempt to handle such phenomena in terms of *properties* (p. 191). We pick up the differing properties, rather than invoke the differing mental representations. "Property," presumably, is Fodor and Pylyshyn's way of talking about Gibson's "affordance," though this is not made clear.

At this point, we encounter the heart of Fodor and Pylyshyn's argument:

> *Property* is an intentional notion. . . . However, *specification* is an *extensional* notion. . . . Specification cannot, then, explain property pick up (p. 191).

The idea here is that, on the one hand,

> The Morning Star = The Evening Star, but the property of being the Morning Star the property of being the Evening Star (p. 191).

while, on the other hand,

> Specification comes down to correlation . . . and if X is correlated with the Morning Star and the Morning Star = the Evening Star, then, of course, X is correlated with the Evening Star. Which is to say that, on Gibson's notion of specification, it must turn out that whatever specifies the Morning Star specifies the Evening Star too (p. 191).

There are so many missing premises in this argument that it is difficult to know where to begin. The fact of there being missing premises, and perhaps missing conclusions as well, can be seen by considering the equally "true" sentences and equally invalid and pointless argument;

> Mental representation is an intentional notion. Specification is an extensional notion. Therefore, specification cannot explain mental representation.

Furthermore, the potential missing arguments, insofar as we can infer them, appear to be egregiously wrong, not as any deep consequence of the differences between encoding and interactive perspectives, but simply in terms of *any* plausible interpretation of Gibson. (It is for these reasons that we consider this argument in a footnote rather than in the text.)

First, note that the problem of meaningful perception, the problem that is commonly solved in terms of meanings (mental representations) being attached to perceptions (seeing), is *exactly* the problem that yielded Gibson's concept of affordances. Fodor and Pylyshyn do not consider *any* of Gibson's discussion, definitions, or arguments regarding this point. Thus, they are neither addressing Gibson's own position, nor any of his reasoning that led him to that position, nor the fact that their own purported solution is wide open to Gibson's counterarguments, the homunculus argument, among others.

Second, the centrality of the specification relationship in their argument is highly misleading in several senses. Most prominent is that the correlational specification relationship has *never* been sufficient to the pickup of affordances in Gibson's model: If it were, *failure* to perceive the affordances of a layout that were 'specified' in the light would be impossible, even for an infant, and, furthermore, we would automatically perceive hosts of other states of affairs (for example, electron energy level transitions) that are also correlatively 'specified' by the light. Gibson has always focused on learning a specification as necessary to perception. But, if correlational 'specification' per se is *not* sufficient to perceiving in Gibson's model, then Fodor and Pylyshyn's argument is simply irrelevant.

Fodor and Pylyshyn, in fact, appear to contradict themselves on this issue: They point out that correlation is "objective" (p. 158), but that specification is "epistemological" (p. 159), that is, that it depends on the abilities of specific organisms to make use of such correlations (p. 159) and that correlation is thus *necessary* to specification (p. 159), but they later claim that specification is "extensional" and essentially *equate* it to correlation (p. 191). Their first definition of specification does not support their claims regarding Gibson and intentionality, while their second, inconsistent characterization is not adequate to the facts regarding Gibson's model.

The most ironic of the senses in which the advertence to the specification relationship is misleading in Fodor and Pylyshyn's argument is that its irrelevancy is formally similar when the argument is applied to mental representations as when it is applied to affordances. In both cases, the correlational specification relationship is causally mediating for an epistemic relationship; in both cases, more than such potential causal mediation is needed for perceiving to be possible, or to occur; in both cases, correlation (specification) is not sufficient to perception. In the mediated encoding case, what is needed are inferences based on knowledge of such correlations applied to encodings of the light. In Gibson's case, what is needed is the realization of the learned *direct* correlations between the internal states and the layout affordances (or affordance encodings, in a direct encoding interpretation). In either case, if the argument that specification is insufficient applies to Gibson, it applies to Fodor and Pylyshyn just as strongly.

The argument that seems most relevant to Fodor and Pylyshyn's point would seem to be that perception has intrinsically intentional characteristics, while Gibson's whole theory (not just the specification relationship) is intrinsically extensional. This seems prima facie unlikely, given the necessary role of learning in being able to pick up an affordance. It is simply false under an interactive interpretation: the mental representations are there to do the same work that Fodor and Pylyshyn want them to do (though not in the same way), but they are simply not *encoded* representations. It is equally false under a direct encoding interpretation of Gibson; there the representations are even encodings.

