Cognition, 9 (1981) 139–196

How direct is visual perception?: Some reflections on Gibson's "Ecological Approach"*

J. A. FODOR

Massachusetts Institute of Technology

Z. W. PYLYSHYN

University of Western Ontario

1. Introduction

There is always a scientific Establishment, and what it believes is always more or less untrue. Even in the respectable sciences empirical knowledge is forever undergoing reformulation, and any generation's pet theories are likely to look naive when viewed from the perspective of thirty or forty years on. In psychology, however, reformulation tends to be radical. When the dominant paradigm goes, the whole picture of the mind may shift; and, often enough, the scientific consensus about what constitutes a psychological explanation changes too. At such times, to use a phrase of Gibson's, the "old puzzles disappear,"[1] (p. 304) and one may be hard put to understand what on earth one's predecessors thought that they were up to. This has happened so often in the history of psychology that it would surely be unwise to assume that it is not going to happen again; in particular, it would be unwise to assume that it is not going to happen to *us*. Gibson thinks that it has already, and it seems that a substantial minority of the cognitive science community is inclined to agree with him. The purpose of this essay is to examine whether there is anything to that claim. In particular, we will examine the thesis that the postulation of mental processing is unnecessary to account for our perceptual relationship with the world; that if we describe the environment in the appropriate terms we see that visual perception is *direct* and requires only a selection from information present in the ambient light.

The current Establishment theory (sometimes referred to as the "information processing" view) is that perception depends, in several respects presently to be discussed, upon *inferences*. Since inference is a process in which premises are displayed and consequences derived, and since that takes time, it is part and parcel of the information processing view that there is an

*Reprint requests should be sent to Dr. J. Fodor, Center for Cognitive Science, M.I.T., Cambridge, Mass. 02139, U.S.A.

[1] All references are to Gibson, 1979 except as otherwise noted.

intrinsic connection between perception and memory. And since, finally, the Establishment holds that the psychological mechanism of inference is the transformation of mental representations, it follows that perception is in relevant respects a computational process.

What makes Gibson's position seem outrageous from the Establishment perspective is that it is presented as an outright denial of every aspect of the computational account, not merely as a reformulation of parts of it. According to Gibson, the right way of describing perception is as the "direct pickup" of "invariant properties". (More precisely, we are taking Gibson to be claiming this: for any object or event x, there is some property P such that the direct pickup of P is necessary and sufficient for the perception of x.) Now, what is "direct" is *ipso facto* not mediated; in particular, according to Gibson, perception is not mediated by memory, nor by inference, nor by any other psychological processes in which mental representations are deployed. Moreover, Gibson insists upon the radical consequences of his unorthodoxy: "The ecological theory of direct perception... implies a new theory of cognition in general." (p. 263).

In his last book, which will serve as the basis for our discussion, Gibson elaborates on the views he has arrived at after thirty years of research on perception, and on the bases of his disagreement with the Establishment position. The tone of the book, when it comes to Gibson's relation to received psychological theorizing is pretty intransigent:

> "The simple assumption that the perception of the world is caused by stimuli from the world will not do. The more sophisticated assumption that perceptions of the world are caused when sensations triggered by stimuli are supplemented by memories will not do... 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, and there are too many perplexities entailed by this theory. It will not do, and the approach should be abandoned... What sort of theory, then, will explain perception? Nothing less than one based on the pickup of information..." (p. 238).
>
> "The theory of information pickup differs radically from the traditional theories of perception. First, it involves a new notion of perception, not just a new theory of the process. Second, it involves a new assumption about what there is to be perceived. Third, it involves a new concept of the information for perception... Fourth, it requires the new assumption of perceptual systems with overlapping functions... Finally, fifth, optical information pickup entails an activity of the system not heretofore imagined by any visual scientist... (p. 239). Such is the ecological approach to perception. It promises to simplify psychology by making old puzzles disappear." (p. 304).

We will suggest that there is a way of reading Gibson which permits the assimilation of many of his insights into the general framework of Establishment psychological theorizing. Moreover, given this conciliatory reading, much that Gibson says is seen to be both true and important; and it does indeed differ in significant respects from what has generally been assumed by psychologists who accept the information processing framework. But, as should be clear from the preceding quotes, Gibson *does not want* to be read in a conciliatory way. And for good reason: if the program as he presents it were to succeed, it would constitute a conceptual revolution on the grand scale. Many of the deepest problems in cognitive psychology and the philosophy of mind would be bypassed, and the future of research in both disciplines would be dramatically altered. Such a possibility may seem particularly attractive to those who believe that our current understanding of psychological processes has been too much influenced by the achievements of computer technology. And it will appeal, too, to those who feel that the anti-behaviorist revolution in cognitive psychology has gone too far; a sentiment with which Gibson is by no means unsympathetic.

We will argue, however, that Gibson's claim to have achieved, or even to have initiated, such a fundamental reformulation of the theory of mind simply cannot be sustained. The main line of our argument will go like this: Gibson's account of perception is empty *unless* the notions of 'direct pickup' 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 pickup of an invariant, then the identification of perception with the pickup 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 of 'invariant' and 'pickup' can be appropriately constrained only on the assumption that perception is inferentially mediated. This is hardly surprising: Gibson and the Establishment agree that pickup and inference exhaust the psychological processes that could produce perceptual knowledge; hence, the more pickup is constrained, the more there is left for inference to do.

It will turn out, in the case of visual perception, that at least two constraints upon pickup are required. First, nothing can be picked up except a certain restricted class of properties of the ambient light. Second, spatio-temporal bounds on the properties that are picked up are determined by what stimuli turn out to be "effective"; i.e., sufficient to cause perceptual judgements. The consequence of the first restriction is that all visual perception must involve inferences based upon those properties of the light that are directly detected; in particular, all visual perception of features of objects in the

environment requires such inferences. The consequence of the second restriction is that visual perception typically involves inference from the properties of the environment that are (to use Gibson's term) "specified" by the sample of the light that one has actually encountered to those properties that would be specified by a more extensive sample. This sort of inference is required because the causally effective stimulus for perception very often underdetermines what is seen. These two kinds of inference are, however, precisely the ones that information processing theories have traditionally assumed must mediate visual perception. We will therefore conclude that Gibson has not offered a coherent alternative to the Establishment view; indeed, that the Establishment view falls out as a consequence of the attempt to appropriately constrain Gibson's basic theoretical apparatus.

2. The trivialization problem

The easiest way to see that constraints on the notion of invariant and pickup are required is to notice that, in the absence of such constraints, the claim that perception is direct is *bound* to be true simply because it is empty. Suppose that under certain circumstances people can correctly perceive that some of the things in their environment are of the type P. Since you cannot correctly perceive that something is P unless the thing is P, it will always be trivially true that the things that can be perceived to be P share an invariant property: namely, *being* P. And since, according to Gibson, what people do in perceiving is directly pick up an appropriate invariant, the following pseudoexplanation of any perceptual achievement is always available: to perceive that something is P is to pick up the (invariant) property P which things of that kind have. So, for example, we can give the following disarmingly simple answer to the question: how do people perceive that something is a shoe? There is a certain (invariant) property that all and only shoes have—namely, the property of being a shoe. Perceiving that something *is* a shoe consists in the pickup of this property.

It is quite true that if you do psychology this way, the old puzzles tend to disappear. For example many psychologists have wondered how somebody like Bernard Berenson managed to be so good at perceiving (i.e., telling just by looking) that some painting was an authentic Da Vinci. This problem is one of those that disappears under the new dispensation, since there is obviously some property that all and only genuine Da Vincis share; namely, the property, *having been painted by Da Vinci*. What Berenson did was simply to pick up this invariant.[2]

[2] See facing page.

Clearly this will not do, and we do not suppose that Gibson thinks it will. Although he never discusses the issues in quite these terms, it is reasonably evident from Gibson's practice that he wishes to distinguish between what is *picked up* and what is *directly perceived.* In fact, Gibson ultimately accepts something like our first constraint—that what is picked up in visual perception is only certain properties of the ambient light array. Gibson is thus faced with the problem of how, if not by inferential mediation, the pickup of such properties of light could lead to perceptual knowledge of properties of the environment. That is: how, if not by inference, do you get from what you pick up about the light to what you perceive about the environmental object that the light is coming from? If Gibson fails to face this difficulty, it is because of a curious and characteristic turn in his theorizing: when he is being most careful, Gibson says that what is picked up is the *information* about the environment which is contained in the ambient array. We shall see that it is close to the heart of Gibson's problems that he has no way of construing the notion *the information in the ambient array* that will allow it to do the job that is required.

Pursuing the main course of Gibson's attempt to constrain the notion of pickup will thus bring us, eventually, to the notion of the "information in the light". There are, however, other passages in Gibson's writings that can also plausibly be viewed as attempts to impose constraints on the notions of pickup and invariance. We will discuss several of these proposals, but we want to emphasize that it is not clear which, if any, of them Gibson would endorse. This deserves emphasis because the constraints are not only non-coextensive,

---

[2]The problem that we are raising against Gibson is, to all intents and purposes, identical to one that Chomsky (1959) raised against Skinner. Chomsky writes: "A typical example of *stimulus control* for Skinner would be the response... to a painting... *Dutch.* (Such responses are said by Skinner to be) 'under the control of extremely subtle properties of the physical object or event' (p. 108). Suppose instead of saying *Dutch* we said *Clashes with the wallpaper, I thought you liked abstract work, Never saw it before...*, or whatever else might come into our minds when looking at a picture... Skinner could only say that each of these responses is under the control of some other stimulus property of the physical object. If we look at a red chair and say *red,* the response is under the control of the stimulus *redness*; if we say *chair,* it is under the control of (the property) *chairness,* and similarly for any other response. This device is as simple as it is empty... properties are free for the asking... (p. 52 in Block, 1980; Chomsky's page reference is to Skinner, 1957)." If one substitutes 'the property picked up in perception' for 'the stimulus property controlling behavior', it becomes apparent how similar in strategy are Skinner's antimentalism and Gibson's. There is, however, this difference: Skinner proposes to avoid vacuity by requiring that the 'controlling stimulus' be physically specified, at least in principle. Chomsky's critique thus comes down to the (correct) observation that there is no reason to believe that anything physically specifiable *could* play the functional role vis à vis the causation of behavior that Skinner wants controlling stimuli to play; the point being that behavior is in fact the joint effect of impinging stimuli *together with the organism's mental states and processes.* Gibson has the corresponding problem of avoiding triviality by somehow constraining the objects of direct perception; but, as we shall see, he explicitly rejects the identification of the stimulus properties that get picked up with physical properties.

they are not even mutually consistent; and none of them is consistent with *all* of the things that Gibson describes as being directly perceived. So this is very much a matter of our reconstruction of Gibson's text. The reason it is worth doing is that we will argue that there is, in fact, *no* satisfactory way of constraining the notions of invariant and of pickup so as both to exclude the sort of trivialization discussed above and at the same time to sustain the thesis of unmediated perception; and to make such an argument one has to consider all the possible ways of interpreting Gibson's views.

2.1. First gambit: Only the ecological properties of the environment are directly perceived

Gibson's last book starts with the observation that "Physics, optics, anatomy and physiology describe facts, but not facts at the level appropriate for the study of perception." (p. xiii). The first section of the book is then devoted to sketching an alternative taxonomy in terms of *ecological* properties of environmental objects and events. Gibson provides many examples of properties that are to count as ecological and some examples of properties that are not. The former include some properties of objects (for example, texture, shape, illumination, reflectance, and resistance to deformation are mentioned). There are also ecological properties of arrangements of objects and of surfaces. For example, being *open* or *cluttered* are ecological properties of what Gibson calls the "layout" of an environment (an open layout is one which consists of just a ground, horizon and sky; a cluttered layout is one that has objects scattered on the ground). Similarly, containing a hollow or an enclosure is to count as an ecological property of a layout.

This list by no means exhausts the examples that Gibson provides, nor are we to assume that the examples he provides exhaust the category of ecological properties. There is, however, one class of ecological properties which requires special mention: the "affordances". Affordances are certain properties of objects which somehow concern the goals and utilities of an organism. So, being edible is an affordance of certain objects, as is being capable of being used as a weapon or tool, being an obstacle, being a shelter, being dangerous or being a potential mate. Roughly, affordances are *dispositional* properties (because they concern what an organism *could* do with an object); and they are *relational* properties (because different organisms can do different things with objects of a given kind).

According to Gibson then, "the environment of any animal (and of all animals) contains substances, surfaces, and their layouts, enclosures, objects, places, events and other animals... The total environment is too vast for description even by the ecologist, and we should select those features of it

that are perceptible by animals like ourselves." (p. 36). When, by contrast, Gibson gives examples of properties that are *not* ecological, they tend also to be properties that things *cannot be perceived to have.* "Perceiving" here means something like telling-by-looking. (Perceiving by the use of instruments does not count as a core case for Gibson). So, properties like being made of atoms, or being a thousand light years away are offered as instances of *non*-ecological properties. This makes it seem as though Gibson has it in mind that "ecological" and "directly perceivable" should be interdefined, as is also suggested by the quotation just cited.

But, of course, that will not work. If the notion of an ecological property is to serve to constrain the notion of direct perception, then it cannot be stipulated to embrace all properties that are "perceptible by animals like ourselves". Consider again the property of being a shoe. This is clearly a property that we can perceive things to have, hence it is a property we can *directly* perceive, assuming that being ecological is a sufficient condition for being perceptible. But this means that introducing the construct 'ecological property' has not succeeded in constraining the notion of direct perception in such a way as to rule out vacuous explanations like "the way that you perceive a shoe is by picking up the property it has of *being a shoe*". If all properties that can be perceived are *ipso facto* ecological, then the claim that perception is the pickup of ecological properties is vacuously true. What we need, of course, is some criterion for being ecological *other than perceptibility.* This, however, Gibson fails to provide.

2.2. Second gambit: Only the projectible properties of ecological optics are directly perceived

We have just seen that if by "ecological properties" Gibson means *all* perceptible properties, then the notion of an ecological property will not serve to constrain the notion of direct pickup. Perhaps, then, only some independently specifiable subset of the ecological properties should count as directly perceptible. In particular, the directly perceptible properties might be the ones that figure in the laws of the science of "ecological optics".

There are, according to Gibson, *laws* about ecological properties of the environment. The laws that get discussed most in Gibson's text are the ones which connect ecological properties with features of the light that objects in the environment emit or reflect. For example, such laws connect certain sorts of discontinuities in the light array with the spatial overlap of surfaces of environmental objects; and they connect flow patterns in the light array with characteristic alterations of the relative spatial position of the observer and the object being observed. Similarly, Gibson presents the following

"tentative hypothesis". "Whenever a perspective transformation of form or texture in the optic array goes to its limit and when a series of forms or textures are progressively foreshortened to this limit, a continuation of the surface of an object is specified as an occluding edge." Presumably, if this hypothesis is true, then the relation between the occlusion and the transformation of the textures is lawful, and the generalization that the hypothesis expresses is a law of ecological optics.

Now, it is generally held that laws of a science are distinguished by, among other things, characteristic features of their vocabulary (see Goodman, 1954). Only certain sorts of predicates can appear in a law, those being the ones which pick out natural kinds in the domain that the law subsumes. We need such a notion of "natural kind" in order to explain a striking difference between laws and mere true generalizations: the former hold in counterfactual cases (hence, they apply to unexamined instances) and the latter do not.

