ALVIN I. GOLDMAN

PERCEPTUAL OBJECTS*

What are the conceptually necessary and sufficient conditions for a person, or organism, to perceive a given object? More precisely, what is the nature of our ordinary thought about perception that gives rise to our willingness or unwillingness to say that *S* perceives *O*? Some form of causal theory of perception is now, I think, widely accepted. Such a theory maintains that it is part of our concept of perception that *S* perceives *O* only if *O* causes a percept, or perceptual state, of *S*. I accept this causal requirement, though with some qualification. The crucial problem for the theory of perception is what must be added to the causal requirement. Not all causes of a percept are said to be perceived. Smith looks at a tree illuminated by the sun, when the sun is not itself in view. His visual percept is caused by the tree, but not only the tree. The percept's causal ancestry includes the sun and the array of light two inches from Smith's eyes. Why do we say that he sees the tree, but not the sun or the array of light? What principles underlie this invidious choice among causes? Similarly, we think of a bat as perceiving objects in his path, but not as perceiving himself. Yet, since he emits the sonar signals these objects reflect, he is himself a cause of his percepts. Why don't we say he perceives himself?

I

My proposed solution focuses indirectly on the information acquired through perception. The primary function of perception is the generation of information, or true belief, about the organism's environment. Organisms deploy their sensory apparatus in ways designed to acquire such information. They focus their eyes toward an object they want to know about; they move their hands, paws, or claws over objects that interest them; they sniff the air, perk up their ears, and extend their antennae when curiosity provokes them.

A *straightforward* information-acquisition theory of perception will not do, however. How would such a theory be formulated? Here is a first

Synthese **35** (1977) 257–284.

try. *S perceives O if and only if O causes a percept of S which gives rise to, or generates, at least one true belief about O.*[1] One difficulty is that this does not provide a necessary condition for perception; or so it seems. There are, or seem to be, cases in which no true belief is formed about a perceptual object. In one kind of apparent counterexample, no belief at all is formed about the perceptual object. In Pitcher's desert traveller case (Pitcher, 1971), a man travelling in the desert sees a distant oasis, but, having been fooled before, thinks he is just hallucinating. Although he doesn't believe there is a pool of water, or anything like that, ahead of him, he does see the oasis. This example is not conclusive. Pitcher would reply that the man forms at least a suppressed belief (or suppressed inclination to believe) that water is ahead of him. If we construe 'belief' broadly, the case may not be a counterexample. Perhaps other cases, however, would show that the analysis is too strong. They may be cases in which *some* belief about *O* is generated, but no *true* belief. Seeing an object in a house of distorting mirrors, you may mistake its location, shape, color, and all other properties. Hearing a muffled sound, and misjudging its direction, you may form no true belief about its source. Glimpsing an object briefly, with a misleading mental set, your perception of it may be wholly false.

I shall not try to show conclusively that true belief formation isn't necessary for perception. Even if it is necessary, it isn't sufficient; so the analysis is still in trouble. In the tree-sun case, Smith may well form true beliefs (indeed, 'non-inferential' beliefs) about the sun. By looking at the tree, he may detect the approximate position of the sun in the sky. Since the sun is one of his percept's causes, the analysis mistakenly implies that Smith perceives the sun.

To avert this difficulty, the information-acquisition approach could be revised as follows. *S perceives O if and only if O causes some percept of S which generates more true beliefs about O than about any other cause of the percept.* The rationale for the revision is, of course, that although Smith may form one or two true beliefs about the sun, he will surely form more true beliefs about the tree. But is this necessarily so? If the tree is at some distance, and Smith is almost blinded by the light, he may discern very little about the tree and acquire more information about the sun.

An instructive example for examining the proposed analysis is a slightly elaborated version of Grice's two-pillar case (Grice, 1961). Jones is looking in the direction of pillar *P*, but his view of *P* is blocked by a mirror (of which he is unaware). The mirror reflects the image of pillar *P**, qualitatively identical to *P* but off to the side. Things look to Jones just as they would look if the mirror were not there. As a result of his visual experience Jones acquires a belief, or set of beliefs, that there is a pillar of roughly such-and-such a height, such-and-such a shape, and such-and-such a color directly ahead of him. Suppose further that *P*'s presence is a cause of *P**'s presence. (A mechanical contraption ensures that *P** comes out of the ground into its current position if and only if *P* moves into its current position). Hence, by the transitivity of causation, *P* is a cause of Jones' percept: *P*'s presence causes *P**'s presence, and *P** (via the mirror reflection) causes Jones' percept.

Which pillar does Jones perceive? The correct answer is *P**; but does the proposed analysis yield that answer? More of the descriptive content of Jones' beliefs is satisfied by *P* than by *P**. Not only does *P*, like *P**, have the right height, shape, color, and distance, but also, unlike *P**, it is *directly* ahead of Jones. So it looks as if Jones forms more true beliefs about *P* than about *P**, and this incorrectly implies that Jones perceives *P*.

A way out of the counterexample is to reply that Jones' beliefs aren't *about P*; they are *about P**. That is why the analysis survives the putative counterexample. This reply won't do. We need an account of the conditions in which a belief, or set of beliefs, is *about* a given object, if we are going to be able to apply the proposed analysis. In other words, we need an account of *de re* belief. But an account of *de re* belief, I suspect, presupposes an analysis of perceptual objects. The reason we say that Jones' beliefs are *about P**, for example, is that *P** is the perceptual object of Jones' visual percept. The notion of perceptual objects, then, is conceptually prior to the notion of *de re* belief; so it is illegitimate to use *de re* belief in an analysis of perceptual objecthood.