This leaves as an apparent last possibility the claim that Gibson leaves no possibility in his theory for any differences in mental states (representations) corresponding to differences in affordances picked up, therefore his theory is *necessarily* strictly extensional. We have seen that this is prima facie unlikely and simply false on either an interactive or a direct encoding interpretation, but Gibson's more recent claims that his approach has consequences for cognition, not just for perception (1979), might conceivably be adduced in support of his supposed commitment to extensionality. If the case could be made that Gibson's claims amount to a claim that no mental representations of *any* sort occur anywhere in perception or cognition, then the extensionality point would have at least some basis, though it would still be contradicted by the bulk of his theory *and* his metatheory, again, on either a direct encoding or an interactive interpretation. If Gibson's claims, however, are interpreted as pointing out the fact that his arguments against encodings carry impact for cognition as well as for perception, then Gibson is exactly correct, his claim is 'simply' a further elaboration and realization of his basic interactive insights, and Fodor and Pylyshyn's argument has no basis whatsoever.

20. See Bickhard (1980a) and, for a related point, Noble (1981).

21. Goal directedness is likely to be an additional intermediate level of emergence.

22. This and later related points in the text concerning emergence and explication are also discussed in Bickhard (1980a, 1982).

23. A number of challenges are addressed in Bickhard (1980a). These involve questions such as whether the interactive approach is able to explicate various forms of representation and other psychological phenomena. The ability of the interactive approach to at least programmatically meet such challenges is an important corroboration of the approach, though, of course, no finite number of such positive demonstrations can prove the overall validity of inter activism. A single conclusive failure, conversely, would disprove it. This status of being capable of falsification but not of proof is one that interactivism shares with Turing's thesis; as mentioned earlier, interactivism can in important ways be considered to be a psychological version of Turing's thesis.

There is another important sense, however, in which the interactive approach is not just *as* powerful as Turing machine theory, but is in fact *more* powerful - more powerful in the sense of being able to explicate process phenomena and abilities that Turing machine theory cannot. One partial perspective on this point is obtained by noting that a Turing machine is not, and cannot be, truly interactive in any interesting or important way. It cannot be because its environment, the tape, is static and provides nothing (interactively) interesting to be interacted with.

There is a deeper point behind this one, however, and that is that interactive competence in an environment requires *skill,* and skill requires not just formal 'information processing' or 'symbol manipulating' capabilities, but also timing and temporal coordination capabilities, and Turing machine theory has no natural way of introducing such considerations. There is no natural timing unit in Turing machine theory, nor even any sense in which the processing steps take the same or determinate multiple amounts of time as each other: the steps are simply serially ordered with no metric time considerations at all. There is, for example, nothing equivalent to an oscillator in Turing machine theory, and no way to construct one without adding to the fundamental assumptions of the formalism. A related point is that interactive functions such as multi-system environmental monitoring or coordination intrinsically require simultaneous processing across the various systems involved, while Turing machine theory is intrinsically logically serial and temporally sequential.

24. Such as, for example, the logically general strategy that is instantiated in and is thus a property of, a particular goal-directed interactive system. A strategy consisting of a structure of conditional control flows among possible subsidiary interactions (or substrategies of interactions) is, as abstracted from the particular substrategies involved, an aspectual property of any such interactive system. Aspects of strategy, in turn, include such properties as sequence and iteration (of substrategies) and, thus, (count) number.

25. Furthermore, the approach is not ad hoc: It has other important implications. For example, the Piagetian flavor of the origins of higher order abstractions in the properties of lower order interactions is matched by a consequent Piagetian-flavored developmental-stage structure. In particular, in such a hierarchy of levels of systems as discussed above, it is impossible for a system at any given level to exist without all levels below it also being instantiated—otherwise, at least one system would have no system below it with which to interact,

and thus it could not come into being. Correspondingly, development through such a hierarchy must be in invariant sequence starting at the lowest and thus must manifest a stage structure. Such stages explicate a number of the properties of Piaget's stages, but also differ in critical ways, such as in their boundaries and in the absence of any kind of structures of the whole (Bickhard, 1978, 1980a, 1980b).