Consider, for example, the following two generalizations: *all mammals have hearts* and *all mammals are born before 1982.* The point is that (as of this writing) *both generalizations hold for all the observed cases.* To date, there have been no observations of mammals without hearts, and there have been no observations of mammals born after 1982. The difference between the cases is that, whereas the observation of a large number of mammals with hearts (and none without) is grounds for believing that there *could be* no mammals without hearts, the observation of a large number of mammals born before 1982 (and none born after) provides no reason at all for believing that there could be no mammals born in 1983. The idea, then, is that the property *being born before 1982* fails to subsume a natural kind; it is not the sort of property in virtue of which things enter into lawful relations. Since generalizations about things which happen to have that property are not laws, there is no reason for believing that they will hold in *new* cases. The inductive "confirmation" of such generalizations provides no rational basis for making predictions.

We will borrow a term from the philosophy of science and refer to predicates that appear in laws as "projectible predicates", and we will say that projectible predicates express "projectible properties". To say that a predicate is *not* projectible is thus to say that there are no laws about the property that it expresses. For example, the predicate "is my grandmother's favorite metal" is nonprojectible since, presumably, there are no laws that apply to things in virtue of their being metal of my grandmother's favorite kind. Notice that this is still true even on the assumption that my grandmother's favorite metal is gold and that there *are* laws about gold. This is because *being my grandmother's favorite metal* and *being gold* are different properties,

and the laws about gold would continue to hold even if my grandmother's taste in metals were to change. Coextensive properties may differ in projectibility (see also, footnote 8).

To return to Gibson: the projectible ecological properties would be the ones which are connected, in a lawful way, with properties of the ambient light. It would thus be in the spirit of much of Gibson's text to suggest that it is the projectible ecological properties, and only those, that are the possible objects of direct visual perception. This would at least rule out the direct perception of such properties as having been painted by Da Vinci since, presumably, there are no laws, ecological or otherwise, which subsume objects in virtue of their possession of that property (whereas, on Gibson's assumptions, there *are* laws which subsume objects in virtue of such of their properties as their surface texture—see above).

As will presently become clear, we think that there is much to be said for explicating the notion of a directly detected property by reference to the notion of projectibility. Nevertheless, this move will not do much for Gibson, for the following reasons:

a. Not all projectible properties are directly perceived on Gibson's view. For example, the projectible properties of classical optics are not; that is why you need *ecological* optics to construct a theory of visual perception. That classical optics fails to taxonomize properties in the ways that a theory of direct visual perception requires is, in fact, among Gibson's favorite themes. So, then, if the distinction between directly perceptible properties and others is to be explicated by reference to the projectible *ecological* properties, and if the explication is to be noncircular, we need a principled way of distinguishing between ecological laws and laws of other kinds. This, however, Gibson does not provide. Rather, insofar as Gibson is explicit about the matter at all, the notion of an ecological law is introduced by reference to the notion of an ecological property (e.g., ecological laws connect ecological properties to properties of the ambient light). But, as we have seen, the notion of an ecological property appears to be characterizable only by reference to the notion of a property that is directly perceivable (e.g., by "animals like ourselves"). And, of course, it was precisely the notion of direct perception that needed explication in the first place.

b. Not all of the properties that Gibson wants to be directly perceptible are plausibly considered to be projectible; in particular, affordances usually are not projectible. There are, for example, presumably no laws about the ways that light is structured by the class of things that can be eaten, or by the class of writing implements, though being edible or being a writing implement are just the sorts of properties that Gibson talks of objects as affording. The best one can do in this area is to say that things which share their affordances often (though, surely, not always) have a characteristic shape (color, texture, size, etc.) and that there are laws which connect *the shape* (etc.) with properties of the light that the object reflects. But, of course,

this consideration does Gibson no good, since it is supposed to be the affordances of objects, not just their shapes, that are directly perceived. In particular, Gibson is explicit in denying that the perception of the affordances of objects is mediated by inference from prior detection of their shape, color, texture, or other such "qualities".

In short, if we assume (as we should) that being a Da Vinci (or a pencil, or a shoe) is *not* projectible, we are in need of an explanation of how people perceive that some paintings are Da Vincis (or that some objects are shoes). The natural view would be: the Da Vincihood of an object (or its shoehood) is inferred from those of its (projectible) properties that are directly perceived. But this is the Establishment solution; precisely the one that Gibson is pledged to avoid.

As is customary with dilemmas, Gibson's has two horns. Either you trivialize the notion of a projectible property by stipulating that all perceptible properties are projectible; or you assume that some perceptible properties are not projectible, in which case you need to say how the perception of these nonprojectible properties is possible. The Establishment story is that the detection of nonprojectible properties is inferential, but that is the route that Gibson has eschewed. In either case, projectibility is not doing the job that Gibson needs done: *viz.* to provide a notion of direct perception that is simultaneously nonvacuous and compatible with the doctrine that perception is immediate.

2.3. Third gambit: Only phenomenological properties are directly perceived

Introspection suggests that the world is perceptibly accessible under some descriptions but less so under others. A landscape, for example, is readily seen as containing fields, trees, sky, clouds, houses, shrubs, and stones. But it takes special training to see those properties of a landscape which a convincing *trompe l'oeil* painting would have to duplicate; typically, properties which depend on a fixed locus of observation. It is a matter of considerable significance that properties of the world that seem to be perceptually accessible are generally ones that children learn early to recognize and to name.

Suppose we call these relatively accessible properties of things their *phenomenological* properties. Then much of what Gibson says can be construed as suggesting that it is phenomenological properties, and only those, that are directly perceived. This may be what is at issue in Gibson's injunction that the environment must be described in *meaningful* terms: "...the qualities of the world in relation to the needs of the observer are experienced directly", whereas "sensations triggered by light, sound, pressure and chemicals are merely incidental." (p. 246).

Phenomenological properties are accorded a similarly central role in Gibson's discussion of ontogenesis. "...the infant does not begin by first discriminating the qualities of objects and then learning the combinations of qualities that specify them. Phenomenological objects are not built up of qualities; it is the other way around. The affordance of an object is what the infant begins by noticing. The meaning is observed before the substance and the surface, the color and the form, are seen as such." (p. 134).

If we go by introspection alone, the identification of the perceptually accessible properties with those that are directly perceived certainly seems plausible: phenomenological properties are precisely the ones which strike one as "given" rather than inferred. Gibson says such things as that "the perceiving of the world entails the coperceiving of where one is in the world and of being in the world at that place" (p. 200) and "the environment seen-at-this-moment does not constitute the environment that is seen." (p. 195). And these remarks (with which, by the way, Husserel would have been entirely comfortable) seem true enough in light of introspections of perceptual salience. There is a scale of phenomenological accessibility, and locations, objects, and affordances are high on this scale. Contrariwise, the "sensory properties" which function as the bases of perceptual inference in, for example, Helmholtzian versions of the Establishment theory, do seem to be very low in phenomenological accessibility.

There are, however, three objections to the proposal that we take the phenomenological properties to be directly perceived. The first is internal: the proposal fails to include some of Gibson's own favorite examples of ecological invariants. For example, the slant of surfaces, the gradients and flows of textures, the amount of texture occluded by interposing objects, the moving occluding texture edge, etc., are *not* phenomenologically accessible. Witness the fact that it requires delicate experimentation to discover the central role that the detection of such properties plays in perception. Roughly, the present proposal has difficulties complementary to those of the suggestion that the object of direct perception is the projectible properties of ecological optics (see above); whereas the projectibility criterion leaves the affordances out, the phenomenological criterion lets almost only the affordances in. This is not surprising; you would not really expect the properties in virtue of which objects satisfy laws to be coextensive, in the general case, with those which are phenomenologically accessible. If such a general coextension held, doing science would be a lot easier than it has turned out to be.

Second, it seems at best a dubious strategy to infer direct perception from phenomenological salience: perhaps the latter is *solely* a phenomenon of conscious access and tells us nothing about the nature of perception *per se*. This is, in any event, a familiar claim of Establishment theories, and it is

often rendered persuasive by experimental demonstrations that the perception of phenomenologically salient properties of the stimulus is causally dependent upon the detection of features whose phenomenological accessibility is negligible; properties of the stimulus which may, in fact, entirely escape the conscious notice of the subject. For example, Hess (1975) has shown that a variety of highly accessible perceived properties of faces—including their relative attractiveness—depends on the detection of the relatively *in*accessible property of pupilary diameter. In the light of such findings, Gibson cannot, in any event, establish the identification of directly perceived properties with phenomenologically salient ones by fiat; he cannot simply assume that what is most readily reported is what is noninferentially perceived.

Finally, we are going to need a *mechanism* for the direct perception of phenomenological properties, and it is hard to imagine one that will work in the case of properties like the affordances. It is, for example, not good enough merely to *say* that we directly perceive that a rock can be used as a weapon; we need an account of how the apprehension of such a property *could* be noninferential. We will see, presently, that Establishment theories do propose mechanisms for the direct pickup of certain sorts of stimulus properties; but it is a consequence of the Establishment proposal that affordances (and, indeed, most phenomenologically salient properties) are inferred rather than directly perceived. Gibson sometimes speaks of the perceptual mechanism as "resonating" to the values of ecological parameters that they are "tuned" to. But since a more detailed account does not appear to be forthcoming, the resonance metaphor amounts to little more than whistling in the dark. We shall return to this issue further on.

2.4. Fourth gambit: What is directly perceived is whatever "perceptual systems" respond to

It is a point that we will presently make much of—and that Gibson is reasonably clear about—that *all* theories of perception must acknowledge the direct pickup of *some* properties. In Establishment theories, what is directly picked up is often taken to be the properties to which *transducers* respond. There is a circle of interdefined notions here, a directly detected property being one to which a transducer responds, and a transducer being a mechanism that responds directly to the properties that it detects. One way that Establishment theories have of breaking out of this circle is by specifying —typically by enumeration—which organs are to count as transducers; for example, the retina in the case of vision and the tympanic membrane in the case of audition.

We shall have more to say about how the notion of transduction can be constrained presently, and we will argue that such specification by anatomical enumeration is inadequate. The present point is that Gibson recognizes that to specify what is to count as a perceptual organ is implicitly to constrain what a theory says is directly picked up. For example, if you think that the organ of visual transduction is the retina, and if you can show that the retina responds only to such properties as the wavelength and intensity of light, then you are committed to the view that only those properties are directly detected. Consequently, other properties of the light (and, *a fortiori*, all visual properties of distal objects) are apprehended only *indirectly*, presumably *via* processes that involve inference.

Gibson believes that the perceptual organs have been misidentified by Establishment theorists. Correspondingly, he claims that if one individuates the perceptual organs correctly, one gets a new and better census of the immediately perceived properties. So, "Helmholtz argued that we must deduce the causes of our sensations because we cannot detect them... The alternative is to assume that sensations triggered by light, sound pressure, and chemicals are merely incidental, that *information* is available to *perceptual systems*, and that the qualities of the world in relation to the needs of the observer are experienced directly" (p. 246, emphasis added). It is a moral of *The Ecological Approach to Visual Perception*, and it is the main point of *The Senses Considered as Perceptual Systems* (Gibson, 1966) that the "perceptual system" for vision is the entire complex consisting of "first, the lens, pupil, chamber and the retina... Second, the eye with its muscles in the orbit... Third, the two eyes in the head... Fourth, the eyes in a mobile head that can turn... Fifth, the eyes in a head on a body..." (p. 245). It is the discriminative capacity of this system—*and not the discriminative capacity of the retina*—which determines what vision can, in principle, detect.

We can certainly grant that the class of properties to which this complex system is specifically "tuned"—the class of properties it can "directly respond to"—may not be the class of properties that Establishment theories have usually taken to be visually transduced. (It is far from clear that it will be the class of ecological properties either. But as we remarked above, the criteria we are ascribing to Gibson for selecting candidate objects of direct visual perception are not, in general, coextensive.) So, Gibson is right to claim that reparsing the system of perceptual organs provides for, or at least permits, a new census of directly detected properties. It follows that *if* Gibson had a motivated criterion for deciding what is to count as a perceptual system, he would *ipso facto* have a principled way of constraining the notion of direct pickup.

But Gibson provides *no* criterion for identifying perceptual systems, or even for circumscribing which organs can in general be regarded as parts of the same perceptual system. For example, it is notable that Gibson's enumeration of the parts of the visual system does not include the brain. Inasmuch as Gibson emphasizes that perceptual systems can overlap (different such systems may share anatomically individuated organs) this exclusion seems, to put it mildly, unmotivated. If, however, the brain *is* included as a component of the visual system, then presumably the properties that the visual system can pick up would *ipso facto* be coextensive with the properties that people can visually perceive and we are back where we started. We still want independent characterizations of "perceive" and "pick up directly" if the identification of perception with direct pickup is to amount to an empirical hypothesis.

It is clear from Gibson's discussion of perceptual systems that he intends to individuate them functionally rather than anatomically, a decision which we applaud. The problem is that the proposed criteria of individuation are so flexible that the notion of "perceptual system" actually provides *no* constraint on what might count as a "directly detected" invariant. According to Gibson, there are five overlapping perceptual systems, each of which can "...orient, explore, investigate, adjust, optimize, and come to an equilibrium...". The functioning of these systems is explicitly *not* limited to the transduction of impinging stimulation. Rather, the responses of perceptual systems are "specific to the qualities of things in the world, especially affordances" (p. 246). Furthermore, the nature of the information which such systems can pick up "becomes more and more subtle, elaborate and precise with practice". Given the unbounded scope of the activities that perceptual systems can perform, there would seem to be nothing in the notion that prevents the detection of shoes, grandmothers, genuine Da Vincis, performances of Beethoven's Kreutzer Sonata, or authentic autographs of George Washington all being possible "achievements of perceptual systems". It looks as though whatever is perceived is *ipso facto* the proper object of a perceptual system, and whatever is the proper object of a perceptual system is *ipso facto* perceived directly; we have, in particular, no independent constraints on the individuation of perceptual systems that will permit us to break into this chain of interdefinitions.

The moral of all this is that to define the directly perceivable in terms of what perceptual systems respond to is merely to shift the locus of trivialization from the former notion to the latter. It puts the same pea under a different shell. 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

*non*transductive, hence presumably inferential. This is Gibson's characteristic dilemma, and we claim that he has no way out of it.

2.5. *The problem of misperception*

In much of the preceding discussion we have emphasized the undesirable consequences of interdefining "pick up", "invariant", "ecological property", and "directly perceive", but that is not the only difficulty with Gibson's approach. Part of an adequate theory of perception ought to be an account of perceptual *errors*, and it is hard to see how this requirement can be squared with the claim that perception is direct on *any* of the interpretations that Gibson's text suggests.

People who have tried to understand the nature of the mind, at least since Plato, have been particularly worried about the problem of false belief. In the present context, this is the problem of explaining how *misperception* is possible. The standard approach to this problem within Establishment theories is to connect misperception with failed inference. Your perception that something is edible, for example, is said to depend upon inferences from the appearance of the thing (e.g., from its smell, taste, texture, shape, color, and so forth). These inferences depend upon generalizations gleaned from past experience, and the generalizations are themselves nondemonstrative, and hence fallible. So, for these and other reasons, the (perceptual) inference from appearance to edibility sometimes goes wrong, with consequences that are typically painful and occasionally fatal.