To avoid the *de re* construction, the analysis might be revised once more. Suppose that percepts give rise to (or are constituted by) sets of (*de dicto*) beliefs in existential propositions, the matrices, or open sentences, of which can be *satisfied by* one or more objects. The analysis, then, may be rewritten as follows. *S perceives O if and only if O causes a percept of S which generates existential beliefs whose matrices are satisfied more by O than by any other cause of the percept.* The defectiveness of this analysis is evident. In the two-pillar case *P* satisfies more of the relevant matrices than *P** does. But *S* sees *P**, not *P*.

I conclude that a *direct* information-acquisition approach to perceptual objects is misguided. I propose a more complicated approach, in which true belief plays an important but *indirect* role.

II

Perception is always perception *in a modality*. We never *just* perceive; we either see, or hear, or smell, etc. A theory of perceptual objects must therefore be a theory of our perceptual modality concepts. The heart of my theory is briefly encapsulated in the following rational reconstruction of such concepts.

Suppose someone *S* (or some organism, or group of organisms) has percepts of kind *Q*, where *Q* is defined by 'qualitative' similarity. These percepts characteristically generate beliefs concerning objects in the environment, and frequently the percept-generated beliefs are uniquely satisfied by an environmental object. For example, a given *Q*-percept of *S* generates a belief that there is currently an object at location *L* with properties *F*, *G*, and *H*; and there is precisely one object at *L* having *F*, *G*, and *H*. Now we consider all such cases with which we are acquainted. We notice that in the great majority of these cases the objects uniquely satisfying the *Q*-generated beliefs bear a certain relation to the perceiver, or to his percept, at the time he has the percept. Since this relation seems to play a special role in producing true *Q*-generated beliefs, we select this relation as *constitutive* of a modality concept. We introduce a modality term '*M*', and we say that all and only objects bearing this relation to a *Q*-perceiver, or to his *Q*-percept, are *M*-perceived by him. All and only objects in this relation are *M*-perceptual objects, or *M*-objects, of the *Q*-perceiver. Objects not in this relation are not *M*-objects, even if they are causally relevant to a *Q*-percept and even if they satisfy (or uniquely satisfy) a *Q*-generated belief. The nature of the selected relation will be discussed below. (Actually, it will turn out that a family of three relations is involved: environmental, counterfactual, and causal. For simplicity, though, this sketch talks of a single relation.) What should be emphasized here is this. The rationale for choosing a given relation as constitutive of a perceptual modality is its general (though not necessarily universal) correlation with unique belief-satisfaction. But since it is logically possible for this relation to hold between an object, *O*, and a perceiver, *S*, even when *S* forms no *Q*-generated belief which *O* satisfies (uniquely or otherwise), it is logically possible for *S* to *M*-perceive *O* without forming any beliefs which *O* satisfies.

Now in certain cases a *Q*-percept may generate two (or more) beliefs each of which is uniquely satisfied by a different object, where these objects are in different relations to the percept. A single visual percept, for example, may generate a belief that a certain tree uniquely satisfies and another belief that the sun uniquely satisfies, though the tree and the sun, on this occasion, bear significantly different relations to the percept (or to the perceiver). How do we decide, then, which relation is to be constitutive of the *Q*-percept modality? We focus on *Q*-percepts in general, not on any specific case. If *Q*-generated beliefs are satisfied (or uniquely satisfied) far more frequently by objects in one relation than by objects in any other relation, this relation is selected as the sole constituent of the *M*-modality. If *Q*-generated beliefs are satisfied very frequently by objects in each of two (or more) relations – say, in roughly comparable amounts – then *both* relations are constituents of the *M*-modality. Thus, an object in *either* of these relations is said to be *M*-perceived.

Let me illustrate and elaborate the foregoing model by reference to vision. Consider the beliefs that are typically generated by visual percepts. When these beliefs are uniquely satisfied by particular objects, which objects are they? That is, what is the relation of these objects, at the time in question, to the subject of the *Q*-percept (the perceiver)? In the vast majority of cases, the *environmental* relation between each object and the perceiver can be characterized roughly as follows. Consider the conical region that extends indefinitely outwards from the perceiver's (open) eyes, in the direction of ocular fixation. (The sides of the cone are the boundaries of his peripheral vision). For each line of the cone traced outwards from his eyes, consider the first, or nearest, opaque object in that line. Next consider the set of all such nearest objects. Finally, consider the subset of these objects which are illuminated at the time in question. I shall say that each member of this subset, or all members collectively, bear the environmental relation '*R*' to the perceiver (or to his visual percept). In other words, roughly speaking, an object *O* bears *R* to *S* at *t* if and only if (a) *O* is illuminated at *t* (or is a source of illumination), (b) *S* has a visual percept at *t*, (c) *S*'s eyes are open at *t*, and (d) *O* is a first opaque object in a line within the conical region defined above. (Notice that an 'environmental' relation such as *R* does not merely pertain to the organism's environment but also, for example, to the orientation of certain of the organism's bodily organs.)

Now in the vast majority of cases in which a visually generated belief is uniquely satisfied by some object, the object is in relation *R* to the subject of the visual percept. For this reason, relation *R* is selected as a constituent of the visual modality.[2] As a first approximation, a person is said to *see* an object (at *t*) if and only if the object bears relation *R* to him (at *t*). (This is by no means my final account. It will be modified and supplemented in several respects). In other words, *O* is a visual object for *S* at *t* if and only if *O* bears *R* to *S* at *t*.

Notice that although we have selected *R* because of its connection with (unique) belief-satisfaction, there is nothing in the conception of *R* that is related to belief-satisfaction, unique or otherwise. Thus, it is logically possible for someone to see an object which does not satisfy (uniquely or non-uniquely) any visually-generated belief of his.