26. The general form of the arguments presented is relatively indifferent to whether the basic elements to be encoded are taken to be entities, events, features, propositions, and so on. The most sophisticated versions take propositions as basic elements, and that move is an improvement over the alternatives, but it does not solve the basic problems of encodings.

27. Again the classical approach ends up as an asymptotic limiting case of the interactive approach. The classical approach assumes that the internal "predicates" differentiate uniquely and knowably down to some specified, and thereby referred to, environmental conditions. The interactive approach recognizes that such unique knowable differentiations are impossible.

28. Learning within such a perspective, learning what is being differentiated and implicitly represented, is in large part a matter of learning that an encounter with Pi indicates the possibility of Pj. Vast structures of such relationships among potentialities constitute knowledge. These issues are explored somewhat more fully in Bickhard (1980a).

Such an approach to knowledge, as structured and defined from within the system', has a flavor of coherence approaches to truth and representation, and there are similarities, but it differs in the multitudinous implicit definitional anchors that, though not referential, do connect to the environment.

29. Such procedures are essentially subroutines, but the word is at times avoided because it so strongly carries the connotations of standard computer-language subroutines, which operate on symbolically encoded data—and the status of encodings is exactly what is at issue. The essential idea of a subroutine, however,—that of a semiautonomous task performer that can be called upon as needed—is not restricted to such a symbol-manipulation environment.

30. Generalized strategy procedures are themselves special cases of kinds of procedures, here called *themes.* To understand themes, consider the progression: First, some procedures will decompose given actions or goals directly into strategies of subactions or subgoals—this is the familiar paradigmatic case; second, some procedures will decompose given actions or goals similarly to the previous case, but something about the nature or manner of that decomposition is determined not by the primary procedure, but by some procedure operating (logically) parallel to it—the primary procedure is sufficient to the decomposition, but it can be influenced in aspects of that decomposition by other sources; third, some procedures may be involved in the decomposition, computation, and specification of actions or goals, but they are insufficient to that decomposition and specification task by themselves—such procedures are concerned only with *aspects* of the computation of actions and goals (for example, an influencing procedure as mentioned in case two above); and fourth, some procedures are involved in such determinations of aspects of actions and goals together with other relevant procedures, *all* of which deal only with such aspects—no single procedure suffices to specify the action or goal, but two or more together, by specifying an exhaustive set of aspects, *can* be sufficient to the ultimate specification. The aspects that are in this manner differentiated and the procedures that differentiate them are called themes.

As an example, consider walking. It might be decomposed into the components of individual steps, for example, or, for a different example, it might be decomposed into the *aspects* of frequency and forward-backward. Neither frequency nor forward-backward suffices alone to specify an action, and neither is a component of an action, but together they can specify the action of walking (ignoring turning). The procedures that specify such characteristics are themes, as are the characteristics. The cockroach, for a rather mundane instance, seems to walk via precisely these themes (Gallistel, 1980).

Note that themes do not function like components. They do not retain their individual identities within the larger whole that they make up; they are not like building blocks. Instead, they merge and blend indivisibly. They are not individuated by the usual action and action-component boundaries of space and time. Instead, their differentiation is more strictly logical and functional. They are not necessarily easily recognizable as units. Instead, there may be multiple-potential such differentiations, and the most obvious is likely to be a decomposition into components rather than into themes. Furthermore, abstract and generalized themes may be manifested across a large number and variety of interactions, thereby making their identification perhaps easier, but simultaneously making their nature more obscure—so long as themes are not recognized as a different principle of organization from components.

Organization with respect to themes is to be found throughout psychology. Abstract cognitive strategies and procedures of many kinds are not in themselves sufficient to specify or compute interactions, and thus they constitute themes. Meanings of all kinds, including linguistic meanings, merge and blend and differentiate like themes more commonly than they add and subtract like components. Similarly with meanings in attitudes, personality, psychopathology, the 'unconscious,' and so on: Meanings are the paradigm case of organization and interaction as themes, rather than as components. Clearly, development will take place in large part in terms of themes. Themes constitute an ubiquitous and virtually unexplored subject matter in psychology.