Now consider how a noninferential story about misperception might go. Here we get a first glimpse of a dilemma that emerges, in various guises, throughout Gibson's text. If "directly perceive that" is taken to be factive, then by stipulation "x directly perceives that y is edible" will entail that y is edible. It follows that what goes on when one misperceives something as edible cannot be the direct perception of edibility. If, on the other hand, "directly perceive that" is *not* taken to be factive, then it is logically possible to, as it were, directly *misperceive* that something is edible. But Gibson will then need an account of what has gone wrong when misperception is direct. Notice, in particular, that he *cannot* account for such cases by saying that what you pick up when you directly misperceive the edibility of a thing is the property of *apparent* edibility. For, things that are misperceived to be edible *do* have the property of being *apparently* edible, and the problem for a theory of misperception is to explain how things could be taken to have properties that in fact they do *not* have. (A way out would be to say that you pick up the apparent edibility and *infer* the edibility from that; but this just *is* the Establishment way out and, of course, it is closed to Gibson.)

Probably the line that Gibson wants to take is that *if* an affordance is correctly perceived, *then* it is perceived directly; and that is, of course compatible with the factivity of "directly perceive". Notice, however, that such an approach does not help with the problem of misperception, since it does not tell us how we are to describe the cases where the antecedent of the hypothetical is *false*. We will return to this sort of difficulty. Suffice it at present to say that the problem of constraining "directly perceive" so as to provide a nonvacuous construal of the claim that perception is noninferential, and the problem of providing a coherent account of misperception without recourse to the notion of perceptual inference, are two sides of the same coin. No wonder Gibson is so unhappy about the role that appeals to illusions have played in the confirmation of Establishment theories of perception.

If a theory of perception is to be tenable it must not only address the most common (veridical) cases, but also the ones in which perception fails to be veridical and leads to false beliefs. The relative infrequency of the latter sorts of cases does not alter this principle (and, in fact, they are arguably not all that infrequent; only they tend to escape our notice except when the consequences are serious). Gibsonians sometimes urge that we should take very seriously the fact that perception works most of the time (see Reed and Jones, 1978), and it is true that this fact is of central importance for epistemology. But the goal of psychological theory construction is not to predict most (or even all) of the variance; it is to explicate the underlying mechanisms upon whose operation the variance depends. It seems quite inconceivable that the psychological mechanisms of perception and the psychological mechanisms of misperception are different *in kind*.

This problem is such a serious one that it sometimes drives Gibsonians to truly desperate recourses. For example, Turvey and Shaw (1979) suggest that we should cope with the issue of perceptual error by "tak(ing) perception out of the propositional domain in which it can be said to be either right or wrong... and relocat(ing) it in a nonpropositional domain in which the question of whether perception is right or wrong would be nonsensical". (p. 182). Apparently, this means either that we should stop thinking of perception as eventuating in beliefs, or that we should stop thinking of beliefs as having truth values. Turvey and Shaw describe this proposal as "radical", but "suicidal" might be the more appropriate term.

Perhaps the most characteristic Gibsonian move in this area is to identify misperception with failure to pick up 'all the relevant information'; (the bird flies into the window because it failed to pick up the ambient information that specifies *window*). But, of course, pick up of the very light structures which failed to specify *window* for the bird might be adequate to specify *window* for *us*. From a mentalistic point of view, this is not surprising; we

know a lot more about windows than birds do. So, the form of the problem for Gibson is to explain how pick up of the very same state of affairs that constitutes an adequate sample of information for one organism could constitute an inadequate sample for another. The Establishment account has an answer: *viz.* that what you perceive depends not only on the ambient information picked up, but also on the mental processes deployed in processing that information. It is far from clear what alternative the Gibsonian position could propose.

3. The problem of direct detection in establishment theories

Our argument thus far has been that unless the notions of pickup and invariant are constrained, it will always be trivially true that there is an invariant property whose pickup is necessary and sufficient for the perception of any object: *viz.* the property of being that object. We have also argued that some doctrines of Gibson's which can plausibly be construed as attempts to provide the relevant constraints do not succeed in doing so.

Though these considerations raise problems for Gibson's theory, it is important to understand that all other theories, including Establishment theories, have problems of a corresponding sort. This is because 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.[3]

To admit this is not, however, to endorse any "foundationalist" view of epistemology: to say that the pickup of some properties must be noninferential is not to say that our knowledge of these properties is infallible, or that the justification of perceptual beliefs depends upon assuming that the mechanisms of direct pickup are epistemologically privileged. Many philosophers have held that the deliverances of direct perception must figure centrally in

[3]There is, nevertheless, a sense in which all perceptual processes, strictly so called, might be inferential. Perception is usually taken to affect what the organism knows, and it is conceivable that transducer-detected properties are epistemically inaccessible to the organism and subserve no purposes except those of perceptual integration. (Cf. Stich's (1978) discussion of "subdoxastic" states.) In that case, these non-inferential processes are nonperceptual, as it were, by definition. In deference to this consideration, we have generally avoided talking of transduced properties as directly *perceived*, preferring the less tendentious "directly picked up". Of course, this terminological issue does not jeopardize the observation in the text that processes of perceptual inference must begin from premises that are not themselves inferred. The present question is just whether the noninferential processes of pickup which make such premises available should themselves be referred to as perceptual. (See also the discussion in section 7.4.)

the arguments which justify our perceptually based knowledge claims, but it is quite unnecessary to read this sort of moral from Establishment perceptual psychology.

The psychologist's topic is the causation of perceptual judgements, not the establishment of epistemic warrant in justificatory arguments. One can perfectly well hold—as in fact we are inclined to do—*both* that matters of epistemic warrant are typically determined by "inference to the best explanation" *and* that the causation of perceptual judgements typically involves inferences from premises which are not themselves inferred. The causal chain in perception typically flows "inward" from the detection of those properties to which peripheral transducers respond. But the flow of epistemic justification typically goes in every which way since the justification of perceptual knowledge claims is heavily constrained by principles of conservatism, parsimony, and coherence. In what follows, then, the epistemological issues will be put completely to one side: we make no assumptions about the epistemological role of whatever is directly detected[4]; for us, "direct" means only "noninferential".

One can distinguish at least two proposals that Establishment theories have made about how to draw the line between what is directly detected and what is inferentially mediated. On some views, especially the older, epistemologically based theories, the distinction between direct detection and inferential mediation is taken to be coextensive with the distinction between "sensory" properties and the rest. Typically, the sensory properties are characterized by introspective availability, and often enough it is assumed that the deliverances of introspection are infallible; hence the putative connection between perceptual immediacy and epistemic warrant that we noted in the preceding paragraph. Gibson holds, and we think that he is right about this, that the appeal to introspection will not do the job. In fact, as we saw when we discussed the "phenomenological" criterion for direct detection, what is introspectively accessible is typically not the traditional sensory properties (color, two-dimensional form, etc.) but rather "meaningful" properties like the affordances. When Gibson says that "phenomenological

[4]Some Gibsonians apparently want to read a sort of epistemological Realism as one of the morals of theories of direct perception (see, for example, Turvey (1977), but that would seem quite unjustifiable. On the one hand, *every* theory will have to acknowledge the fact of at least *some* misperception, and if one is going to run skeptical arguments in epistemology, that is the premise one needs to get them started (e.g., "If you admit that perception is sometimes fallible, what reason is there to suppose that it isn't always wrong?..." etc.). If you find such arguments persuasive, the idea that perception is direct *when it is veridical* will do nothing to soothe the skeptical itch, since that idea is compatible with the possibility that perception is *never* veridical. Correspondingly, an inference based theory of perception is perfectly compatible with a Realistic account of the information that perception delivers. All that is required for a perceptual inference to yield knowledge is that it should be sound. Gibson's views have philosophical implications, but not for epistemology.

objects are not built up of qualities; it is the other way around" (p. 134) he is quite right about the deliverances of introspection. Since, however, traditional theorizing is precisely concerned to treat properties that are on the level of the affordances as *inferred*, it very early abandoned the identification of what is directly detected with what is introspectively available. If, however, the sensory properties are *not* identifiable with the ones that are introspectively available, it does not help much to say that sensory properties are what we detect directly, the former notion being as unclear as the latter.

Recent versions of the Establishment theory have sought to constrain the notion of direct detection by identifying the properties that are available without inferential mediation with those to which transducer mechanisms are sensitive. This transfers the problem of constraining "directly detectible property" to the problem of constraining "mechanism of transduction" and, contrary to assumptions that appear to be widely made, specifying what is allowed to count as a transducer for the purposes of cognitive theory is a nontrivial problem. For example, transducers are technically defined as mechanisms which convert information from one physical form to another. But this definition is entirely compatible with there being transducers for *any* pattern of stimulation to which the organism can respond selectively since *whole organisms* are, in that sense, transducers for any category to which they can reliably assign things; e.g., for sentences, or shoes, or, in Berenson's case, for Da Vincis. This is precisely Gibson's problem as it arises in the context of Establishment theories, and to fail to grasp its seriousness is to fail to understand the challenge that Gibson poses to the Establishment. The theory that perception is typically direct is empty barring an independent construal of pickup; *but so too is the theory that perception is typically inferential.* On the other hand, it should be borne in mind that the Establishment does not accept Gibson's condition on the solution of this problem; *viz.* that the objects of direct detection (transduction) must be so specified that no perceptual judgements turn out to be inferentially mediated. We think that Gibson's position is hopeless precisely because pickup can be constrained only if that condition is abandoned.

Some theorists in the Establishment tradition hold that the way to decide what transducers we have is by appealing to neurophysiology—for example, by finding out what biological mechanisms serve to convert ambient stimulation into the electrical energy of nerve impulses. There are, however, several difficulties with this sort of approach. In the first place, it fails to rule out the whole nervous system as a transducer since, after all, converting ambient energies into neural excitations is a good part of what the nervous system does. Moreover, the class of mechanisms that would count as transducers by this criterion involves many which perform no function that is of significance

for the theory of perception. This is because not all stimulus events that affect the character of nerve impulses are *ipso facto* available for the causation of perceptual judgements. Uttal (1967) refers to those neural events that are functionally relevant as *signals* and those that are not as *signs*, precisely in order to emphasize this distinction. This consideration suggests that the identification of transducers will have to advert not, in the first instance, to their neurological structure but to their role in the cognitive processes that they subserve. Like Gibson, we assume that the individuation of perceptual mechanisms is primarily functional rather than physiological.

Finally, it might be argued that whether a device (including a neurophysiological mechanism) counts as a transducer depends, at least in part, on its psychophysical characteristics; on the way that its output depends upon its input. As will become clear, we think that some proposal of this general kind is probably correct. Notice, however, that it does not follow that the sort of evidence that is collected in standard psychophysical experiments will resolve the issue. This is because such evidence does not, in the general case, directly reflect the behavior of isolated components of the perceptual system. Psychophysical curves reflect patterns of judgements produced by *whole organisms*, and are typically affected not only by stimulus parameters, but by the utilities, expectations, and beliefs that the organism entertains.

We will assume, in what follows, the identification of what is "picked up" with those properties that transducers respond to. Our problem will thus be to find some satisfactory alternative to the ways of constraining transduction that we have just discussed.

4. The first constraint on pickup: What is picked up in (visual) perception is certain properties of the ambient light

We begin by considering a fundamental construct in Gibson's theory, the notion that states of affairs can *contain information about* one another. The basic idea is that the state of affairs S1 contains information about the state of affairs S2 if and only if S1 and S2 have correlated properties. Suppose that S1 consists of a's having property F and S2 consists of b's having property G. Then if, in general, x's having property F is correlated with y's having property G, then S1 contains information about S2.

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. In particular, information-cum-correlation is not something that is encoded, or transmitted, or stored; though it is, according to Gibson, "picked up" whenever anything is perceived.

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 S1 specifies S2, S1 and S2 are correlated,[5] and S1 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. Features of the light are correlated with features of the layout in virtue of the regularities expressed by laws of ecological optics. The structure of the light therefore contains information about the character of the layout; and, since organisms actually use that information in the perceptual identification of layout features, the structure of the light is said to specify the character of the layout.

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—barring, perhaps, the occasional ecological optician—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. Since the two are correlated you could, in principle, determine the structure of the light given the facts about the layout (and about the correlations) just as you can, in principle, determine the structure of the layout given the facts about the light (and about the correlations). And this raises a problem, though not one that Gibson discusses in these terms: *viz. what determines the direction of specification?*

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. If we had transducers for layout and no transducers for light, then any specification relation that held between the two would have to go in the opposite direction. The moral is: if we are in a position to say what the direction of specification is for a given organism, then that fact

[5]Strictly, S1 and S2 are tokens of correlated types. We will not be explicit about the type token relation except where the intention is not clear from context.

constrains our attribution of transducer mechanisms to the organism. The attribution of transducers must serve (*inter alia*) to explain the facts about the direction of specification for the organism.

So we have a constraint on transduction. But how is this constraint to be applied? In particular, how do you tell which sorts of states of affairs serve as specifiers for a given organism? Given correlated states of affairs, how do you tell which specifies which? The answer is sufficiently obvious. What you do is, you break the correlation experimentally (you set up a case in which the correlation fails) and then you see what happens.[6]

Consider the following simple examples. How do we know that the light specifies the layout and not *vice versa*? Well, we can create paired situations in one of which we preserve the features of the light without the corresponding layout, and in the other of which we preserve the features of the layout without the corresponding light. The presentation of a hologram would be an example of the first kind; turning out the lights would be an example of the second kind. There is no dispute about what would happen in such experiments. You can vary the layout as much as you like; so long as the properties of the light are unaffected, the perceptual judgements of the organism are unaffected too. On the other hand, leaving the layout intact does you no good if the structure of the light is changed. In the extreme case, take the light away, and the organism cannot see.

In short, the way you determine which of a pair of correlated states of affairs specifies the other is by applying the "method of differences," in which one determines which of two factors is the cause of some effect by setting up a situation in which only one of the factors is operative. In the present case, we have a pair of correlated states of affairs and a perceptual judgement in which they eventuate. We assume that the light contains information about the layout, but we have still to show that the information in the light serves to *specify* the layout; *viz.* that the perception of layout features is causally dependent upon the detection of the information in the light. The hypothesis that the light does specify the layout implies two predictions corresponding to the two ways of breaking the correlation between light features and layout features: since the detection of the light is causally *necessary* for the (visual) perception of the layout, we predict that the organism sees nothing in the layout-without-light setup. Since the

[6] Of course, knowing the physical/physiological structure of the organism can provide some constraints upon the assignment of transducers, since if there is *no* mechanism that is differentially sensitive to a given form of input energy, then that form of input *cannot* be a specifier for that organism. However, as we remarked above, this consideration does less than might be supposed since, in the general case, practically any form of ambient energy is likely to have *some* effect on the organism's neurological condition, and it is functional considerations which must decide which such effects are to count as transductions.

detection of the light is causally *sufficient* for the perception of the layout, we predict *layout illusions* in the light-without-layout setup.