Consider another relation, R^*. An object bears R^* to *S* at *t* if and only if it is outside the conical region that we used to define *R*. (We could define R^* simply as the complement of *R* relative to *S*. This would include objects *behind* a nearest opaque object but within the conical region. The proposed definition, however, makes examples a bit more intuitive.) Now objects in R^* are not said to be *seen*, at least not in primary cases of seeing. (I shall come to secondary cases momentarily.) These objects are not seen even if they uniquely satisfy some visually-generated belief. The reason for this, once again, is that *R* is a much more fecund source of true visually-generated belief than R^*, so we select *R*, but not R^*, as constitutive of the visual modality.

If you need to be persuaded that *R* is a much richer source of true visually-generated belief than R^*, consider the following. Suppose you want to get as much and as detailed information (true belief) as possible, by visual means, about some object. Wouldn't you try to get it into relation *R* to you, rather than R^*? If the object is not a source of illumination, there is virtually nothing you could learn about it by visual means, if it is in R^* to you. But there are innumerable things you could reliably learn about it if it were in R to you (at least if it were close enough and there were enough illumination). Even if the object were a lamp, however, or other source of illumination, R would be a much more fecund source of information than R^*. You can detect much more about even a lamp's shape, size, texture, and relative position by looking at it than by looking at things it illuminates. That it is one lamp rather than two, or that it is a *lamp* at all (rather than a naked bulb), is much more reliably ascertainable in R than in R^*.

We cannot, of course, rest content with the suggestion that R is the exclusive constituent of the visual modality concept. An object's being in R to S at t is not *necessary* for S's seeing it at t. You can see objects in mirrors or on television, when they are not in R to you. R is also *insufficient* for vision, for reasons that will emerge in different parts of the paper. To begin to remedy the situation, let us consider another relation, a *counterfactual dependence* relation. (I use 'counterfactual dependence' in the sense of Lewis, 1973.)

Associated with relation R is a certain counterfactual dependence, a dependence of one's visual percept on the aggregate of objects in R. Suppose a certain set of objects, with specific properties, bear R to S at t. If hypothetical changes were made in these objects at t, there would be certain sorts of changes in S's visual percept. If one of the objects were much larger than it actually is, or had an entirely different shape, there would be corresponding differences in S's visual percept – at least if S were *attending* to the relevant portions of his visual field. If the objects in R were different objects than they actually are – e.g., cars and skyscrapers rather than cows and silos – there would again be corresponding differences in the visual percept. Needless to say, not all differences in the objects, or their properties, would be reflected in differences of percept. But a wide range of differences would be so reflected.

Give the label 'D' to the counterfactual dependence relation associated with R. Now although D is associated with R, it is not associated with R *only*. There are other circumstances in which the same counterfactual relation – or a very similar one – holds between an object and a perceiver (or percept), specifically, cases in which a mirror is appropriately situated vis-a-vis the object and the perceiver, or cases in which the object is shown on television. Admittedly, the counterfactual dependences in reflections and television are not exactly the one associated with *R*. They include left-right reversals, loss of vividness, or loss of the third dimension. Nonetheless, they may be regarded as degenerate cases of relation *D*. Now relation *D*, so construed, is an integral part of our concept of vision. In those cases where *D* obtains but *R* does not, we still say that the object in *D* is *seen* by the perceiver. Thus, relation *D* is one constituent of the visual modality concept. When an object bears both *R* and *D* to *S*, we have a *primary* case of seeing. When an object bears *D* but not *R* to *S*, we have a *secondary* case of seeing.

I shall not try to describe the counterfactual dependence that comprises relation *D*. Not only is *D* too complex to describe, but all of us (who are sighted) are sufficiently acquainted with it to make such description unnecessary. Moreover, the only way of describing differences in percepts is in terms of the stimulus conditions on which they depend. So any such description is likely to be unhelpful. To pick out dependence *D*, however, it helps to contrast it with another counterfactual dependence involving visual percepts. This is dependence D^*, the counterfactual dependence of a visual percept on objects in relation R^* to the perceiver (objects not reflected in mirrors, etc.)

To illustrate the difference between *D* and D^*, consider a lamp of a certain size, shape, and color. Imagine it being directly in front of you, and your having the appropriate visual percept. Then imagine the differences in the percept you would have if, counterfactually, the lamp were much larger, or very different in shape, or very different in color, or much farther away. Now change the case. Imagine the lamp being behind you. Then make the same hypothetical changes in its size, shape, color, and distance. Although some of the latter hypothetical changes would result in differences in your visual percept, the *kinds* of differences would be entirely different from the differences associated with the first set of hypothetical changes. This difference in kinds of differences illustrates the difference between *D* and D^*.

It must be admitted that the idea of a *single* counterfactual dependence *D* is a simplification. What a percept is like at a given moment depends on more than current environmental stimuli, or receptor stimulation. It depends as well on memories, categories, and expectations that filter incoming stimulation: these contribute to the 'construction' of the percept from the initial stimulation. People with different categories, schemata, or expectations (as well as different degrees of sensitivity in receptive organs) will undergo different changes in percepts for the same hypothetical changes in stimulus conditions. I think, however, that there is sufficient homogeneity among perceivers that, in the interest of simplicity, we may speak of a 'single' relation D. The identity of D, after all, is to be defined in terms of its *contrast* with other visual dependences, such as D^*. Such a contrast exists for all visual perceivers; indeed, roughly the same contrast.

III

As developed thus far, my theory says that our visual modality concept has two constituents: environmental relation R and counterfactual relation D. But neither R nor D, it will be noted, is an explicitly *causal* relation (though one view of causation, to be discussed later, says that the counterfactual dependence guarantees a causal relationship). In what sense, then, is my theory a version of the causal theory?