31. Note that an important aspect of strategy in terms of subordinate goals is sequence (of subordinates), and an important aspect of sequence is iteration. As general procedures for these aspects became differentiated, one expectation would be for the development of one (or more) generalized iteration computation procedures, at least up to the number of iterations commonly found useful in strategy computations. A generalized iteration computation procedure is a counting procedure, and, correspondingly, number is a cognitive theme.

32. Note that a generalized and autonomous strategy procedure, when carried to the extreme case of only one such procedure for each logically differentiable strategy, constitutes a Piagetian structure. It is unlikely in the extreme, however, that a developing system would discover such an abstract-level strategy in the first place, or that it would discover such a strategy's full scope of application even if it had one available. The heuristics for discovering such procedures can be expected to function in terms of *locally* stable and successful structures, with consequent partial overlap and duplication of effort, just like the rest of development. An abstract similarity or commonality between two strategy procedures is a potential discovery for the system to make and exploit, not an a priori constraint on the organization of the system. Conversely, an important kind of process in development is likely to be the differentiation of a strategy that originates within a particular procedure, and with respect to a particular domain, into a separate procedure that can operate on encodings from arbitrary domains. Such a process involves a concomitant and consequent expansion of the scope of application of that strategy as the heuristics that try it out progressively learn of its usefulness.

33. These points concerning the relationships between perceptual development and general strategy development are not specific to *perceptual* development.

34. Furthermore, most of the arguments against logically independent encodings apply directly against each level in this presumed sequence, not just against the initial 'transductions'; the sequence, thus, not only collapses from lack of foundation, but it also disintegrates throughout. With respect to the arguments against the sufficiency and the necessity of logically independent encodings, for example—that there is not any level of atomic representations

suitable for the atomic encodings, that representational adequacy requires an ad hoc proliferation of types of encodings, that there are skills, which constitute knowledge that cannot be captured as encodings, that encodings require interactive interpreters (all against sufficiency), and that encodings are logically eliminatable in favor of interactive systems (against necessity)—are all easily relativized to any particular level of representation being considered.

With a sufficiently narrow perspective, however, the arguments against coherence and against any possible origin of basic encoding elements (a close relative of the coherence argument) can appear to be avoided at a particular higher level of representation by considering the encodings at that level to be dependent on the encodings at levels closer to perception (or transduction). This, of course, is not commonly the way higher level encodings are considered or defined, but, in any case, the move does not solve the problems; it simply passes them down the line.

35. The word "functionalism" is used in many ways in various literatures. It is used here in a broad sense to refer to positions that recognize the importance of abstract process patterns as distinguished from the physical realizations of those patterns and, in a narrower sense, to refer to models involving operations on encoded symbols. In both of these senses, Piaget, for example, is a functionalist.

The word is also commonly used in psychology to refer to a kind of modeling methodology that is often, though not always, associated with information-processing approaches. (Information-processing approaches per se are themselves versions of functionalism in both of the senses of the above paragraph.) This methodological functionalism is a kind of ad hoc empiricism in which the basic concern is to construct low- level, narrow-scope models that can "account" for the data in particular studies and in which the possibility of intrinsic constraints on such models, be they structural, evolutionary, developmental, or broader parsimony constraints, is relatively ignored. This version of functionalism does not explain anything by itself: The "accounting for the facts" (when it even manages that) is strictly ad hoc. Such an approach is doomed to sink under the burden of its mass of particularistic dust mote models. It is not the kind of functionalism addressed in the text. (See Beilin, 1981, which contrasts this functionalism with Piagetian structuralism; Piaget is *not* a functionalist in this sense.)

36. The implicit dualism in this sentence can also be removed, but the discussion would take us too far afield at this point.

37. This point explains the intuition that, for example, visual inputs can be identified as optic-tract action potentials, as retinal stimulations, or as ecological arrays of light (or other possibilities). From the perspective of a physical analysis, all of them are correct: They all constitute points in the path of visual interaction that provide explanatory understanding of how those interactions and their epistemic functions are realized. Epistemically, however, none of them is an input, except perhaps in the somewhat metaphorical sense that one of them might provide deeper and perhaps more encompassing explanations of how the epistemic functions are served.

38. Actually, he concludes that it will always either underselect or overselect, but there seems to be little content to the overselection argument.