It is the latter consideration which accounts for the centrality, in perceptual psychology, of experiments which turn on the creation of perceptual illusions. An illusion is simply a case in which the specifying state of affairs is brought about without the occurrence of the correlated state of affairs that it normally serves to specify. To produce an illusion is thus to demonstrate a direction of specification. It is characteristic of Gibson's break with the tradition that he disapproves of psychological theories which appeal to perceptual illusions as data, Gibson's point being that the laboratory illusion is an 'ecologically invalid' happening. So it is—by definition—since, as we have seen, you construct an illusion precisely by breaking a correlation that holds *in rerum natura*. Our point is, however, that the theoretical pertinence of facts about illusions is an immediate consequence of taking the specification relation seriously. If saying that S1 specifies S2 implies that the perception of S2 is causally dependent upon the detection of S1, and if causal dependence implies causal sufficiency, then one is committed by the logic of the case to the prediction that S1 presentations can engender S2 illusions. It is notable that Gibson himself (tacitly) accepts this form of argument. When he cites evidence in support of particular empirical claims regarding the identity of specifying stimuli, he frequently appeals to the standard kinds of experimental data about illusions; e.g., cases where one can produce illusions of motion by providing subjects with simulations of optical flow patterns. It seems that some illusions are ecologically more valid than others.

The state of the argument is now as follows: when S1 specifies S2, the perception of S2 is causally dependent upon the detection of S1. Since the direction of specification is determined by the transductive capacities of the organism, it follows that S1 specifies S2 only if the organism has transducers for S1. The notion that the facts about transduction determine the direction of specification thus serves simultaneously to constrain the notion 'object of detection' (only specifiers are directly detected) and the notion 'mechanism of transduction' (only mechanisms which respond to specifiers are transducers). The method of differences gives us a way out of the threatened interdefinition of 'transducer' with 'object of direct detection' since we have *empirical* tests for whether a stimulus is a specifier.

Here, then, is the proposal in a nutshell. 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.[7]

[7] See overleaf.

It is, of course, condition (b) that does the work. For, if a state of a system is correlated with a property, then it will typically also be correlated with any property with which that property correlates. Specifically, 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. However, only the former correlation will be counterfactual supporting in the required way; visual transducers are unaffected by manipulation of the layout unless the manipulations affect the properties of the light. Hence, by our criterion, *only* properties of the light are transduced in visual perception.

Another way of stating this condition is to say that a system which is functioning as a detector of P is in a certain sense illusion-free with respect to P. This is not, however, because detection is, in some puzzling way, infallible; it is only because, by assumption, the validity of P-*perception* depends upon situational correlations in a way that the validity of P-*detection*, by assumption, does not. To say that a property is detected is to say that the property would continue to have its psychological effect in circumstances in which correlated properties were suppressed. But P-illusions are possible only where the perception of P is mediated by the detection of one of its correlates, the illusion 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.

We have seen that the counterfactual-support condition on transducers has the consequence that only properties of the light are transduced in visual perception. It should be emphasized, however, that not *all* properties of the light can be so transduced if that condition is to be honored. Consider, for example, the (relational) property that the light has if and only if it is caused by the layout being arranged in a certain way. This is a perfectly good property of the light, but it is not one that can be directly detected according to the present view. For, this property has its effect on perception only *via* the effects of such correlated light features as wavelength, intensity, color discontinuities, etc. That is, the perceptual effects of the former property are preserved only in those circumstances in which the latter properties are detected. (We make this claim on empirical rather than *a priori* grounds; we

[7] It is of prime importance that the employments of the method of differences should be *relevant* since, of course, there are *some* counterfactual conditions in which P will *not* produce S_t even if S *is* a transducer: e.g., the universe blows up, the organism dies, and so forth. The counterfactual supporting generalizations about transducers are thus like most counterfactual supporting generalizations in science in that they must be relativized to assumptions of 'normal background conditions'. Perhaps only the fundamental laws of microphysics are exempt from such relativization, these being assumed to hold, literally without exception, for all segments of space-time.

assume that it is what the relevant employments of the method of differences would show to be true.) The property of having-been-caused-by-such-and-such-a-layout-feature is thus a property that the light may have, but it is not a detectable property of the light.

Because the counterfactual support condition is not satisfiable by such properties, the illusion freedom condition is not either. It will always be possible, at least in principle, to construct minimal pairs of light arrays such that one of them has the property and the other does not; and the organism will be unable to distinguish between such pairs within the limits of the experimental procedure. That is what happens when we construct an object that looks like a shoe but isn't one; if it structures the light in a way sufficiently like the way that a shoe does, the subject cannot tell by looking that the light structure lacks the property of having been caused by a shoe. Similarly, *mutatis mutandis*, when one fakes a Da Vinci. So, then, on the one hand, nothing but the properties of the light can be directly detected in visual perception; and, on the other hand, there are (infinitely) many properties of light that cannot be so detected.

We shall presently return to the bearing of all this upon the main question of whether perception ought to be considered to be an inferential process. First, however, it may be worth considering some further implications of the counterfactual-support condition. We believe that the tacit acceptance of this condition upon detection explains a number of intuitions theorists have had concerning what can count as a transducer.

For example, it is frequently assumed that detectors are sensitive only to *physical* properties (i.e., to such properties of states of affairs as can be expressed in the vocabulary of the physical sciences). On this view, we could, in principle, have detectors for wavelength, intensity, pressure, or even chemical composition, but not, say, for being expensive, being nutritious, being causally related to some past event (e.g., being a genuine Da Vinci), or being a sentence of English. We suggest that these intuitions about which properties are transducible are shaped by the theorist's implicit allegiance to the counterfactual support condition *via* the following considerations.

It is usually assumed that the only empirical generalizations which support counterfactuals are laws. This is practically tautological since a law just is a generalization that holds in all physically possible worlds in which the relevant background conditions are satisfied; i.e., across all relevant employments of the method of differences. Suppose that this assumption is correct. Then, since generalizations which specify the relation between detector output states and detected properties must be counterfactual supporting, it follows that such generalizations must be lawful. However, as we have seen, the vocabulary of laws is restricted to predicates which express projectible

properties. In short, then, the following theoretical decisions ought to go together: (a) the decision as to whether a property is detectible; (b) the decision as to whether the property is projectible; (c) the decision as to whether a generalization which involves that property is a law; (d) the decision as to whether the generalization is counterfactual-supporting; (e) the decision as to whether a mechanism which is sensitive to the property can count as a detector for that property.

Now, many theorists have held, more or less explicitly, that the only laws there are are laws of the physical sciences, hence that the only properties that can be subsumed by counterfactual supporting generalizations are physical properties. If you believe this, then given the considerations just reviewed, you ought also to hold that there can be detectors only for physical magnitudes. And, whether or not you believe that all laws are laws of physics, there is presumably nobody who believes that there are laws about, say, grandmothers *qua* grandmothers or about genuine Da Vincis *qua* genuine Da Vincis, though there may, of course, be laws about coextensive kinds. (Remember that coextensive properties may nevertheless differ in projectibility.)[8] The suggestion is that the intuition that there are no laws about the property *grandmother* is what explains the intuition that there cannot be grandmother detectors. The moral is: the decision about what detectors there are is linked to the decision about what laws there are. A world in which there were laws about the property *shoe* would be a world in which there could be detectors for shoes. After all, a law about *shoe* would, presumably, connect the shoe property to other sorts of properties, and then things which have properties of these other sorts would *ipso facto* be available for service as shoe detectors.

In the light of these considerations, we can now understand at least one of the moves that Gibson makes. The fact that Gibson holds *both* that there is detection of ecological parameters and that there are *laws* of ecological optics are seen to be linked decisions. If you hold that nonphysical parameters can be detected, and if, by definition, the states of detectors are lawfully connected with the properties they detect, then you must also hold that there are laws which involve nonphysical magnitudes. In this respect, at least, Gibson's doctrines are mutually consistent.

---

[8] For our purposes, a world in which there were laws about grandmothers would be one in which some effect is a consequence of something being a grandmother, regardless of what other properties it may have. But, surely, this is not true in our world. Suppose, for example, that there are true empirical generalizations of the form $(\forall x)(\exists y)$ (x is a grandmother $\rightarrow$ Fy). Then it seems enormously plausible that such a generalization holds only because there is some property H *other than being a grandmother*, such that the generalization $(\forall x)(\exists y)(Hx \rightarrow Fy)$ is true; and moreover that it is the latter generalization which supports counterfactuals in the critical cases. That is, if *a* were a grandmother but *Ha* was false, then the former generalization would not hold for $x = a$.

5. The "information in the light"

The main point of our discussion was to establish some conditions on the notion *detection* (transduction). We needed to do this because we doubted that the notion *could* be appropriately constrained consonant with the doctrine that perception is, in the general case, not inferentially mediated. We are now in a position to see one of the ways in which the conflict arises; indeed, one of the respects in which the Gibsonian model of visual perception is after all committed to inferential mediation, just as Establishment models are.

The first point to notice is that Gibson actually agrees with much of what we have been saying, although the terminology he employs sometimes obscures the consensus. Gibson makes a distinction (largely implicit, and not invariably honored) between what he describes as "directly perceived" and what he describes as "picked up". The latter locution is usually reserved for features of the light, while the former is usually used for features of the layout. Moreover, Gibson seems to agree that picking up features of the light is causally necessary for "directly perceiving" features of the layout. Notice that, in this respect, Gibson's view is simply indistinguishable from the Establishment theory. 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. Thus far, the differences are merely terminological. The important fact is the agreement 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 that the former relation is causally dependent upon the latter.

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.* The story is that you infer the latter from the former 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. It is all very well to call your epistemic relation to layout features "direct perception", but if it is agreed that that relation is dependent upon an epistemic relation to properties of the light, "direct" certainly cannot be taken to mean "unmediated". 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. In the normal case, states of mind like the first are causally necessary to bring about states of mind like the second (and they are normally causally sufficient for organisms which have had appropriate experience of the ways in which light

states and layout states are correlated). 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. The point is that Gibson has done nothing to avoid the need to postulate such a process; it arises as soon as "direct detection" is appropriately constrained. And he has suggested no alternative to the proposal that the process comes down to one of drawing perceptual inferences from transducer outputs; in the present state of the art that proposal is, literally, the only one in the field.

What obscures this problem in Gibson's presentation is that, instead of speaking of picking up properties of the light, he talks about picking up *the information about the layout* that the light contains. This certainly appears to be an alternative to the Establishment idea that layout features are inferred from light features. But, in fact, if one bears in mind the character of the theory of information that Gibson has actually provided, one sees that the appearance is illusory. Remember that "information" is a defined construct for Gibson; S1 contains information about S2 if, and only if, they are correlated states of affairs. The problem is that while Gibson gives no hint of any notion of information other than this one, it is hard to see how this account can sustain talk of information pickup.

Given that "contains information about" just means "is correlated with", what could it mean to say that an organism picks up the information that S1 contains about S2? The obvious suggestion is that you pick up some property of S1 that you know to be correlated with some property of S2, and you use your knowledge of the correlation to infer from the former property to the latter. But this cannot be what Gibson has in mind, since this is just the Establishment picture; we learn about the layout by inference from the detected properties of the light. That is, what we detect is not the information *in* S1 but rather the informative properties *of* S1. Then what we learn about S2 in consequence of having detected these informative properties depends upon which inferences we draw from their presence.

Perhaps, then, Gibson's idea is that detecting the information that S1 contains about S2 is detecting the correlation between S1 and S2. But a moment's thought shows that this cannot be right either. To say (loosely) that S1 is correlated with S2 is to say that S1 and S2 belong to correlated *types* of states of affairs (see footnote 5). But, surely, you find out about correlations between types of states not by "detecting" the correlation but by processes of nondemonstrative (e.g., inductive) inference.

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").

This is such an easy slide that it is essential to bear in mind that Gibson has no notion of information that warrants it. Information, in Gibson's sense, is not the sort of thing that can affect states of perceptual systems. What *can* function causally is *informative properties* of the medium, properties of the medium which are *de facto* informative because they are correlated with properties of the layout. So, for example, the frequency of the light can cause a state of a detector, and the frequency of the light can be *de facto* informative about the color of reflecting surfaces in virtue of a correlation that holds between frequency and color. But the fact that the frequency of the light is correlated with the color of reflecting surfaces cannot itself cause a state of a detector, and appeal to that fact exhausts Gibson's construal of the notion that the light contains information about the color of surfaces. So we are back in the old problem: 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? The function of the Establishment notion of perceptual inference is, of course, precisely to answer this question.

In short, "picking up the information in the light" must, given Gibson's account of information, come down to picking up features of the light that are correlated with features of the layout. Since the correlation is empirical (*via* the laws of ecological optics), 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. In this case, picking up the information in the light will *not* lead to perceptual knowledge of the environment. Since this can happen (and does in innumerably many cases; see the discussion of the bird and the window in section 2.5), the theorist must face the question: what more than the pickup of *de facto* informative medium properties is required to mediate perceptual knowledge? The notion of inference may provide an answer; but, in any event, Gibson's notion of information does not. Information explains only what correlation explains, and the existence of a correlation between two states of affairs does not, in and of itself, explain how the detection of one of them could eventuate in perceptual knowledge of the other.

There is, we think, a deeper way of putting these points, and it is one that we will return to in the last section of our discussion. The fundamental difficulty for Gibson is that "about" (as in "information *about* the layout in the light") is a semantic relation, and Gibson has no account *at all* of what it is to recognize a semantic relation. The reason this is so serious for Gibson is that it seems plausible that recognizing X to be about Y is a matter of mentally representing X in a certain way; e.g., as a premise in an inference from X to Y. And it is, of course, precisely the notion of mental representation that Gibson wants very much to do without. We have here a glimmer of Gibson's ultimate dilemma: the (correlational) notion of information that he allows himself simply will not serve for the identification of perception with information pickup. Whereas, the semantic notion of information that Gibson needs depends, so far as anyone knows, on precisely the mental representation construct that he deplores. The *point* of the inferential account of perception is to spell out what is involved in taking proximal (or ambient) stimulation as containing information about its distal causes. One cannot provide an alternative to that theory merely by assuming the notion of information as unexplicated, though that is, to all intents and purposes, just what Gibson does.

To summarize: Gibson has no notion of information over and above the notion of correlation. You can, no doubt, pick up properties of S1, and, no doubt, some of the properties of S1 that you can pick up may be correlated with properties of S2. But you cannot pick up the property of *being correlated with S2*, and it is hard to see how the mere existence of such a correlation could have epistemic consequences unless the correlation is mentally represented, e.g., as a premise in a perceptual inference. We can put it in a nutshell: sensible constraints on visual direct detection make properties of light its natural object. And then the question "how do you get from an epistemic relation to properties of the light (*viz.* pickup) to an epistemic relation to properties of the layout (*viz.* perception)?" seems to have only one conceivable answer: by inferential mediation, like the Establishment says.

The moral of all this is that when Gibson says that we perceive the layout "directly", one must not take him to be claiming that the perception of the layout is unmediated. Gibson, in fact, accepts that visual perception of the layout is mediated at least by the detection of properties of the light, and we have argued that he has suggested no alternative to the idea that such mediation also involves inference. Thus, if we want to find a disagreement between Gibson and the Establishment, we shall have to look to something other than the question whether the perception of distal visual layout involves inference from proximal visual stimulations; both sides agree that it does, albeit with unequal explicitness.

6. The second constraint on pickup: Only properties of "effective stimuli" are directly detected

Since when Gibson says that perception is "direct" he is clearly not saying that it is unmediated, the question arises what alternative construal might be placed upon his claim. The following suggestion seems to be compatible with much of the text: Although perception of the layout is causally dependent upon pickup of properties of the medium, still the information about the layout that the medium makes available is so rich that the pickup of that information is, as it were, *tantamount* to the perception of the correlated layout properties. To all intents and purposes, this comes to the claim that a given configuration of the medium (e.g., of the ambient optical array) specifies a corresponding configuration of the layout uniquely.