I am not certain that the causal theory of perception is correct if it embodies the claim that concepts of perception have *always* included a causal component. It seems clear, though, that our *current* concept of perception includes such a component. Arguably, it has always been the case that the ascription of perceptual achievement requires *some* causal connection between object and percept, without specification of the details of the causal connection. But it is also true that, over time, some degree of causal detail has been built into each of the various perceptual modality concepts.

A rational reconstruction of the development of a modality concept might go as follows. Initially, a modality concept is tied to an environmental relation and a counterfactual relation. It is natural, however, to inquire into the causal mechanisms that underlie the counterfactual dependence. Part of the mechanism is normally part of the environmental relation itself, i.e., that part that involves bodily organs. Other causally relevant features are suggested by other conditions involved in the environmental and/or counterfactual relations, e.g., the fact that visual percepts are affected by conditions of illumination, auditory percepts by the condition of the air, etc. Gradually, hypotheses are accepted concerning the details of the causal mechanism, and some of these hypotheses become part of the everyday conception of the modality. It becomes part of the concept of vision, for example, that it involves the transmission of light from objects to the perceiver's eyes.

We have already seen how a counterfactual relation can supersede an environmental relation, how, for example, *D* comes to be a more decisive constituent of vision than *R*. Similarly, once the notion of a certain physical medium, or other causal mechanism, becomes imbedded in a modality concept, it can sometimes override a counterfactual relation in importance. Whether we use a specific modality term to describe the accomplishments of a non-human organism depends more on the causal-mechanism factor than the counterfactual dependence. (In part this is because our access to the percepts of other organisms, and hence their perceptual dependences, is very limited.) Why do we say, for example, that *frogs* see? The visual processing equipment of frogs, as Lettvin, Maturana, McCulloch, and Pitts (1959) have shown, is quite different from ours: it is much more specialized. The counterfactual dependence associated with their 'vision' must therefore be significantly different from *D*, the dependence I have claimed to be (partly) constitutive of vision. Still, we have little hesitation in saying that frogs see (even once we learn of the relevant physiological facts). Apparently, the counterfactual relation is less important than the causal-mechanism relation. Since the frog has organs of photoreception, whose construction is fairly similar to ours, and since he makes *some* of the same discriminations we make with our visual apparatus, we say that he too 'sees'.

Should we conclude that the counterfactual dependence relation drops out of our concept of vision once we have enough information about causal mechanisms? I don't think so. Consider the following case. You begin to have visual experiences that co-vary in a *D*-like way with objects in a room three thousand miles away. When chairs and tables are in the room, you have visual experiences of an appropriate sort. When people are in the room, you have corresponding visual experiences. Etc. Investigation shows that this co-variation is due to some physical medium *other* than light-transmission, and stimulation of organs *other* than your eyes. Despite this difference in causal mechanisms, wouldn't we be inclined to say that you *see* the objects in the room? We might prefer calling it 'clairvoyance'; but, etymologically, clairvoyance is clear *vision*. That we are still inclined to speak of vision, or something like vision, shows that the counterfactual component sometimes dominates the causal-mechanism component.

My conclusion, then, is this. Each modality concept is constituted by a family of three relations. The first member of the family is an environmental relation (such as R), the second a counterfactual relation (such as D), and the third a causal-mechanism relation (e.g., light-transmission culminating in stimulation of photo-receptors). Paradigm, or primary, cases of perception in a modality are cases where all three members of the family are exemplified. When some members of the family are not exemplified, as in reflections or infra-human perception, a decision of whether to attribute perception *in the modality in question* sometimes emphasizes certain family members and sometimes others.

How does my theory apply to the two-pillar example? Which pillar (if any) does Jones see, according to the theory? Well, Jones is not in relation R to pillar P (since P is not a first opaque object in his line of vision). Nor does Jones' visual percept depend counterfactually in manner D on pillar P. It is true (given our hypothesis about the contraption connecting P with P^*) that if P weren't present at all, things would look different to Jones, since pillar P^* wouldn't be where it is. But a change in color of P (we may suppose) would not have any effect on P^*, and hence no effect on Jones' visual percept. And similarly (we may suppose) for differences in shape and texture of pillar P's surface. Finally, there is no process of light-transmission from P to Jones' eyes. For these reasons, Jones does not see P. What about P^*? P^* is not in R to Jones, but this alone does not preclude his seeing it. P^* *does* bear the other relations to Jones, or to his percept. His visual percept depends counterfactually on P^* in manner D (or roughly D); and light is transmitted from P^* to Jones' eyes. So on our theory, Jones sees P^* but not P.

Let me give a rough summary of the theory as presented thus far. (The theory will not be complete until an important addition is made in Section VII.) The theory is two-tiered. The first tier contains the conditions in which object O is perceived by organism S. O is perceived by S (at t) if and only if there is a perceptual modality concept M composed of an environmental, a counterfactual, and a causal-mechanism relation such that O bears enough of (the most important of) these relations to a percept that S has (at t). The theory's second tier explains why *certain* environmental, counterfactual, and causal-mechanism relations, and not others, are selected as constitutive of a modality concept. The relations selected as constitutive of a modality are those that most commonly obtain (or are believed to obtain) between objects and percepts when those objects uniquely satisfy beliefs which the percepts generate (or constitute).

IV

Let us test my theory of perceptual objects on other perceptual modalities, and on one outstanding problem concerning vision. First, consider the tactual, or haptic, modality. When you grasp a rod and touch various objects with it, e.g., a stone, you are said to 'feel' these objects. Why do we say that you feel these objects, as well as the rod itself? One point, of course, is crucial: you characteristically acquire true beliefs about the distal object, just as you acquire true beliefs about objects actually *touched* with the hand. Indeed, once the rod becomes familiar to you, you can acquire *as much* true belief in this fashion as by unmediated contact with an object. What fundamentally underlies our use of the term 'feel', therefore, is true belief acquisition. Beyond this point, however, there are two ways of dealing with the case.