39. There are serious problems with this particular example—economic variables are too interdependent to be manipulated with the required freedom—but the basic point remains.

40. How otherwise are we to understand the knowledge of those truth conditions, and the truth conditions of that knowledge of truth conditions, and so on?

41. In either case, whether practical or in principle, the argument for the impossibility of a semantics for the symbols seems prima facie contradictory to the Fodor and Pylyshyn (1981)

argument, discussed in Chapter 3, concerning transduction as the means by which such a semantics is provided. If semantics, thus transduction, *is* impossible, then what is supposed to be the impact of Fodor and Pylyshyn (1981) on Gibson?

42. Dreyfus might well argue, of course, that it cannot be so approached, but his arguments against Artificial Intelligence are not germane to that position.

43. In some cases, such as Dennett (1978a) and Haugeland (1981), the focus is on conditions for the *ascription* of meaning, intentionality, and so on by an outside observer, rather than on the *explication* of meaning, intentionality, et cetera. This is not an uninteresting question, but it ignores the issues of what it is that is being ascribed and how to account for the existence and nature of that which is being ascribed.

Dennett (1978a) has a flavor of objecting to this latter issue with the position that all that is at stake are manners of speaking ("the tactic of adopting the intentional stance," p. 8) and the conditions for using those manners, with the position that the issue of the nature of what is being ascribed in such manners of speaking is somehow void (but he appears to himself introduce considerations of what is being talked *about* in discussions of such issues as conditions of evolutionary, rational, optimal design and of language). In order for there to be nothing to talk about, the structure and conditions of the manner of speaking of the intentional stance would have to impose *no* constraints on the organization of the system at issue (no constraints more specific than "optimal design"). Dennett (1978a) does not explicitly address this presupposed lack of constraint, but in other discussions (for example, 1978b), he seems to assume that it does *not* hold.

In any case, that there should be no such constraints on what is being talked about is implausible in the extreme, and the burden of proof is clearly on anyone proposing such a position. This general position concerning "manners of speaking" assumes that conditions of speaking are superordinate to what we are speaking about, that those conditions of speaking, of application and prediction, exhaust what we mean in that speaking and that further knowledge of the inner organization of a system would never (correctly) lead us to decide that our manner of speaking in a particular case were wrong or metaphorical, so long as the predictions still worked. This general later Wittgensteinian position is far from consensual, but it is not pursued further now.

44. Similarly, most of Searle's (1981) examples are avoided. There are a few of his examples that are interactive in some sense, but the intuitions of their absurdity are much weaker, and the extent that they do seem to apply, they seem to appeal to phenomena of learning, emotions, or consciousness, not just to knowing. It does take more than just a simple interactivism to begin to account for these higher emergent phenomena, and it is in this sense that Searle's interactive counterexamples capture a valid point about interactivism (though not one that he delineated). For approaches to these phenomena within an interactive perspective, see Bickhard (1980b).

45. Issues concerning the relationships between syntax and semantics are strongly affected by interactive conceptions of the nature of semantics (Bickhard, 1980a), but these issues have a scope far exceeding those somewhat more parochial conflicts within generative grammars (for example, Seuren, 1974).

46. The holistic approach to language meaning that Davidson (1980) espouses (along with Quine), in which meaning is approached via the myriad "points of contact" between reality and the pattern of a whole language (instead of via building-block constructions out of references and denotations), can escape some of the particular problems of supposed encoded word meanings, but unless those points of contact (generally conceived of in terms of sentential truth conditions) can themselves be rendered in nonencoding terms, the most basic issues remain unchanged. No such nonencoding rendering is currently to be found.

47. It is not uncommon to find the assumption (often implicit) that model theory captures the essence of meaning and truth for formal logics, with the problem being one of extending those insights to natural language. Implicit in such a position is a failure to realize that there are nonmodel-theoretic approaches to logic and semantics. Examples include combinatorial logic (Fitch, 1974) and algebraic logic (Craig, 1974; Grandy, 1979; Henkin, Monk, & Tarski, 1971; Quine, 1966). These approaches to logic are more amenable to adaptation within interactive approaches to thought and language than are standard approaches. Furthermore, Tarskian model theory itself, thus even standard approaches to logic and mathematics, can be rendered in terms of differentiations within abstract relational patterns, rather than in terms of encoding maps (Resnik, 1981). Differentiations within relational patterns is precisely the approach to linguistic meaning that is required within the interactive perspective (Bickhard, 1980a).