There is a stronger and a weaker version of this claim. The stronger version is that (so long as we focus on the right properties of the medium) the information we find there is, under normal circumstances, *almost invariably* sufficient to specify the ecologically relevant properties of the layout. The weaker claim is that, although *some* of the perceptually available properties of the layout are uniquely specified by properties of the medium, it is left open that other such properties may not be. Gibson's exposition makes it clear that he intends the former of these claims, as indeed he must if there is to be a difference between his views and those of the Establishment.

The Establishment theory takes the connection between the distal layout and the states of the transducers to be something like this: certain properties of the medium are causal determinants of the output of the detectors. Some of these medium properties are, in turn, causally determined by properties of the distal layout. Since the relation of causal determination is transitive, the detector output is itself normally contingent upon particular features of the layout. If, as has usually been assumed, this relation between layout properties and detector outputs is more or less one-to-one (i.e., the mapping is reversible), then this view is entirely compatible with the weaker version of Gibson's claim. In both the Gibsonian account and the Establishment view, it is part of the explanation of the veridicality of perception that, in ecologically normal circumstances, many of the directly detectable properties of the light are specific to properties of the layout which cause them.

In short, the weak version of Gibson's claim is that there are *some* visual properties of the layout which are, to a first approximation, causally necessary and sufficient for properties of the light, which latter properties are themselves directly picked up. Our point has been that the Establishment theories say that too; in particular, the Establishment theories provide precisely that account in the case of the sensory properties of the layout, taking

these to be, by stipulation, the properties of the layout which are causally responsible for the properties of the medium that transducers respond to.

Since the weak version of Gibson's claim does not distinguish his position from the Establishment's, let us consider the stronger version, which is that the light contains information that is specific to just about all the visual properties of the layout. By contrast, according to Establishment theories, the sensory properties are a very small subset of the visual properties, and it is *only* for the sensory properties that the medium-to-layout mapping is assumed to remotely approach uniqueness. Hence, the Establishment theory is *not* compatible with the strong version of Gibson's claim. Given this incompatibility, the next question is whether Gibson's claim is plausible on its strong construal. It turns out, however, that before we can raise this question, we have to face yet another trivialization problem.

Consider the claim that, for each visually perceptible property of the layout, there is a corresponding property of the light which is, in some non-vacuous sense, directly detectable and specific to the layout feature in question. For the moment, let us not worry about how the notion of a directly detectable property is to be constrained, and concentrate instead on the issue of specificity. Once again there is a way of trivializing Gibson's claim. The trivializing alternative arises if, among relevant properties of the medium, we allow properties of *arbitrary spatio-temporal cross sections of the light* (for example, the distribution of the light across the entire inhabited universe throughout some arbitrary segment of history, including the arrangement of the light reflected from all the pages of all the books in all the libraries, and all the dials on all the apparatus in all the scientific laboratories). If we take such an arbitrarily bounded sample of the light, then it may well be that its structure does uniquely specify every perceptible property of the corresponding layout. Indeed, it may be that the arrangement of all the light specifies a unique layout of all the objects, perceptual or otherwise. Whether this is true may be of considerable epistemological interest, since an epistemologist might well wonder whether there is enough data in the medium to determine a unique best theory of the world. The trouble is that, either way, the issue has no implications for the psychology of perception.

The important psychological question is whether the claim of specificity can be maintained for *appropriately bounded* samples of the ambient optic array, and the interest of the Establishment contention that perception is typically inferential depends in large part on the claim that the answer to this question is "no". The Establishment holds that there must be inference from medium to layout; but, as we have seen, that must be admitted by anyone who accepts the principle that only properties of the medium are detected directly. What is more contentious is the Establishment claim that

layout properties are typically inferred on the basis of *relatively fragmentary* information about the structure of the medium; hence that the patterns of transducer outputs which serve as "premises" for perceptual inferences in general significantly *underdetermine* the percepts to which they give rise.[9] This is an issue on which Gibson and the Establishment certainly disagree. Our point has been that before it can be assessed, we need some independent criterion for what is to *count* as an appropriately bounded sample of the optic array.

We assume that the following is—or, anyhow, ought to be—untendentious: The goal of a theory of perception includes characterizing the sample of the ambient array which causes each percept. It is true that, in his most recent writings, Gibson sometimes seems to be saying that perception should not be viewed as caused by stimulation. Indeed, he appears to want to do away with the notion of the stimulus altogether. "I thought I had discovered that there were stimuli for perception in much the same way that there were known to be stimuli for sensations. This now seems to me to be a mistake. I failed to distinguish between stimulation proper and stimulus information." (p. 149). But, whatever this distinction may come to, it surely does not provide an argument against there being environmental causes of perception. And, as long as it is assumed that there are, giving a causal account of perceptual phenomena is surely one of the central aims that psychology ought to pursue. In particular, what we want is a specification of the sample of the ambient array which causes each distinct perceptual episode.

Just as the goal of specifying the environmental causes of percepts survives disagreements over whether what causes a percept is stimulation or stimulus information, so it also survives disagreements over how percepts ought to be described. Gibson says: "I should not have implied that a percept was an automatic response to a stimulus, as a sense impression is supposed to be. For even then I realized that perceiving is an act, not a response..." (Ibid). An adequate psychology might provide mappings from segments of the ambient array onto percepts, or onto some larger events, or patterns of behavior; in either case, perceptual episodes will be viewed as caused and the problem of specifying bounds on the causally efficacious sample of the ambient array will have to be faced.

[9] For purposes of this discussion, we will usually speak of the epistemic states arising from perception as *percepts*. One could equally talk of perceptual beliefs, perceptual judgements, or any other epistemic state that an organism is in as a logical consequence of having perceived—as opposed to having guessed, deduced, remembered or otherwise concluded—that P.

Talk of percepts, as opposed to beliefs about the world that do not arise directly from perception, implies a distinction between perceptual and cognitive processes (Dretske, 1978). Empirical grounds for drawing this distinction are discussed in, for example, Hochberg (1968). The present point is that all theories have to draw it somewhere, and the question about the richness of the ambient array arises specifically for the causation of perception, however it may be defined.

In fact most of these worries are, in the invidious sense, academic. Gibson and the Establishment agree on what constitutes some clear cases of perceptual phenomena to be explained. When one performs an experiment by setting up certain displays and finds that subjects report seeing certain things, this is *prima facie* a relevant datum for perceptual theory. For example, Gibson (p. 190) makes much of an experiment by Kaplan involving progressive deletion or accretion of a random texture. He reports: "What observers saw was an edge, a cut edge, the edge of a sheet, and another surface behind it." This, then, is a perceptual phenomenon which everyone agrees requires causal explanation; and this agreement presupposes no general consensus about the ontological status of percepts.

On any account, then, percepts have causes, and among the causes of a percept will be some bounded spatio-temporal segments of the ambient optical array. Let us call such a segment the *effective stimulus* for the percept that it causes. Thus for every percept there is some effective stimulus which is its immediate environmental cause. Given this notion we can now ask the critical question: Is it true, in the general case, that each effective stimulus is uniquely correlated with the structure of a corresponding layout? We take it that this is the appropriate way to ask the question whether, in the general case, the structure of the medium specifies the structure of the layout uniquely.

When, however, we put the question this way, it seems obvious that the answer is "no". The mapping of layouts onto effective stimuli is certainly many-to-one, for it has been repeatedly shown in psychological laboratories that percepts can be caused by samples of the ambient medium which demonstrably underdetermine the corresponding layout. Nor is this phenomenon specific to vision. Consider, for example, the "phoneme restoration effect" (Warren, 1970) in psycholinguistics: Take a tape recording of an English word, and delete the part of the tape corresponding to one of the speech sounds. (For example, one can start with a recording of "determine" and produce a recording of "de#ermine".) Now record a cough-sound and splice it into the gap. The resulting tape is heard as containing the ungapped original word ("determine") with a cough-sound "in the background". The experiment thus demonstrates that an acoustic array which serves as an effective stimulus for the perception of a cough when heard in nonspeech contexts, can also serve as an effective stimulus for the percept /t/ when heard in the context "de#ermine", for the percept /k/ when heard in the context "es#ape", etc. The mapping from effective stimuli onto layouts is thus one-many in at least one case.[10]

[10] See facing page.

Gibson is, of course, aware of such results, but he deprecates them on the grounds that providing a *richer* sample of the ambient array is often sufficient to change the organism's perception of the layout; as, for example, when one destroys the Ames' room illusion by allowing the viewer to move freely through the experimental environment or when, in the case of the phoneme restoration effect, one slows down the tape enough to hear what is "really" going on. (The latter case shows, by the way, that the "richer" stimulus—the one which leads to true perceptual beliefs—is by no means always the ecologically normal stimulus; there are illusions which occur in the normal situation, and in these cases it is the ecologically pathological stimulus which is required to produce the veridical percept). But Gibson's criticism of these results is irrelevant once one accepts the condition, enunciated above, that an adequate theory must account for the effects of *all* perceptually effective stimulations. True, we can alter the initial percept by adding to the input (supplying context); but it remains to be explained how the original "ecologically invalid" percept was caused. In effect, Gibson's criticism is telling only if one accepts the trivializing construal of his claim that the medium contains information sufficient to specify the layout, thereby avoiding the serious issue which is how much information the *effective stimulus* contains.

If, by contrast, we take the effective stimulus constraint seriously, the facts seem to be clear: percepts are often caused by effective stimulation which is not specific to a layout. In such cases, the properties of the medium that are picked up underdetermine the layout that is perceived. So we are in need of an answer to the question what processes *other than* the pickup of medium properties are implicated in the causation of percepts? The Establishment theory has an answer; *viz.* the occurrence of certain perceptual inferences. In particular, inferences *from* the detected properties of the fragmentary stimulus *to* the properties that a richer sample of the ambient array would reveal. Gibson has, thus far, provided no reason for rejecting that answer, nor has he shown how an alternative might be formulated.

---

[10] It might occur to a Gibsonian to avoid this conclusion by reanalyzing the effective stimuli. Whereas we assumed that the effective stimulus was *cough* (which, occurring in isolation is heard as a cough but occurring in speech-context is heard as a phone) a Gibsonian might want to argue that the isolated stimulus is actually *#silence-cough-silence#*. Since that stimulus is never presented in the speech condition, the appearance of a one-many stimulus-to-layout mapping is dissipated. This would be a typically Gibsonian tactic of appealing to context to avoid the problem of ambiguity.

The disadvantages of the tactic are, however, clearly revealed in this case. For, the stimulus #noise-*cough*-noise# *will*, in general, be heard as containing a cough; and this fact is rendered a mystery on the assumption that the right way to describe the effective stimulus for a cough perceived in isolation is as #silence-cough-silence#. Quite generally, what you gain vis à vis ambiguity by enlarging the effective stimulus, you lose vis à vis the perception of similarity. This is because the perception of similarity is so often mediated by the recognition of partial identity of the internal structures of the stimuli. See below, section 7.2.

To summarize: the claim that there is enough information in each sample of the light to specify a unique layout is empty without some constraint on what is to count as a sample. Gibson provides no such constraint, but it is fairly clear how one ought to do so: since the goal is a theory of the causation of percepts, the appropriate sample must be what we called the "effective stimulus". For, by stipulation, the effective stimulus just *is* an arrangement of the medium that is sufficient to cause a percept. But then the claim that, in the general case, effective stimuli uniquely specify layouts is patently false on empirical grounds. Note, finally, that the fact that the perception of the layout is generally veridical does *not* require that effective stimuli specify uniquely; indeed, the inferential account of perception is precisely an attempt to show how veridical perception could occur without unique specification. Since perception depends on ambient stimulation *together with inference*, the veridicality of perception requires only that the principles of inference should be truth preserving most of the time.

7. What properties of the effective stimulus are directly detected?

In view of the preceding discussion, one might well ask what remains of the ecological approach to visual perception. We will now see that Gibson can be construed as making a number of plausible objections to standard Establishment assumptions about which properties of the light are directly detected. Recall that we started by recognizing that the claim of immediate perception was vacuous unless constrained. We suggested that one such constraint is the requirement that generalizations which relate the inputs and outputs of detectors should be counterfactual supporting (i.e. that they should survive appropriate applications of the method of differences). A second constraint was that immediately detectable properties of the light should be properties of effective stimuli (i.e. of light samples which cause percepts). The present point is that within these constraints there is a real empirical issue as to what the correct inventory of detected properties is, and here Gibson departs in important ways from the assumptions that Establishment theorists have often made.

The version of Establishment theory against which Gibson typically pits his approach takes instantaneous point-intensities of the light impinging on a retinal surface as the only properties that are primitively detected in vision. Gibson's point is, at a minimum, that this decision cannot be defended on *a priori* grounds and that describing the directly detected stimulus as an instantaneous mosaic has profound and implausible implications for the rest of one's theory of perception. Both these points are well taken, and it is

worth emphasizing that the pointillist view of the directly detected stimulus is not entailed by the constraints on transduction that we have proposed. There is nothing in these criteria to rule out there being directly detected properties of light distributed over much longer segments of space-time than has been widely assumed by Establishment theorists. Nor, as Gibson points out, do the physiological facts about the size of retinal receptive fields imply the pointillist view since, as we have repeatedly remarked, decisions as to what counts as a transducer—and hence decisions as to what properties are transduced—must ultimately be made on functional grounds.

The assumption of highly local detectors has often lead Establishment theorists to conclude that the perception of patterns over space and time *necessarily* involves construction and inference from information stored in memory. This conclusion is one of Gibson's favorite targets (see, for example, Gibson 1966). He maintains that memory is not needed if we allow features of larger segments of space-time to be detected. And, strictly speaking, this is correct so long as the effective stimulus constraint is observed. There would, for example, be no need to posit the construction of such things as texture gradients and flows from snapshot memories if the organism could detect spatially and temporally extended light patterns. In the spatial case, such detection could proceed by the use of devices like spatial frequency filters and templates. Analogous methods are available in the temporal domain where tuned filters can be made to play the role of templates.

Similarly, as Gibson has frequently pointed out, it is incorrect to assume that the only way that one can perceive change is by detecting and comparing two instantaneous states, at least one of which is retained in memory. In the first place, there is nothing to prevent such mechanisms as speedometers, accelerometers or frequency meters from being primitive in the required sense. Yet all of these can be viewed as detecting change: A speedometer detects rate of change of distance, an accelerometer detects rate of change of velocity, and so on. In the second place, as the preceding examples suggest, whether one is detecting a given magnitude or a temporal derivative of some other magnitude is sometimes a matter of how one chooses to describe things. A transducer for acceleration (say in the vestibular system) can be viewed as detecting a magnitude (a force vector), or the rate of change of a magnitude (velocity), or even the second derivative of another magnitude (distance). The same applies to the visual modality, where the detection of the second derivative of a magnitude (number of photons incident on the receptor) can be viewed as the transduction of intensity change. In this case, there is no more need to posit memory to account for the detection of intensity *change* than there is to account for the detection of intensity *value*. Either could, in principle, be taken as primitive.

Even the absolute time over which a change can be primitively detected is an empirical question and cannot be settled in advance of the data. This is true for the same reason that rate and magnitude detection can be interchanged in certain situations. We saw above that when certain rates and magnitudes are nomologically related (e.g. force and acceleration), the question which one of them a mechanism detects may be without empirical import. The same principle holds over long time spans. Whenever an equivalence class of histories of some system is nomologically connected to the current state of that system, detecting the latter can be tantamount to detecting the former. Thus, the issue of how long a time period of changes can be primitively detected must be settled by actually studying the functional capacities of the organism in question. In particular, the spatial limits of the immediately detected visual properties may extend beyond the retinal field and their temporal limits may extend beyond the measured refractory period of the visual system as neuroanatomically defined. But if they do, of course, then the retina and the neuroanatomical structures involved do not constitute the whole transducer mechanism for these properties.