First, one might say that in the rod-mediated case, the stone stands in the *same* (or similar) environmental, counterfactual, and causal relations to you as in the unmediated case. This will mean, of course, that these relations have to be construed sufficiently broadly. The environmental relation, for example, must not require that the perceptual object be *in contact* with some bodily part, but only in contact with some sequence of objects the final member of which is in contact with some bodily part.[3]

An alternative approach is to employ a principle that I mentioned in the original sketch of my theory. This is the principle that if classes of objects in two different (families of) relations to a given kind of percept satisfy percept-generated beliefs in roughly equal amounts, then *both* (families of) relations become constituents of the modality. Objects in *either* family of relations to a percept are said to be *M*-perceived. Now we might say that a stone in a rod-mediated case is *not* in the same family of relations to the perceiver as in an unmediated case. Yet *both* families of relations are 'approved' perceptual relations; both are constituents of the tactual modality.

This principle seems to be required to handle audition and olfaction. In both of these modalities, a person is said to perceive objects that are (apparently) in different families of relations to the perceiver. A person is said to hear a *bell*, and also the *sound* of the bell. He is said to smell a *skunk*, and also the *odor* of the skunk. It seems plausible to say that the bell and its sound, and the skunk and its odor, are respectively in different families of relation to the percept. *Each* family, however, is a constituent of the auditory and olfactory modalities, respectively, because objects in each family satisfy percept-generated beliefs in roughly equal amounts as objects in the other family of relations to the same kind of percept. In other words, we form roughly equal amounts of true belief about sounds and sound-sources, and about odors and odor-sources.[4] (Here and elsewhere I use the preposition 'about' for convenience. It can be replaced by our official terminology where necessary.)

Let us return to a remaining problem in the visual modality, the problem of proximal arrays of light. Why aren't we said to see these arrays, e.g., arrays of light two inches from our eyes? A large part of the answer is, of course, clear. We seldom form percept-generated beliefs that are satisfied uniquely by these proximal arrays. We seldom have beliefs of the form: 'There is an array of light at such-and-such a distance from me with properties *I*, *J*, and *K*'. The beliefs that spontaneously arise from, or are constituted by, visual percepts are beliefs concerning *distal* objects (normally physical bodies), not *proximal* ones (at least not proximal rays of light).

Why, it may well be asked, do we standardly form beliefs concerning odors and sounds, but not concerning proximal arrays of light? Why do we, as perceivers, think of odors and sounds as mediating our access to their sources, but do not think of light as mediating our access to physical objects? (Of course we *do* so think of light in our scientific moments, but this does not pervade our intuitive thinking very much.) The answer, I think, runs as follows. Odors can often be smelled after their sources have departed, and the same holds, to a lesser extent, for sounds. An odor of a skunk remains after the animal has disappeared. The echo of my voice is heard after I have finished calling. These cases are evident to common observation, and it forces us to countenance mediating 'objects' of audition and olfaction, objects other than physical bodies. We have no comparable experience of objects that mediate our visual access to physical bodies. Admittedly, the same phenomenon occurs when we receive the light from a distant star, which may since have exploded. But common experience gives us no awareness of this phenomenon. So we do not intuitively conceive of our visual access to the star as mediated by another object.

Another kind of problem about light, however, emerges at this juncture. Although proximal arrays of light are not said to be seen, we do sometimes see (distal) light, e.g., the aurora borealis. A problem for my theory is this. It would seem that a proximal array of light is in the same family of relations – environmental, counterfactual, and causal – as a distal array, such as the aurora borealis; it is only *closer*. According to my theory, therefore, we should either be said to see both or be said to see neither. Why, then, are we said to see the aurora borealis, but not the proximal array?

To answer this question, I need a more precise rendering of my theory. In trying to explain the sorts of perceptual attributions we make, what is strictly relevant is not what environmental, counterfactual, and causal relations *actually* obtain, but what relations we standardly *believe* to obtain. Whatever the scientific facts concerning proximal and distal arrays of light, we do not ordinarily conceive of, say, the aurora borealis as being in the same set of relations to us as arrays of light two inches from our eyes. We think of the aurora borealis as being constituted by colored light in certain locations, and of our visual percepts as depending, in a *D*-like way, on these colors. We do not so think of the light two inches from our eyes (in general, we do not think of it at all). It is certainly arguable that our thought on this subject ought to be reformed, and our language of perception along with it. But the primary task here is to understand our intuitive conceptual scheme, not to reform it.

Let me examine a few more cases to confirm my general theory. First, recall our bat example. Why do we think of the bat as perceiving objects in his path, but not himself? Given the bat's flight behavior, it is natural to attribute to him beliefs that are satisfied by obstacles in his path, such beliefs as 'There is a dense object occupying such-and-such a region', 'There is an insect at such-and-such a spot', etc. Since these beliefs are elicited by the sonar signals that are caused, in part, by these belief-satisfying objects, we take the family of relations between this class of objects and the bat to be constitutive of his echolocation modality. Since the bat does not bear those relations to himself, we do not think of him as perceiving himself.

It might be pointed out that we *could* construe the bat as forming beliefs about his own location or change in location relative to the environment. These would be beliefs *he* satisfies. But it doesn't matter what beliefs we *could* attribute to him; what matters is only the beliefs we *do* (generally) attribute to him. I contend that the kind of objects we select as perceptual objects is determined (ultimately) by the beliefs we *actually* ascribe to perceivers. The bat example sustains this contention as long as the beliefs we actually attribute to the bat are primarily beliefs of the sort satisfied by external objects, not by himself. This seems to be correct. The beliefs we attribute to the bat are ones in which objects are located *at some distance* from a (tacit) point of origin, viz., himself. These beliefs are satisfied (if at all) by objects other than himself.