48. This retention is clear in the case of Austin, but somewhat controversial in the case of Wittgenstein (Kenny, 1973). In any case, it *is* clear that Wittgenstein did not develop an interactive approach to language (Bickhard, 1980a).

49. Motor and conceptual skill representations are intuitively representable directly as control structures. Semantic memory takes on a distinctly different form from the perspective of an interactive approach to language. Episodic memory is one of the more seemingly problematic cases for an interactive approach: The intuition of stored, encoded snapshots or movie sequences is very strong. However, the interactive equivalent of a movie sequence is the continuous, temporal trajectory of apperception; an organism has reason to develop the ability to reconstruct such trajectories because not all possible apperceptive consequences can be computed in the initial case and because some unanticipated consequences may be useful later on; the storage of the necessary indicators and heuristics for reconstructing such temporal, apperceptive trajectories *is* episodic memory.

50. The focus of investigation is primarily on what gets constructed, what the construction heuristics are, and what the selection criteria are.

51. An early and brief version of this explication is presented in Bickhard (1980b). The model has been considerably deepened and extended since then, including with respect to differentiations between and within positive and negative emotions, but it has not yet been prepared for publication. It should be noted that such a metainteractive model is in principle fully consistent with possible evolutionary (thus genetic and physiological) specializations of forms of emotional reactions to certain generic situations (for example, Plutchik, 1980).

52. For example, the differing relationships among representation, motivation, and competence within the interactive perspective or the fact that the standard distinctions among syntax, semantics, and pragmatics cannot be coherently defined within the interactive perspective (Bickhard, 1980a).

53. The internal organization of the system is the only possible realization of knowledge precisely because the external relationship to the world is epistemically implicit, thus not epistemically directly available to the system, contrary to encoding presuppositions. Specification must occur within that organization because there is no other domain of knowledge or representations within which to specify. Specification must occur via functional differentiation because organizations of interactive pattern indications have no existence other than a functional one, and a functional existence is strictly a relational existence. That is, there are no properties by which such representational entities could be identified, encoded, or named except by their relational location within the overall functional organization. But to specify them by tracing down to their location within the web of functional relationships *is* to differentiate them. (For discussion of the problem of the initial location from which the differentiations begin, see Bickhard, 1980a.)

References

Anderson, J. E., & Bower, G. H. *Human associative memory.* New York: Wiley, 1973.

Atkinson, J. W., & Birch, D. *The dynamics of action.* New York: Wiley, 1970.

Austin, J. L. *How to do things with words.* New York: Oxford, 1962.

Beilin, H. *Piaget and the new functionalism.* Invited address to the Eleventh Symposium of the Jean Piaget Society, Philadelphia, May 1981.

Berkeley, G. *A new theory of vision and other writings.* New York: Dutton, 1709/1922.

Bernstein, R. J. *Praxis and action.* Philadelphia: University of Pennsylvania Press, 1971.

Bickhard, M. The nature of developmental stages. *Human Development,* 1978, *21,* 217-233.

Bickhard, M. *Cognition, convention, and communication.* New York: Praeger, 1980. (a)

Bickhard, M. A model of developmental and psychological processes. *Genetic Psychology Monographs,* 1980, *102,* 61-116. (b)

Bickhard, M. Automata theory, artificial intelligence, and genetic epistemology. *Revue Internationale de Philosophie,* 1982, vol. 36, no. 142-43, 549-66.

Block, N. Introduction: What is functionalism? In N. Block (Ed.), *Readings in philosophy and psychology* (Vol. I). Cambridge: Harvard, 1980. (a)

Block, N. Troubles with functionalism. In N. Block (Ed.), *Readings in philosophy of psychology* (Vol. I). Cambridge: Harvard, 1980. (b)

Boden, M. *Jean Piaget.* New York: Viking, 1979.

Brainerd, W., & Landweber, L. *Theory of computation.* New York: Wiley, 1974.

Bruner, J. S. The ontogenesis of speech acts. *Journal of Child Language,* 1975, 2, 1-19.

Carlson, N. R. *Physiology of behavior.* Boston: Allyn & Bacon, 1981.