In short, the methodological constraints on transduction that we have discussed so far considerably underdetermine a census of the directly detected properties, and they are compatible with a view of these properties quite different from the pointillism that Gibson deplores. The pressing issue is thus to understand what sorts of empirical considerations are operative in deciding which properties of the light are transduced, and what the empirical consequences of such decisions are. In what follows, we briefly review three such considerations. For expository purposes, we will take some of our examples from phonetics rather than vision. The question of which (if any) phonetic properties are transduced is a classical problem in psycholinguistics, and one in which the theoretical consequences of the various options reveal themselves with particular clarity.

7.1. Productivity

There is a *prima facie* assumption that productive properties of the effective stimulus are not directly transduced. Roughly, a productive property is one which determines an infinite equivalence class of (actual and possible) discriminable stimuli, such that the organism is, at least in principle, capable of identifying arbitrary novel stimuli which belong to the class. We will call this set of stimuli the *associated set* for a productive property. The property of being a token of an English sentence is, in this sense, a productive property, and its associated set contains all and only the actual and possible token utterances of English sentences.

The reason that productive properties are *prima facie* not transduced is that, in many of the most interesting cases, membership in the associated set is inferred from a prior identification of the internal structure to the stimulus. For example, we presumably infer the sentencehood of a token from a prior analysis of its internal lexical and syntactic structure. Since inferential processes are *ipso facto* not transductive, cases in which productive properties are recognized *via* the assignment of internal structure are *ipso facto* not cases of direct detection.

There are, however, lots of examples where productive properties *are* transduced, and these raise theoretical issues of considerable importance. It is, for example, perfectly possible to build a detector which transduces—or, to borrow Gibson's term, "resonates to"—the frequency of ambient sound or light, so the fact that there are indefinitely many physical displays that the organism can recognize as tones or colors is no evidence against the direct detection of colorhood and tonehood. In general, where the associated set that a property determines can be specified in terms of the values of some physical parameter, it seems possible—in principle, if not in fact—that the mechanism for the recognition of that property might be a resonator which is "tuned" to the relevant parameter values.

These considerations suggest an argument which, though rarely explicit, may well underlie the suspicion that the resonator model (and hence the direct detection theory) might after all provide a *general* account of perception. It may also help explain the emphasis on analogue models for perception that one finds in so many Gibsonian texts (see, for example, Prendle, Carello and Turvey (1980), Runeson (1980). We digress to consider this argument.

Consider a productive property like token sentencehood which appears, at first blush, to be nontransducible.

1. It seems plausible that every token sentence, including a merely possible one, is identical with some actual or possible physical object. It presumably follows that, for every token sentence there is some or other *identifying physical description*; i.e. a description that is true of that token and no other.
2. We can therefore, in principle, specify a physical property P that is *counterpart* to the property of being an actual or possible sentence token. P is specified by forming the (presumably infinite) disjunction whose disjuncts are the identifying physical descriptions for each of the actual and possible sentence tokens.
3. P is a physical property; that is, it is specified solely in terms of physical parameters.
4. Therefore, it is possible (at least in principle) to build a resonator for P.

5. Therefore, it is possible (at least in principle) to build a resonator for token sentencehood.
6. If there could be a resonator for token sentencehood, then perhaps we detect that property by resonating to it. Similarly for any other productive property.

This argument is important because it seems to show that the mere fact that a class of stimuli is a class of physical objects implies the possibility-in-principle of a resonator for the property with which the class is associated. And it is a short step from that conclusion to the claim that the postulation of inferential processes in perceptual recognition is always *heuristic* in the sense that, if the associated class can be identified by an inferential system, it can also be identified by a resonator (albeit, in some cases, a very complicated one).

The usual objection to this line of argument is to challenge the inference from 3 to 4. Notice that P will be a transducible property only if there exists a nomologically possible device (a detector) whose states are causally dependent on the presence of P. But the nomological possibility of such a device is by no means implied just by the assumption that P is physically specifiable. Though we have every reason to suppose that there must be some physical property which is common to all and only the possible sentence tokens, we also have every reason to believe that it is a highly disjunctive and arbitrary property; one which may perfectly well fail even to be finitely specifiable. But if P is disjunctive and arbitrary, then presumably P is not projectible (there are no laws about P). And if P is not projectible, there is no reason to suppose that a device whose states are lawfully related to the presence of P—*viz.* a P-detector—is empirically possible.

We believe this rebuttal to be well taken. However, at first glance it may seem to be denying the following obvious fact. An English speaker is, presumably, a physical device; and English speakers are capable of distinguishing arbitrary token sentences from arbitrary token nonsentences. So it looks as though the very existence of English speakers constitutes a constructive proof of the possibility of physical systems that are selectively sensitive to P, and hence of the possibility of transducers for P. But if P is transducible and coextensive with token sentencehood in all possible worlds, then token sentencehood must be transducible after all.

This argument is attractive but fallacious, as is the argument from (1) through (6). In particular, you cannot infer the transducibility of P from the fact that English speakers are selectively sensitive to token sentencehood, and you cannot infer the transducibility of token sentencehood from the assumption that there is a resonator for P (as in the inference from (4) to (5)). Both inferences fail, and for the same reason: they both assume that

the ability to detect membership in the associated set for a property implies the ability to detect the property. And this is wrong because it is perfectly possible for *distinct* properties to have the same associated set.

It is extremely tempting to assume that if two properties are coextensive for *both actual and possible* cases (as are, by assumption, P and token sentencehood) then the properties are identical, and the detection of one is the detection of the other. But this assumption is false. Distinct properties can be coextensive in all possible worlds, since properties that are coextensive in all possible worlds may nevertheless differ in their (higher order) properties.[11] This is by no means a quibble; it converts directly into differences in the empirical consequences of the hypothesis that English speakers respond selectively to token sentencehood, on the one hand, or to P on the other.

Consider the property of being a token sentence of Pig Latin (the property L_1). L_1 is very similar to English token sentencehood since you can pair any token English sentence with something which has L_1 by employing a trivial algorithm. But now consider the property P_1 which is the physical counterpart of L_1, in the same way that P is the physical counterpart of English token sentencehood. Suppose that $L_1 = P_1$ (i.e. that being a token sentence in Pig Latin is the same property as P_1). Then, since tokens of English and tokens of Pig Latin are similar in respect of their linguistic properties, and since we are assuming that their linguistic properties are *identical* to their physical properties (in particular, that being an English token *is* being P), it follows that tokens of English and tokens of Pig Latin must be similar in respect of their physical properties. But this consequence is false; a token of an English sentence may bear very little physical relation to a token of its Pig Latin translation.[12] Hence, token sentencehood cannot be the same property as P. To put the point quite generally, if two things are *similar* because one has A and the other A′, but *dissimilar* because one has B and the other has B′, it cannot be the case that A = B and A′ = B′.

---

[11] This point is obvious, and widely admitted, independent of the present examples. Consider the properties of being an equilateral triangle and being an equiangular triangle, which, though distinct, are coextensive in all possible worlds.

For an argument which runs along lines similar to those in the text, see Sober (1980).

[12] In general, linguistically natural transformations of sentences tend to be acoustically arbitrary, and *vice versa*. This is easy to understand on the assumption that P and token sentencehood are distinct properties. Roughly, a thing has the latter property in virtue of the identity and order of its segments; whereas a thing has the former property in virtue of the character of its acoustic analysis (e.g. in virtue of its formant structure). Not surprisingly, the linguistically natural transformations are (a subset of) the ones that substitute for segments or rearrange them, and the acoustically natural transformations are (a subset of) the ones that operate on formant structure. It is these considerations that the argument in the text relies on.

What we are saying is that even if token sentencehood and P are coextensive in all possible worlds, they nevertheless have different locations in the space of property similarities, and are thus *different properties.*[13] Hence the hypothesis that English speakers are responsive to token sentencehood leads to quite different predictions than the hypothesis that they are responsive to P. On the former hypothesis, it ought to be quite easy (having once learned to distinguish sentences from nonsentences) to then learn Pig Latin. Whereas, on the latter hypothesis it ought to be quite easy (having once learned to distinguish P from non-P) to learn to detect the property of being just like P except for some simple acoustical transformation—like having the acoustic spectrum inverted or reversed. It is entirely clear which of these hypotheses is true of English speakers. Pig Latin is easy to learn, but very simple transformations of the physical signal can entirely destroy the intelligibility of speech. So when we respond selectively to tokens, it is sentencehood, and not P, that we are perceiving.

Since P and token sentencehood are not the same property, the assumption that P is transducible would not at all imply that it is possible to directly perceive sentencehood, and the fact that we can respond selectively to the associated set for sentencehood does not imply the nomological possibility of a transducer for P. It may be that the associated set for any perceptible property is a set of actual or possible physical objects. It does not follow that there could, even in principle, be a resonator for every perceptible property. Probably the right thing to say is that the only properties that can be transduced are the projectible one: If a property is projectible, then by definition there are things whose states are lawfully connected to the presence of that property; i.e. things that detect that property. But it is by no means the case that all, or even most, of the properties that we can perceive are projectible.

As for P and token sentencehood, the right thing to say is probably that *neither* is transduced. We do not transduce P because P is not projectible, hence not the sort of physical property to which empirically possible detectors can be selectively tuned. And we do not transduce token sentencehood because the token sentencehood of a stimulus is inferred from the prior identification of its parts and their arrangement. The point to bear in mind is that *both* of these claims (and their conjunction) are compatible

[13]Of course, difference in similarity relations is a sufficient, but not a necessary condition for the nonidentity of token sentencehood and P. The reader who does not accept the claims about similarity may nevertheless agree that the two properties are distinct since the former, but not the latter, determines a natural kind; there are presumably scientific generalizations about sentencehood—*viz.* in linguistics—but there are surely no such generalizations about P.

with the assumption that we are able to identify the members of the stimulus set that is associated with both token sentencehood and P.

7.2. *Internal structure and generalization gradients*

As we saw in the preceding section, the reason that many productive properties cannot be viewed as transduced is that recognition of items belonging to their associated sets depends upon inference from assignments of internal structure: the items are recognized *via* a prior identification of the character and arrangement of their parts. In short, cases where perceptual recognition depends upon analysis of internal structure are cases where the direct perception model fails to apply.

The present question is how this criterion for nontransduction is to be applied; that is, how we can tell whether perceptual recognition is mediated by the assignment of internal structure. One answer is that, in such cases, the generalization gradient for the stimulus is often predictable from relations of similarity and identity defined over its parts.

As an illustration of what is at stake, consider the problem of deciding whether the perception of speech involves the direct detection of phones, or phonemes, or syllables or perhaps some higher level entities like words or sentences. The view that what is directly detected is one of the smaller units (e.g. phones) has as a consequence the prediction that segments which share a phone ought to be perceived as more similar than contrasted segments whose phonetic transcriptions do not overlap. By contrast, the view that larger units (such as syllables or words) are transduced would not lead directly to such predictions. Notice that this predictive asymmetry between the models holds even if the acoustic and phonetic structure of a speech segment are isomorphic, so that the phonetic similarity of segments guarantees their accoustic similarity. Suppose that a pair of segments have acoustically identical parts wherever they have phonetically identical parts. Even so, phonetic similarity warrants predictions of generalization only on the *further* assumption that the segments are identified *via* the prior identification of their parts—i.e. that the identification of the acoustic structure is nonholistic. On holistic assumptions (e.g. that recognition is accomplished by the application of a segment-length acoustic template), even the fact that a pair of segments exhibits partially identical acoustic substructures provides no grounds for predicting generalizations from one to the other.

Gibson appears to recognize the importance of the issue of internal structure in percepts, since he frequently refers to the perceptual objects in the environment as being "nested". However, he denies that the detection of such nested units is cascaded in the sense that the identification of the higher

units is dependent upon the prior identification of the lower ones. He *must* deny this because the mechanisms in virtue of which the identification of the former is contingent upon the identification of the latter could only be inferential. Perhaps Gibson's view is that units at all relevant levels are simultaneously and independently detected. But, in and of itself, this assumption does little to explain the facts about generalization and similarity structure among percepts. What is required for that purpose is not merely that the information about lower level units be *recorded*, but also that the hierarchical and combinatorial relations among the various units should be among the properties of the stimulus which the organism registers in the course of perceptual identification. Thus, in order to account for the fact that /ba/ and /pa/ are perceived to be related in the same way as /bo/ and /po/, and that /ba/ and /ab/ are perceived to be related in the same way as /bo/ and /ob/, we need to assume the detection, not only of syllables and their constituent phones, but also of the relevant relations of order and inclusion.

Our point is not, of course, that making the appropriate predictions about generalization is *incompatible* with assuming the direct detection of higher-order units; only that their reconciliation often depends upon postulating ancillary mechanisms that the componential approach can do without. For example, one could assume holistic perception and account for generalization gradients by postulating a further process in which the internal structure of higher level units is retrieved by lookup (as in models of speech detection where syllabic identity is recovered first and phonetic structure is assigned by accessing a syllable dictionary). Such postulation is sometimes justifiable. For example, there is sometimes evidence that the higher units can be identified faster than their components, and such results provide *prima facie* (though by no means univocal) evidence for a "top down" order of processing. In any event, the need to account for the subject's perception of similarity and generalization structures among stimuli is real, and it provides a source of constraint on decisions about which of the properties of the stimulus should be viewed as directly detected.

7.3. Cognitive penetrability

We have, in effect, been taking transducers to be devices whose output is lawfully dependent upon the character of their input. The output of a perceptual mechanism, by contrast, may show simultaneous effects of the character of its input and of the inferential operations that it performs. We now note an important consequence of this view of transduction: the character of a transducer's response is not, in general, sensitive to the beliefs and

utilities of the organism. We refer to this as the "cognitive impenetrability" of transduction. (For general discussion, see Pylyshyn, 1980).

It is notorious that the expectations and utilities of an organism selectively affect what it sees and hears. The subject biased to expect a picture of an old woman is not likely to see the ambiguous stimulus as a picture of a young girl; the hungry organism responds at low thresholds to food that the well fed organism misses altogether. The general consideration is that the range of "perceptual properties"—the range of properties that the perceptual system can respond to—is practically unbounded, and the mechanisms of perception can be selectively biased to very nearly any of them. This sounds surprising only if we forget that perception involves the integration of current inputs with background information; how the perceptual apparatus is cognitively tuned is largely a question of which such background information is being deployed. Whereas, since the output of a transducer is, by assumption, nomologically dependent upon the properties that it responds to, and since the number of properties that are lawfully connected with the output states of such devices is, in any event, very much smaller than the number of perceptible properties, the possibilities for the cognitive retuning of a transducer are correspondingly restricted relative to the possibilities for cognitive retuning in perception.[14]

The conclusion is that the right kind of cognitive penetrability is a counterindicant of transduction. The "right kind" of penetrability is exhibited when the property to which a mechanism responds proves to be arbitrarily sensitive to the content of the subject's goals, beliefs, and utilities.[15] We have qualified the claim in this way because we want to admit the possibility of such relatively undifferentiated effects of goals and beliefs on the behavior of transducer mechanisms as would be exhibited in cases of centripetal damping; (as, for example, when the impedance of the ear or the aperture of the eye is altered as a mechanism of selective attention). Notice, however, that such examples typically concern modulation of the *amplitude* of a transducer's

[14]We stress that the relevant dependencies are those between stimulus properties and *output states* of the device. As we mentioned above, there is probably a large number of types of ambient energy to which the states of any physiological mechanism are responsive; however, most of these responses have no functional significance and hence do not count as outputs.