Another example that confirms our theory involves rheumatic aches. Rheumatic aches are sometimes caused by drops in atmospheric pressure, drops that characteristically precede rain. A person who first experiences such aches would not be said to 'perceive' anything, at least no external state of the world. But if he gradually associates the aches with the onset of rain, we shall come to think of him as *perceiving* or *sensing* the onset of rain. He may not know enough to have beliefs about changes in atmospheric pressure. But he has beliefs like: 'Something is happening that will lead to rain'. The event that satisfies this belief, and also causes his ache, is the drop in atmospheric pressure. Hence, we say that he *perceives* such drops. We will say that he perceives a fall in atmospheric pressure even if his belief, on a particular occasion, happens to be false (because it doesn't rain). The atmospheric pressure is in the right (perceptual) relation to his ache, even though, on this occasion, it doesn't satisfy the ache-generated belief.

V

I have made a few *obiter dicta* about *why* certain beliefs are or are not formed. These remarks are not strictly part of the theory of perceptual objects. The theory is entitled to take it for granted that certain classes of belief are or are not formed; it need not *explain* these facts. Still, it is empirically interesting to consider why organisms form beliefs that are satisfied by certain classes of objects and do not form beliefs, or form fewer beliefs, that are satisfied by other classes of objects. Why, for example, do visual percepts generate more beliefs that are satisfied by R-and-D-related objects than beliefs that are satisfied by R^*-and-D^*-related objects? And why don't we form any beliefs (at least no spontaneous or 'non-inferential' beliefs) that are satisfied by proximal arrays of light?

Let us concentrate on people's superior ability to form beliefs that are satisfied by R-and-D-related objects, as opposed to beliefs that are satisfied by R^*-and-D^*-related objects. Part of the explanation is purely physical, as opposed to physiological. It concerns the relations between the classes of objects in question and the energy that impinges on the receptors. A necessary condition for *systematic* formation of true perceptual belief – not merely occasional lucky guesses – is that differences among objects in the environment be mapped into differences of retinal excitation. Suppose there are a number of alternative possible states of an environment, each with equal prior probabilities. A person cannot reliably make a doxastic choice among these alternatives on the basis of background beliefs. He needs input from the environmental state that actually obtains, input that distinguishes this state from the other possible alternatives. In short, this state must produce distinctive receptor stimulation. If, however, the person's position relative to the environmental objects is such that the same receptor stimulation would result from any of the alternative states of the environment, there is no reliable way for him to form a true belief concerning which of these states actually obtains.

Now this is roughly a person's position vis-a-vis objects that bear the R^*-and -D^* relation to him. For any given retinal stimulation, large numbers of alternative sets of R^*-and-D^*-related objects are compatible with this stimulation. Objects that don't emit light, and have no distinctive light-reflectant properties, make equivalent contributions to retinal stimulation. The retinal pattern would be the same, or essentially the same, whether the objects are chairs or tables, elephants or wildebeest. Thus, S cannot *reliably* form a belief, as a result of retinal stimulation, that a table or a wildebeest is in some R^*-position relative to him. Even sources of illumination present problems. The ambient array of light impinging on S's retinae may be the same whether the R^* region contains two lamps or three lamps with the same total luminance. S has no reliable basis for belief about the number of lamps in his (R^*) vicinity.

In short, and to be more general, the ability of an organism to form beliefs that are (frequently or systematically) satisfied by objects in a certain relation to him depends on the *mapping* from classes of objects in this relation to receptor stimulation. The '*richer*' this mapping – that is, the smaller the number of different possible states that are mapped into one pattern of stimulation – the greater the potential for true belief formation.[5]

Differences at the receptor level, however, do not ensure differences at the level of the percept. It is only at this higher level that beliefs are generated (or constituted). No actual perceptual processing system can retain, or make use of, all differences at the receptor level. As Cornsweet stresses, perceptual processing always involves the *loss* of information. (Cf. Cornsweet, 1970, p. 379.)

Visual processing is a good case in point. The phenomena of brightness contrast and brightness constancy imply the loss of information about absolute light intensity. Take brightness contrast, for example. Due in part to lateral inhibition, experienced brightness does not depend so much on the absolute light intensities at regions of the retinae, but on differences or ratios of intensity between one retinal region and adjacent retinal regions. The apparent brightness of an R-related point in space does not depend so much on the amount of light it transmits, but on the amount it transmits *in comparison with* adjacent points in space. Since differences in reflectance typically occur at the edges or borders of physical objects, there is a consequent 'enhancement' at the level of the percept of the edges or borders of R-related objects. This is helpful for detecting the shapes and shape-related properties of R-related objects, because the contours will tend to stand out. But the loss of information about absolute light intensity (at the level of the percept) is ill-suited for detecting properties of R^*-related objects. Information about absolute light intensity is one kind of information that could help generate true belief about light sources in R^*. The relevant physiological mechanisms of vision are ill-suited for this end.[6]

A physiological mechanism like lateral inhibition is adaptive, in the sense that beliefs about, or discriminations among, physical objects are biologically or ecologically more significant for the organism than beliefs about, or discriminations among, absolute states of illumination. It is biologically useful to sacrifice information about absolute levels of illumination to increase the information – i.e., the *registering* of the information – about physical objects. That visual systems are primarily evolved for registering information about physical objects, rather than about light per se, is indicated by the presence of another kind of photo-receptive mechanism that *coexists* with eyes in reptiles, fish cephalopods, and amphibians. (Cf. Bower, 1974.) These so-called 'third eyes' function to respond to light, and nothing else; real eyes do not function this way at all.