Chomsky, N. *Aspects of the theory of syntax.* Cambridge: Massachusetts Institute of Technology, 1965.

Cicourel, A. *Cognitive sociology.* New York: Macmillan, 1974.

Craig, W. *Logic in algebraic form.* Amsterdam: North-Holland, 1974.

Davidson, D. Semantics for natural languages. In D. Davidson & G. Harman (Eds.), *The logic of grammar.* Encino, Cal.: Dickenson, 1975.

Davidson, D. Reality without reference. In M. Platts (Ed.), *Reference, truth and reality.* London: Routledge, 1980.

Davis, M. *Computability and unsolvability.* New York: McGraw-Hill, 1958.

Dennett, D. Intentional systems. In D. Dennett (Ed.), *Brainstorms.* Montgomery, Vt.: Bradford Books, 1978. (a)

Dennett, D. Artificial intelligence as philosophy and as psychology. In D. Dennett (Ed.), *Brainstorms.* Montgomery, Vt.: Bradford Books, 1978. (b)

Dik, S. C. *Functional grammar.* Amsterdam: North-Holland, 1978.

Donnellan, K. Reference and definite descriptions. In S. Schwartz (Ed.), *Naming, necessity and natural kinds.* Ithaca, N.Y.: Cornell, 1977.

Dore, J. Holophrases, speech acts and language universals. *Journal of Child Language,* 1975, 2, 21-40.

Dreyfus, H. From micro-worlds to knowledge representation: AI at an impasse. In J. Haugeland (Ed.), *Mind design.* Cambridge: Massachusetts Institute of Technology, 1981.

Eilenberg, S. *Automata, languages, and machines* (Vol. A). New York: Academic, 1974.

Field, H. Tarski's theory of truth. In M. Platts (Ed.), *Reference, truth and reality.* London: Routledge, 1980.

Field, H. Mental representation. In N. Block (Ed.), *Readings in philosophy and psychology* (Vol. II). Cambridge: Harvard, 1981.

Fillmore, C. J. The case for case. In E. Bach & R. Harms (Eds.), *Universals in linguistic theory.* New York: Holt, 1968.

Fitch, F. *Elements of combinatory logic.* New Haven: Yale, 1974.

Fodor, J. *The language of thought.* New York: Crowell, 1975.

Fodor, J. Methodological solipsism considered as a research strategy in cognitive psychology. In J. Haugeland (Ed.), *Mind design.* Cambridge: Massachusetts Institute of Technology, 1981.

Fodor, J., & Pylyshyn, Z. How direct is visual perception? Some reflections on Gibson's ecological approach. *Cognition,* 1981, 9, 139-196.

Foss, D., & Hakes, D. *Psycholinguistics.* Englewood Cliffs, N.J.: Prentice-Hall, 1978.

Gallistel, C. R. From muscles to motivation. *American Scientist,* 1980, *68(4),* 398-409.

Geldard, F. A. *The human senses.* New York: Wiley, 1972.

Gibson, J. J. *The perception of the visual world.* Boston: Houghton Mifflin, 1950.

Gibson, J. J. *The senses considered as perceptual systems.* Boston: Houghton Mifflin, 1966.

Gibson, J. J. New reasons for realism. *Syntheses, 1967, 17,* 162-172.

Gibson, J. J. The myth of passive perception: A reply to Richards. *Philosophy and Phenomenological Research,* 1976, *37,* 234-239.

Gibson, J. J. The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Gibson, J. J. *The ecological approach to visual perception.* Boston: Houghton Mifflin, 1979.

Gibson, J. J., & Gibson, E. J. Perceptual learning: Differentiation or enrichment? *Psychological Review,* 1955, *62,* 32-41.