[15]For these purposes, we have treated the cognitive impenetrability of transducers as simply a consequence of the fact that their inputs are nomologically sufficient for their outputs. There is, however, a deeper point. 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. We can then view cognitive penetration as a process which makes additional premises available to such arguments—e.g. premises which express background beliefs, goals, and utilities. But if that is right, if that is what cognitive penetration *is*, then it follows that there cannot, in principle, be cognitive penetration of noninferential processes like transduction.

response rather than of the stimulus property to which its response is specific; and one would scarcely expect that such modulation would exhibit a *detailed* specificity to the content of the organism's cognitive states.

7.4. *Cognitive impenetrability and "compiled detectors"*

It is important to notice that, though cognitive penetrability generally suggests nontransduction, the inverse does not hold. There are, in all probability, many cases of what might be called cognitive automatisms: computational processes which, though inferential, are nevertheless quite rigidly insensitive to modulation by beliefs and utilities. In fact, some of the most important work inspired by the Gibsonian tradition can be viewed as the identification of levels of mental representation which are computed by such processes.

We have seen that the general issue of specifying *which* properties of an effective stimulus are directly detected and which are constructed or inferred *via* processes that access the organism's tacit knowledge, is one of the fundamental problems of perception. Gibson is to be credited with promoting the view that many of the presuppositions of classical theories of perception were based upon untenable assumptions about how this question should be answered. Such assumptions, while not an inherent part of the Establishment view, dominated much of the early theorizing about vision which, for example, assumed that the input to a visual system should be described in terms of point-intensities on the retina at an instant of time. What such theorizing failed to appreciate, and what Gibson helped to dramatize, is that (a) there is no theoretically *neutral* description of a perceptual stimulus—the form that one's theory of perception takes is extremely sensitive to exactly what one assumes are the inputs to the perceptual system, and (b) what ought to be considered to be the input is an empirical question. Largely due to the efforts of Gibson and his coworkers there has been a revival of interest in viewing the input to the visual system as consisting of spatially and temporally distributed properties of the light array. For example, there has been much fruitful work on the perception of form and of motion which relies on the existence of reliable correlations between spatially extended optical texture gradients at the retina and such properties of the environment as the depth, slant and shape of surfaces (e.g., Stevens, 1979); between relative movements of elementary retinal features and the three dimensional shape of a moving rigid form (e.g., Ullman, 1979) or the movement of the observer (e.g. Prazdny, (1981); or between the retinal disparity of primitive features and the perception of depth (Marr and Poggio, 1976).

Although such correlations have long been known to be important for vision, the recent work has added an important new dimension to early

speculative theories. Many of these investigations have located the precise features of the dynamic retinal pattern involved in the correlation, and have shown that there is a specifiable mathematical basis for its reliability. They have, in other words, succeeded in specifying *which* properties of the light constitute good indicators of perceived environmental properties, and *why* they do so. Thus, to take one example from the above list, Ullman has shown that if one assumes (a) that the organism can reliably determine which optical retinal features arise from the *same* point on a moving stimulus across time and (b) that the stimulus features maintain their relative three dimensional distances (i.e. that they are located on a rigid body), then given a sufficient sample of the dynamic retinal pattern, a computable one-to-one mapping onto the three dimensional form can be defined. Furthermore, the few assumptions made by this model have been verified empirically. It might thus be argued that the organism need not make complex, knowledge dependent inferences in order to determine the three dimensional shape of an object; it need not, in other words, go through any such process as that of using highly partial and sketchy cues to frame a hypothesis about the identity of the object, and then using its general world knowledge to infer the shape under observation. At least in this case, it appears that specific properties of the ambient array are able, as Gibson would put it, to specify the three dimensional shape uniquely.

Such results have helped kindle the current interest in Gibson's theory. Many investigators have even adopted some of Gibson's terminology, and describe these cases as illustrating the pickup of invariants. It is important, however, to see that such a terminological policy merely pays tribute to Gibson's attack on the pointillist and static snapshot presuppositions of early Establishment theories. There is really nothing in the recent research that is at odds with the Establishment story (or, for that matter, with the Helmholtz story) about perceptual *inference*. In particular, the fact that we can, in some cases, provide a precise account of the locus of a reliable correlation between light features and layout features—and, indeed, even show the conditions under which a perfect correlation is possible—does nothing to remove the need for postulating computational processes in accounting for the perceptual capacities of the organism.

Recall that we claimed that there are two aspects of the inferential model: the claim that properties of the layout are inferred from properties of the light, and the claim that the directly detected properties of the light generally underdetermine features of the layout. It is patent that nothing in the new research challenges the former claim; so long as the organism is detecting light patterns and not layout patterns, there is no way to avoid the conclusion that perceptual knowledge of the latter is inferred; this remains true no

matter how perfect the correlation between light and layout may prove to be. What such results as Ullman's do is explain why such inferences are *sound* in specified circumstances.

More to the point, however, is the fact that such results *do not imply that the properties of the ambient optic array which correlate reliably with surface orientations, shapes and motions of objects are the ones that are directly detected.* True, the findings suggest that the construction of certain levels of representation may be cognitively impenetrable by, for example, the subjects' prior beliefs about the stimulus that they are viewing. But, as we remarked above, while some notion of cognitive impenetrability provides a necessary condition for direct detection, it certainly does not provide a sufficient condition.

In particular, the research leaves open the possibility that the detection of such properties as, for example, the texture flow of the light when it is accelerating in a certain direction while systematically changing its textural density, is mediated by the detection of more primitive properties such as, perhaps, the relative magnitude of texture densities at various retinal locations, the relative rates of change of these densities in relation to their relative locations, the existence and locations of sharp discontinuities in these or other light features, and so on. The existence of such mediation in fact seems quite likely. As we argued earlier, we want to be able to account for the origins of percepts in terms of the properties of their effective stimuli, and this necessitates that we be able to detect subfeatures of the light like the ones just mentioned since they themselves can and do give rise to percepts. Since, moreover, the inference from the light to the layout depends upon details of the arrangement of such subfeatures in relation to one another and to global properties of the texture flow, it is, as we argued above, vastly implausible that the global feature detectors should be primitive. There are, after all, arbitrarily many ways in which such subfeatures can be arranged in the global array, and each of these arrangements has different implications for the organization of the layout and for the generalization gradients among percepts. It thus seems obvious that, barring the possibility of a primitive "resonator" that is selectively sensitive to these various proximal arrangements, the best assumption is that the output of the complex property detector must be *constructed* from primitively detected subfeatures.

What we have is, in effect, a conflict between the demands of the productivity and internal structure criteria, on the one hand, and the cognitive penetrability criterion, on the other. The discovery of cases where these criteria conflict may well be the most important contribution of the Gibsonian tradition to Establishment theory. Though the empirical research we have alluded to offers no support for direct perception of any kind, it does

provide suggestive evidence for the existence of levels of complex detection which are autonomous, stimulus bound, cognitively impenetrable, and hence unmediated by a certain kind of deliberate deduction from generally available knowledge (e.g., the sort of deduction which enters into explicit question answering).

The notion of an autonomous computational "reflex", compiled from elementary constituents but impervious to general cognition, is one which finds widespread application. For example, there is good evidence that much of lexical identification, and even of syntactic analysis, may be of this type; and it may be that some of what goes under the name of "imagistic processing" belongs to this category. (See, however, Pylyshyn 1980, for an argument that much imagistic processing is cognitively penetrable, hence not compiled). Similarly, Marr (1976) has argued for a level of complex representation in form recognition which he calls the "primal sketch" and which is assumed to be largely or entirely stimulus driven. The argument is that the computational processes eventuating in a primal sketch can be empirically demonstrated to be 'reversible'. Suppose that the primal sketch S arises from the ambient array A *via* the computational process C. It turns out that, if C is, in effect, applied backwards so as to generate an ambient array A′ from S, A′ will prove to be very nearly indistinguishable from A. This suggests that the mapping from stimulus arrays onto primary sketches must be very nearly one-to-one. If, however, a primal sketch is uniquely determined by an ambient pattern, then the processes which eventuate in its integration cannot be cognitively penetrable.

If the existence of processes of "compiled detection" is borne out by further research, that will indeed tell us something interesting about the modularity of perceptual systems. But it should be emphasized that such findings would have no particular bearing on the Gibsonian proposal of direct perception either of layout or of global features of the light. Transduction, as we have seen, is not to be inferred from cognitive impenetrability alone. Indeed, the only connection which a compiled detection theory of perception would have with the Gibsonian view is that they would share certain adversaries. They are both opposed to the old sense data theories, to the excesses of the "new look" in perception (cf. Bruner, 1957), and to those expectation-driven models, popular in certain quarters of artificial intelligence, which sometimes seem to deny that any significant amount of stimulus bound processing occurs after the activation of the rods and cones. But that is a fairly remote connection; the compiled detector story may be a revolution in the psychology of perception, but it is not a *Gibsonian* revolution.

Finally, it is worth bearing in mind that, while many perceptual processes may well be compiled, there is no reason to believe that all of them are.

Object recognition, for example, is a perceptual process *par excellence*, and it appears to be cognitively penetrable through and through.

8. Conclusion: The problem of intentionality

Perception is interesting but cognition is more interesting still. It is, as we have seen, no small part of the importance of Gibson's revisionism that it is supposed to extend, not just to the theory of perception, but to cognitive processes at large. In this section, we argue that focusing on the problem of perception led Gibson seriously to underestimate the difficulty of constructing a cognitive psychology that dispenses with the mental representation construct. Our argument will be that (a) the prototypical perceptual relations (seeing, hearing, tasting, etc.) are extensional (and even where they are not; Gibson, in effect, treats them as though they were); (b) whereas, on the contrary, most other prototypical cognitive relations (believing, expecting, thinking about, seeing as, etc.) are intentional; and (c) the main work that the mental representation construct does in cognitive theory is to provide a basis for explaining the intentionality of cognitive relations. Our moral will be that one has not made a start on developing a representation-free cognitive psychology until one has (at least the outline of) a representation-free theory of intentionality; and that Gibson's concentration on perception led him to overlook this crucial point.

Compare *recognizing* with *recognizing as*. If you recognize the man on the white horse, and the man on the white horse is Tonto's best friend, then you recognize Tonto's best friend. Similarly with any other description true of the man on the white horse; if Tonto's best friend is The Lone Ranger, and you recognize Tonto's best friend, then you recognize The Lone Ranger. However, suppose that you recognize Tonto's best friend *as* the man on the white horse. Then, even though the man on the white horse is The Lone Ranger, it does *not* follow that you recognize the man on the white horse as The Lone Ranger. ("Who was that man on the white horse?" "That was *The Lone Ranger.*")

Roughly, *seeing* works like *recognizing*, and *seeing as* works like *recognizing as*. If you look up at the sky and see the Pole Star, then, in doing so, you see: a certain very large ball of hot flaming gasses; the star that the ancients used to steer by; a star that is not visible from South of the Equator; ...etc. Whereas, if you look up at the sky and see the Pole Star *as* the Pole Star (i.e. you see the Pole Star and take it to be the Pole Star) it does not follow that you see the Pole Star as a large ball of hot flaming gasses, or as the star that the ancients used to steer by, or as a star that is not visible in the Southern

Hemisphere. What you see when you see a thing depneds on what the thing you see *is*. But what you see the thing *as* depends on what you know about what you are seeing. Contexts that work like *see* and *recognize* are called *extensional contexts*. Contexts that work like *see as* and *recognize as* are called *intentional contexts*.

We are not doing ordinary language analysis, so we do not care whether it is *precisely* true that the English verb "see" is uniformly extensional. The point is, rather, that a psychology which limits itself to considering only the extensional relations misses something that appears to be essential to explaining how our perceptual transactions with the world affect what we know and believe. Here is Smith at sea on a foggy evening, and as lost as ever he can be. Suddenly the skies clear, and Smith sees the Pole Star. What happens next? In particular, what are the consequences of what Smith perceives for what he comes to believe and do? Patently, that depends on what he sees the Pole Star *as*. If, for example, he sees the Pole Star as the star that is at the Celestial North Pole (plus or minus a degree or two), then Smith will know, to that extent, where he is; and we may confidently expect that he will utter "Saved!" and make for port. Whereas, if he sees the Pole Star but takes it to be a firefly, or takes it to be Alpha Centuri, or—knowing no astronomy at all—takes it to be just some star or other, then seeing the Pole Star may have no particular consequences for his behavior or his further cognitive states. Smith will be just as lost after he sees it as he was before.

If we want to make predictions from what someone perceives to what he does, or to the cognitive consequences of the perception, we must be able to distinguish between *merely* seeing the Pole Star and seeing the Pole Star *as* the Pole Star. In particular, since merely seeing the Pole Star and seeing it as the Pole Star have different psychological consequences, a cognitive theory must distinguish the state that the organism is in when it does the one from the state that it is when it does the other. And it must make the distinction *in the right way*. Whatever state the theory describes the organism as being in in consequence of having seen the Pole Star as such-and-such must be the right kind of state to explain the psychological consequences of seeing it that way. (If seeing the Pole Star as the Pole Star leads the astronomically sophisticated to jump with joy, then the state that the theory assigns to an organism which sees the Pole Star as the Pole Star must be such as to contribute appropriately to explaining the ensuing glee). Our point is, then, that a theory of the intentional relations must be at the very heart of a cognitive psychology insofar as the psychologist seeks to derive predictive consequences from his claims about what the organism has perceived. To do cognitive psychology, you must know not just what the organism perceives, but how it takes what it perceives.

This is where the mental representation construct does its main theoretical work. In effect, it allows us to understand *seeing as* in terms of *seeing* and *mentally representing*. It thus comes to grips with the fact that the cognitive consequences of perception depend not just on *whether* the world is seen, but also on *how* it is seen. Just how the mental representation theory is supposed to work in this area is, of course, a matter of intense disagreement among Establishment theorists. But the general line is clear enough. To see the Pole Star as the Pole Star is (a) to see it; and (b) to take what one sees to satisfy some such representation as, for example, the open sentence "...is the Pole Star". It is perfectly possible for someone to see the Pole Star and *not* take it to satisfy a representation of that sort. For example, one might see it and believe only that it satisfies some such representation as "...is a firefly" or "...is some star or other". Since all these representational states are compatible with seeing the Pole Star, it is not surprising that seeing the Pole Star can have different consequences for different people or for the same person at different times. The cognitive (and hence the behavioral) consequences of what you see depend on how you represent what you see, assuming that the Establishment theory of the intentional relations is true.

Perhaps, however, it is *not* true. Our point is only that you need some theory or other to work in this area, and that the representational account is an open option. Conversely, if—like Gibson—you propose to do without the mental representation construct, you need a workable alternative to the representation account of the intentional relations. And this Gibson does not have.

We are now in a position to understand Gibson's basic strategy in some depth. To put the point as neutrally as we can, what *everybody* has to deal with, *vis à vis* the problem of intentionality, is the fact that stimuli enter into the causation of behavior under many different aspects. What one man responds to as the Morning Star, another responds to as the Evening Star, and their responses may, in consequence, be quite different even though the Morning Star and the Evening Star are one and the same astronomical object.