These remarks obviously apply to the question of why we don't form beliefs about proximal arrays of light. The interpretation of retinal stimulation in terms of distal physical bodies rather than proximal light clearly has a biological function. To survive, the organism can profit from an accurate representation of the presence or absence of predators and prey; it doesn't need a representation of the light two inches from its eyes.

VI

Given the importance of beliefs and percepts to my theory, some remarks about their relationship are in order. I have generally spoken of percepts as 'generating' beliefs, and this may suggest a causal relationship. It would be inaccurate, however, to think of percepts as only *causing* beliefs; some percepts themselves have belief-content. This content, moreover, is what plays the really critical role in our theory. I have claimed that vision, for example, does not generate beliefs that are satisfied by proximal arrays of light. But that is not strictly true: physicists and psychologists of perception can and do form beliefs – fairly detailed beliefs – that are satisfied by such arrays of light. The point is that these are only 'inferential' beliefs. They are not *part of* the visual percepts in the sense that it 'looks as if' there are arrays of light two inches before the eyes. (I do not preclude the possibility that beliefs of that sort could arise. If this happened, it would be natural to begin speaking of these persons as seeing the proximal arrays of light.) It is important to our theory, therefore, that the beliefs in question are *'embedded'* in percepts, not merely caused by them. If beliefs can be embedded in percepts, however, how shall we conceptualize percepts? Are percepts *simply* sets of beliefs, as some philosophers maintain? Or do they have some other status?

A view of percepts as beliefs or inclinations to believe has been defended by Armstrong (1961) and Pitcher (1971), and goes back to the Idealists (cf. Hirst, 1959, p. 222). I think there is much to be said for this view, although for reasons to be explained I hesitate to accept it outright. Two initial clarifications should be made, however, to give the view even initial plausibility.

An inevitable source of dissatisfaction with the Armstrong–Pitcher line is its apparent attempt to *eliminate* the sensuous or qualitative aspect of percepts. Of course, they would say that this qualitative aspect is not eliminated, but is *captured by* belief content, or at least belief-content coupled with the *mode of acquisition.* The differences between the percepts of the different modalities would be understood, on this view, in terms of the manner of acquisition, i.e., the sense-organs employed. But this will not do. Blind people often have an 'obstacle sense' involving percepts that are qualitatively like pressure against the face. These percepts, as it turns out, are acquired through the ears, just like auditory percepts. (Cf. J. J. Gibson, 1966.) To distinguish them from auditory percepts, one cannot appeal to mode of acquisition; one must appeal to qualitative difference.

A more plausible version of a doxastic theory of percepts would distinguish the *content* of a mental event from the *form* or *vehicle* in which this content is encoded. Any content must have *some* manner of encoding. Perceptual or sensory *qualities* are simply distinctive manners of encoding. (Even 'intellectual' beliefs may have a quasi-perceptual embodiment: either quasi-auditory or quasi-motor, i.e., articulatory.) One can accept such qualities and still maintain that each particular qualitative event has cognitive, or doxastic, content.

Philosophers like Armstrong and Pitcher want to reduce sensory qualities to beliefs because of the ontological and epistemological problems they allegedly pose. But neither the ontological nor the epistemological motivation is well grounded. The ontological status of sensory qualities may indeed be problematic, but no more problematic than the status of beliefs. There is little to be gained by reducing such qualities to beliefs. Concerning the epistemological problem, I believe one can be a direct realist without denying the qualitative distinctiveness of percepts: as long as one doesn't confuse percepts – i.e., perceptual *states* of the organism – with sense-data.

Second, it should be admitted that percepts are not propositional in form. Their content may be translated, with greater or less accuracy, into, say, existential propositions, but they do not really have a quantificational structure. To say they are 'iconic' modes of representation may be correct, but, pending some acceptable rendering of 'iconic', isn't very illuminating. Cognitive psychology has witnessed some interesting arguments in support of non-verbal modes of information-processing (cf. Shepard and Metzler, 1971, and Paivio, 1971.) But the conceptual issues here remain to be settled (cf. Pylyshyn, 1973).

If percepts are not propositional, how can they be belief-like at all? The content of beliefs must be classifiable as true or false, correct or incorrect. But only propositional content can be so classified. This problem goes beyond our purview, but here is one possible answer. Suppose a percept involves, among other things, a set of motor plans or dispositions, dispositions, say, for touching or handling the perceptual object. (Such a view is advocated by 'motor copy' theorists of perception, and is hinted at by Piaget. For a survey, see E. J. Gibson, 1969.) Such plans or dispositions can have a tendency either for 'success' or 'failure': a plan for touching the object might be well-designed or ill-designed for touching it. The correctness or incorrectness of a percept, then, could consist in its propensity to produce successful or unsuccessful motor activity.

Resolution of these two issues is preliminary to a defense of a doxastic view of percepts. What positive considerations favor such a view? Here are a few. First, even if the motor account of percepts is not right in general, it is clear that many, if not all, percepts have a *behavior-guiding* function. (Potential) behavior-guidingness, moreover, is an essential characteristic of belief in general. Consider kinesthetic sensations, which we may regard as a species of percept. When playing the piano, the feeling of the fingers as being certain distances apart guides one in changing their relative position for the next chord. Successful finger deployment does not require a non-sensory *thought* in addition to the kinesthetic *feeling*; the kinesthetic feeling embodies the information about the present position of the fingers. Similarly, the vestibular feeling of the body as being oriented thus-and-so relative to gravity *is* a belief concerning one's orientation; at least it guides one's bodily movements in the way that beliefs in general guide behavior.

Second, many percepts have 'abstract' content. I may see an object not merely as having a certain size and shape, but as a hammer, a table, or a tree. Such abstract categories may be more prominent in the percept than spatial or qualitative features.