Glass, A., Holyoak, K., & Santa, J. *Cognition.* Reading, Mass.: Addison-Wesley, 1979.
Grandy, R. *Advanced logic for applications.* Dordrecht, Holland: Reidel, 1979.
Grice, H. P. Meaning. In P. F. Strawson (Ed.), *Philosophical logic.* London: Oxford, 1967.
Grice, H. P. Utterer's meaning and intention. *Philosophical Review*, 1969, *78,* 147-177.
Grice, H. P. Utterer's meaning, sentence meaning, and word meaning. In J. R. Searle (Ed.), *The philosophy of language.* London: Oxford, 1971.
Haugeland, J. The nature and plausibility of cognitivism. In J. Haugeland (Ed.), *Mind design.* Cambridge: Massachusetts Institute of Technology, 1981.
Helmholtz, J. *Treatise on physiological optics* (J. P. C. Southall, trans.). Rochester, N.Y.: Optical Society of America, 1896/1952.
Henkin, L., Monk, J., & Tarski, *A. Cylindric algebras.* Amsterdam: North Holland, 1971.
Hopcroft, J., & Ullman, J. *Formal languages and their relation to automata.* Reading, Mass.: Addison-Wesley, 1969.
Kenny, A. *Wittgenstein.* Cambridge: Harvard, 1973.
Kintsch, W. *The representation of meaning in memory.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1974.
Kripke, S. Naming and necessity. In D. Davidson & G. Harman (Eds.), *Semantics of natural language.* Dordrecht, Holland: Reidel, 1972.
Kripke, S. Identity and necessity. In S. Schwartz (Ed.), *Naming, necessity and natural kinds.* Ithaca, N.Y.: Cornell, 1977.
Lakoff, B. Linguistics and natural logic. In D. Davidson & G. Harman (Eds.), *Semantics of natural language.* Dordrecht, Holland: Reidel, 1972.
Levine, M. *A cognitive theory of learning.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1975. ,
Melton, A. W., & Martin, E. (Eds.). *Coding processes in human memory.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1972.
Minsky, M. *Computation.* Englewood Cliffs, N.J.: Prentice-Hall, 1967.
Montague, R. *Formal philosophy: Selected papers of Richard Montague* (R. H. Thomason, Ed.). New Haven: Yale, 1974.
Müller, J. The specific energies of nerves. In W. Dennis (Ed.), *Readings in the history of psychology.* New York: Appleton-Century-Crofts, 1838/1948.
Neisser, U. *Cognition and reality.* San Francisco: W. H. Freeman, 1976.
Noble, W. Gibsonian theory and the pragmatist perspective. *Journal for the Theory of Social Behavior,* 1981, *11,* 65-85.
Piaget, J. *The Mechanisms of perception.* New York: Basic Books, 1969.
Piaget, J. *Genetic epistemology.* New York: Columbia, 1970.
Plutchik, R. *Emotion.* New York: Harper & Row, 1980.
Putnam, H. Is semantics possible? In S. Schwartz (Ed.), *Naming, necessity, and natural kinds.* Ithaca: Cornell, 1977.

Quine, W. V. Variables explained away. In W. V. Quine (Ed.), *Selected logic papers.* New York: Random House, 1966.
Resnik, M. Mathematics as a science of patterns: Ontology and reference. *Nous,* 1981, 15(4), 529-550.
Richards, R. J. James Gibson's passive theory of perception: A rejection of the doctrine of specific nerve energies. *Philosophy and Phenomenological Research,* 1976, *37,* 234-239.
Rogers, H., Jr. *Theory of recursive functions and effective computability.* New York: McGraw-Hill, 1967.
Scheffler, I. *Four pragmatists.* New York: Humanities, 1974.
Searle, J. R. *Speech acts.* London: Cambridge, 1969.
Searle, J. R. Minds, brains, and programs. In J. Haugeland (Ed.), *Mind design.* Cambridge: Massachusetts Institute of Technology, 1981.
Seuren, P. A. M. (Ed.). *Semantic syntax.* London: Oxford, 1974.
Shaw, R., & Bransford, J. Introduction: Psychological approaches to the problem of knowledge. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
Silverstein, M. Shifters, linguistic categories, and cultural description. In K. Basso & H. Selby (Eds.), *Meaning in anthropology.* Albuquerque: University of New Mexico, 1976.
Thayer, H. *Meaning and action.* New York: Bobbs Merrill, 1973.
Turner, R. Words, utterances and activities. In R. Turner (Ed.), *Ethnomethodology.* Baltimore: Penguin, 1974.
Ullman, S. Against direct perception. *The Behavioral and Brain Sciences,* 1980, *3,* 373-381.
Wittgenstein, L. *Philosophical investigations.* New York: Macmillan, 1958.

Author Index

Subject Index