The Establishment theory seeks to accommodate such facts by proliferating mental representations. The idea is that the very same object may be represented in many different ways, and someone who responds to Venus as the Morning Star differs from someone who responds to Venus as the Evening Star in respect of the ways that they represent the planet. Specifically, the locus of intentional distinctions, according to the Establishment theory, lies in the consideration that representations which differ in semantic content may nevertheless apply to the same object (Frege, 1949). The hope is, however, that theoretical appeals to the semantic content of mental representations will ultimately prove dispensable; in particular, that identities and

differences among the semantic contents of mental representations will be reconstructable in terms of identities and differences among their functional (e.g., causal) roles. Such a functional account of the semantic properties of mental representations would then round out the Establishment theory of intentionality. (For discussion of this aspect of the relation between functionalism and representational theory of mind, see Fodor, 1980, 1981; Loar, forthcoming; Field, 1978; Pylyshyn, 1980.)

What the Establishment does by proliferating *mental representations*, Gibson proposes to do by proliferating *properties*. Instead of saying that the same astronomical object is represented now as the Morning Star and now as the Evening Star, Gibson says, in effect, that the same object has the two distinct properties of being the Morning Star and being the Evening Star. Which way we respond to the object depends on which of these properties we happen to pick up.

It is not surprising that Gibson has this option. You can do with distinctions among properties a lot of what you can do with distinctions among the semantic contents of representations. This is because, according to at least one standard account, *property* and *meaning* are interdefinable notions: properties are distinct if they are expressed by nonsynonymous representations, and representations are synonymous if they express the same property.[16] Nothing much appears to be gained by chasing around this circle, which is why philosophers like Quine, who are vehemently suspicious about meaning, are equally vehemently suspicious about properties.

So, there is a sense—perhaps a rather uninteresting sense—in which Gibson can make do with distinct properties where Establishment theories postulate semantic distinctions among mental representations. But, of course, you have to pay the piper sometime. The price that Gibson pays is the failure of his theory of specification.

Property is an *intentional* notion in the sense that coextensive sets may correspond to distinct properties. (The Morning Star = the Evening Star, but *the property of being the Morning Star* ≠ *the property of being the Evening Star*. Or so, at least, we must assume if we are to explain differences in responses to Venus by appeal to differences in the properties picked up.) However, *specification* is an *extensional* notion. Specification comes down to correlation (see above), 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. Specification cannot, then, explain property pickup.

[16]It is, of course, possible that Gibson has some other (some nonsemantic) notion of property in mind. One cannot tell because the issue of property individuation is not one that Gibson discusses.

Everybody has to face the issue about intentionality somewhere. For Gibson, push comes to shove with the question: what is it for an event (a configuration of the light, etc.) to specify a *property*? To say that Gibson has no theory of intentionality is to say that he has no answer to that question. Or, to put it the other way around, the failure of Gibson's theory of specification is no minor flaw in his theory. It marks the precise point at which Gibson's treatment of intentionality proves to be bankrupt.

In a nutshell: the move from semantic distinctions among representations to ontological distinctions among properties, in and of itself, buys the psychologist nothing. The problem of substance is to provide an independent account either of the meaning of a representation (Establishment style) or of the specification of a property (Gibson style). The former problem *may* be tractable since it may be that the meaning of a representation can be reconstructed by a reference to its functional role; that is the hope by which Establishment theories live. But Gibson gives no indication at all of how the latter problem is going to be solved. Where the Establishment line offers, anyhow, a pious hope, the Gibsonian line offers only a dead end.

We said that Gibson's basic strategy is to use the (intentional) notion of a property to do what Establishment theories do with the (intentional) notion of the semantic content of a representation. Exegesis is complicated, however, by the fact that Gibson's adherence to this program is only sporadic. To put the point very crudely, it seems clear that one's theory of intentionality will have to postulate two of *something* if it is to account for the two ways of seeing Venus. (In fact, of course, it will be necessary to postulate infinitely many of something since there are infinitely many ways of seeing Venus.) The Establishment proposal is that we postulate two different *mental representations of Venus*. Gibson's proposal is that we postulate two *properties* of Venus to which perceptual mechanisms can be selectively tuned. But there is a third option with which Gibson appears occasionally to flirt: namely, postulate *two Venuses*.

Consider, to vary the example once again, seeing the Pole Star as a distant ball of hot gasses. Seeing the Pole Star that way is a rather late achievement both ontogenetically and phylogenetically. It depends on knowing a lot about stars. Yet seeing the Pole Star *is* seeing a distant ball of hot gasses, so the question arises how you can see it and yet not see it that way. Since, however, Gibson has no account of intentionality, he is faced with the problem of developing a theory of perception which provides an account of *seeing* without raising that sort of issue about *seeing as*. Of the various possibilities, Gibson sometimes appears to want to take what strikes us as clearly the least advisable. He sometimes denies (or so a literal reading of the text suggests) that what you see when you look at the sky are stars.

The idea is, apparently, that there are two kinds of *things*: there are the little whitish things that you see when you look up, which count as *bona fide* ecological objects and hence as *bona fide* objects of perception; and there are also the large hot balls of gas that astronomers describe, which count as astronomical objects and hence as *not bona fide* objects of perception. This gets Gibson out of the need to explain how it could be that the very same object can be seen now as a little whitish thing and now as a (very distant) large ball of gasses. That is, it gets him out of the need for a theory of intentionality *vis à vis* what happens when you see a star. But for this Gibson pays an utterly unreasonable price; if we take him seriously, we will have to say that the astronomer is *wrong* when he claims to have discovered that the Pole Star, though it looks small and chilly and relatively close, is actually large and hot and very far away. He is wrong because, on Gibson's account, the (astronomical) object that is large and far away is not identical to the ecological object that looks small and close by. Gibson does not attempt to say what the relation between the astronomical object and the ecological object is, assuming that it is not identity. In this he is probably well advised.[17,18]

Perhaps, however, you are prepared to pay this price; in effect, to avoid the problem of seeing as by postulating many different things to see where a more plausible theory makes do with many different ways of seeing the *same* thing. Even so, the wriggle will not work. There is a subtle connection between Gibson's misleading talk of direct realism and his refusal to face the problem of intentionality. This point now needs to be addressed.

Remember that, though Gibson sometimes says that we see ecological properties of the layout "directly", it turns out that this sort of direct seeing is mediated by the pickup of information in the light. Gibson admits, in effect, that finding out about the layout depends on first finding out about the light. When we discussed this issue, we emphasized that Gibson's (tacit)

---

[17]Gibson is rather less explicit about all this than we have perhaps made him seem, though there are passages which appear to admit no other interpretation of his views. For example: "The environment of animals and men is what they perceive. The environment is not the same as the physical world, if one means by that the world described by physics" (p. 15). What Gibson ought to have said is "The environment *as perceived* is not the world *as described by physics*". The reason he did not say this is presumably that to do so would have been to make the category *ecological object* overtly intentional —an ecological object would then have been an object as represented in terms of ecological parameters. And Gibson could not do *that* because, as we have seen, he has no theory of intentionality.

[18]Some of our best friends are prepared to quantify over intentional objects (e.g. for purposes of constructing model theoretic interpretations of modal expressions). But that is a far cry from taking intentional objects to be *objects of perception* as Gibson (and, by the way, Brentano before him) appears to be inclined to do. For one thing, if merely intentional (including non-actual) objects can be perceived, we will have to give up the enormously plausible principle that perception is mediated by causal transactions between perceiver and perceivee. Non-actual objects cannot, of course, be actual causes.

concession raises the question of how—by what mental process—finding out about the light eventuates in finding out about the layout. We can now see that this whole issue is implicitly involved with problems of intentionality since *finding out* is itself an intentional relation.

If I find out that (see that, perceive that) the Pole Star is overhead, and the Pole Star is the star I ought to steer by, it does *not* follow that I find out that (see that, perceive that) the star I ought to steer by is overhead. More to the point, suppose that I find out that the light is in a certain configuration, and suppose that the light's being in that configuration is the same state of affairs as the light's being caused by a certain feature of the layout. It *does not follow* from these premises, that I have found out that the light is in a configuration caused by a certain feature of the layout. To get from finding out about the former to finding out about the latter, I have to get from representing the properties of the light in one way to representing them in another way; in effect, I have to make an inference. Missing the point about inference, missing the point about mental representations, and missing the point about intentionality are thus all aspects of missing the same point.

So, for example, Gibson writes that we correctly perceive the unchanging shape of rigidly moving objects "...not because we have formed associations between the optical elements, not even because the brain has organized the optical elements, but because the retinal mosaic is sensitive to transformations as such" (Gibson, 1957, p. 294; quoted by Ullman, 1980). But, to transformations of *what*? Not of the object *per se*, since that would be ruled out by the counterfactual support condition (see above). So, then, presumably, to transformations *of the light*. But the problem of perception is not how we get epistemically related to the transformations of the light; it is the problem of how we get perceptual knowledge of the shape of the object. Gibson must be thinking something like this: the visual system is sensitive to how the light is transformed; for the light to be transformed in a certain way is for it to be reflected from objects with a certain sort of structure; hence the visual system is sensitive to the structure of objects. But, as we have seen, the second premise is not true in the general case (*vide* the illusions) and, more to the present point, even if it were true the conclusion does not follow from the premises. To think it does is to fail to understand the intentionality of such key relations as "being sensitive to". Because he did not take the intentionality of these relations seriously, Gibson greatly underestimated the magnitude of the concession implicit in admitting that only the light is detected directly.

To summarize: Even if all you want is to construct a theory of perception, you cannot do much without encountering problems about intentionality, although it is true that many of the key perceptual relations are more or less

extensional. When, however, you try to construct a theory about how perception relates to cognition at large (or, of course, of the nature of the non-perceptual cognitive processes) problems of intentionality come immediately to the fore.[19] According to the Establishment theory, this is no surprise: the mind is a mechanism for the manipulation of representations, and how what you see affects what you know is primarily a matter of how you represent what you know and see. This is what modern cognitive theory has inherited from the classical tradition in epistemology, and, as we remarked, it may be wrong. But there will be no successful anti-Establishment revolution in cognitive psychology until some alternative to this account is provided. What is finally and fundamentally wrong with the Gibsonian treatment is that it has not grasped that fact.

References

Block, N. (Ed.), (1980) *Readings in Philosophy of Psychology.* Cambridge, Mass., Harvard University Press.

Bruner, J. (1957) On Perceptual Readiness, *Psychol. Rev., 64,* 123–152.

Chomsky, N. (1959) A Review of B. F. Skinner's *Verbal Behavior,* reprinted in Block, N. (Ed.) *Readings in Philosophy of Psychology, Vol. 1.* Cambridge, Mass., Harvard University Press.

Dretske, F. (1978) The Role of the Percept in Visual Cognition, in C. Wade Savage (Ed.) *Perception and Cognition, Issues in the Foundations of Psychology, Minnesota Studies in the Philosophy of Science, Vol. IX,* University of Minnesota Press, Minneapolis.

Field, H. (1978) Mental Representation. *Erkenntnis,* Vol. 13, pp. 9–61.

Fodor, J. (1980) *Representations.* Vermont, Bradford Books.

Fodor, J. (1981) The Mind-Body Problem, *Sci. Amer., 244,* 114–123.

Frege, G. (1949) On Sense and Nominatum. Reprinted in H. Feigl and W. Sellars (Eds.) *Readings in Philosophical Analysis,* New York, Appleton Century Crofts.

Gibson, J. (1957) Optical Motions and Transformations as Stimuli for Visual Perception, *Psychol. Rev., 64,* 288–295.

Gibson, J. (1966) *The Senses Considered as Perceptual Systems.* Boston, Houghton Mifflin.

Gibson, J. (1979) *The Ecological Approach to Visual Perception.* Boston, Houghton Mifflin.

Goodman, N. (1954) *Fact, Fiction and Forecast.* London, The Athlone Press.

Hess, E. (1975) The Role of Pupil Size in Communication. Reprinted in Atkinson and Atkinson (Eds.) *Mind and Behavior, Readings from Scientific American.* San Francisco, W. H. Freeman and Co., 1980.

Hochberg, J. (1968) In the Mind's Eye. In R. N. Haber (Ed.) *Contemporary Theory and Research in Visual Perception.* New York, Holt.

Loar, B. (forthcoming) *Mind and Meaning.* Cambridge, Cambridge University Press.

[19] They crop up a lot earlier if you hold—as it is probably right to do—that every case of seeing a thing involves seeing the thing as something or other. True, you can see the Pole Star and not see it *as* the Pole Star. But, quite plausibly, you cannot see the Pole Star without seeing it as *either* the Pole Star, or a little speck of light, or a firefly, or the star that Granny likes best... or *something.* If this is right, then though *seeing* is an extensional relation, some of the logically necessary conditions on 'x sees y' are intentional.

Marr, D. (1976) Early Processing of Visual Information. *Philosoph. Trans. Royal Soc. London, 275,* 483–534.

Marr, D. and Poggio, T. (1976) Cooperative Computation of Stereo Disparity. *Science, 194,* 283–287.

Prazdny, K. (1981) Egomotion and Relative Depth Map from Optical Flow. *Bio. Cybernetics, 36,* 87–102.

Prendle, S., Carello, C. and Turvey, M. (1980) Animal-Environment Mutuality and Direct Perception. *The Behavioral and Brain Sciences,* in press.

Pylyshyn, Z. (1980) Computation and Cognition, Issues in the Foundations of Cognitive Science. *The Behavioral and Brain Sciences, 3*(1), 111–169.

Pylyshyn, Z. (1981) The Imagery Debate: Analog Media *versus* Tacit Knowledge. *Psycho. Rev., 88,* 16–45.

Reed, E. and Jones, R. (1978) Gibson's Theory of Perception: A Case of Hasty Epistemologizing? *Philosophy of Science, 45,* 519–530.

Runeson, S. (1980) There Is More to Psychological Meaningfulness than Computation and Representation. *The Behavioral and Brain Sciences,* in press.

Skinner, B. F. (1957) *Verbal Behavior,* New York, Appleton Century Crofts.

Stevens, Kent (1979) *Surface Perception from Local Analysis of Texture and Contour,* MIT Ph.D. Thesis.

Sober, E. (1980) Why Logically Equivalent Predicates May Pick Out Different Properties. Mimeog., University of Wisconsin.

Stich, S. (1978) Beliefs and Subdoxastic States. *Philosophy of Science, 45,* 499–518.

Turvey, M. (1977) Contrasting Orientations to the Theory of Visual Information Processing. *Psychol. Rev., 84,* 67–88.

Turvey, M. and Shaw, R. (1979) The Primacy of Perceiving: An Ecological Reformulation of Perception for Understanding Memory. In L. G. Nillson (Ed.) *Perspectives on Memory Research, Essays in Honor of Uppsala University's 500th Anniversary.* New Jersey, Erlbaum.

Ullman, S. (1979) *The Interpretation of Visual Motion,* Cambridge, Mass., MIT Press.

Ullman, S. (1981) Against Direct Perception. *The Behavioral and Brain Sciences,* in press.

Uttal, W. (1967) Evoked Brain Potentials: Signs or Codes? *Perspectives in Biology and Medicine, 10,* 627–639.

Warren, R. (1970) Perceptual Restoration of Missing Speech Sounds. *Science, 167,* 392–393.