Third, even spatial or qualitative representations may be regarded as essentially cognitive. What is fundamental to cognition is classification; and spatial or qualitative representations are simply particular kinds of classification. That the basic function of perception is classificatory may be inferred from physiological and psychological investigation, if not from introspection. (A) Perceptual systems, according to current theory, are composed of various *feature analyzers*, designed to detect specific properties of external stimuli. (The work of Hubel and Wiesel, 1962, is pioneering in this domain.) (B) Perceptual processes operate so as to *organize* sensory input into meaningful categories. (See, for example, Lindsay and Norman, 1972.) (C) There is evidence of pre-established, discrete categories into which continuous sensory data are grouped. Specifically, both hue perception and speech perception are characterized by a division of the continuous physical input into discontinuous categories. (Cf. Bornstein *et al.*, 1976, Eimas *et al.*, 1971.) A single (e.g., visual) percept involves so many classifications that it is impossible to map them all into some accessible set of verbal classifications. But that should not be allowed to obscure the fundamental classificatory nature of perception.

Fourth, percepts are influenced by higher forms of cognition, and sometimes behave in ways akin to higher processes. For example, the perceived segmentation of a discourse depends on one's knowledge of the language and one's expectations about what can or will be said. (Cf. Halle and Stevens, 1962.) Next consider perceptual adaptation. After wearing inverting spectacles, subjects with appropriate learning opportunities undergo reinversion of images. (Cf. Kohler, 1964.) The general cognitive system of the organism seems bent on securing a coherent or unified representation, by coordinating its visual with its tactual representations. This is similar to the readjustment of (non-perceptual) beliefs in the attempt to preserve consistency.

Against these considerations, the following brief rebuttals might be made.

Many of the foregoing arguments could at best show that *some* percepts have cognitive components, not that *all* do, or that percepts *necessarily* have belief-content. Moreover, it is highly problematic whether all percepts have such content, or whether every aspect of every percept has such content. There is the traditional problem of the perceptual experience of the newborn, and of the congenitally blind person after cataract surgery. Although sensory input may provide the potential for classification and organization, it is not clear that perceptual experience at these stages yet manifests this potential. Certainly behavioral control does not set in immediately. And many theorists would argue that the potential for classification and discrimination does not get actualized until the organism learns to *attend* to relevant segments of its sensory data. (Cf. E. J. Gibson, 1969.) Even in the mature individual, not all information that is present at lower levels is grasped in focal consciousness. Current theory views perceptual processing as involving various stages through which 'information' is filtered or channeled. Certain material is fully analyzed, but other, unattended material is only briefly retained and not fully digested. (Cf. Sperling, 1960 and Treisman, 1964. For reviews, see Neisser, 1967 and Norman, 1969.) This attenuated material, I would claim, manifests itself in the percept as a sort of blur, and a subsequent feeling that a lot was 'there' that wasn't fully noticed. Having not noticed, or retained, this material, it cannot be regarded as a *belief*, or as having belief-content. This shows that *some* characteristics of percepts do not have belief-content. Perhaps 'undeveloped' percepts are entirely like this, and therefore have no belief-content at all.

Another kind of criticism of the doxastic theory is that some percepts are not full-fledged beliefs, because what is suggested or intimated by the percept is rejected or ignored at a higher (cognitive) level. When I see a stick half immersed in water, the stick *looks* bent even if I don't *believe* it is bent. This point, however, can be conceded by the doxastic theorist. He can reply that percepts are always kinds of cognitive *input* into a larger cognitive system. This input may sometimes be suppressed, ignored, or rejected, but is still a form of *cognitive* input. Moreover, in ordinary cases it is not *merely* 'input'. Unless it is explicitly suppressed or rejected by the 'central clearing-house', it functions *as* a belief.

It is not crucial to our theory to settle this issue definitively. It is congenial to an information-acquisition theory of perceptual objects to have a doxastic view of percepts. But my theory of perceptual objects does not require the *strongest* form of a doxastic view of percepts. As long as percepts *frequently* contain cognitive content – as long as *enough* aspects of percents are beliefs or belief-like entities – this should safeguard the effort to identify perceptual objects (indirectly) in terms of belief-satisfaction.

It might be suggested that a strong form of the doxastic view of percepts would be *un*congenial to my theory, because it would favor the *direct* information-acquisition theory of perceptual objects which I rejected. This suggestion is based on the recognition that the more belief-content is attributed to a percept, the harder it is to find or conceive of cases in which a percept fails to have *any* belief-content that the perceptual object satisfies.

The most this could show, however, is that belief-satisfaction is a *necessary* condition of being a perceptual object. (Nor does even this much follow. Though percepts may be replete with belief-content, there may be cases in which none of the belief-content is *true of* the perceptual object.) It does not bear on the question of *sufficient* conditions for perceptual objecthood. As we saw earlier, a *direct* information-acquisition theory of perception founders on this problem.

VII

My theory of perceptual objects is not yet complete. Consider a small rock on a distant mountain I am viewing. This rock bears relation R, relation D, and the appropriate light-transmission relation to me. But I don't see it. It is true that the rock is an (R-related) element of a larger entity, the mountain, which I see. Still, I don't see the rock. The general point can be made as follows. Let O^* be the *aggregate* of stimuli that jointly bear the appropriate relation, or family of relations, to S's Q-percept at t, and let O be a sub-unit of O^* that also bears this relation to S. Although our theory correctly implies that O^* is M-perceived by S, it incorrectly implies that O is also perceived by S. To correct this implication, we need to say something more about sub-units.

The right move seems to be the one Dretske makes: a sub-unit is perceived only if it is *differentiated* by S from its immediate environment (Dretske, 1969). If a piece of beige paper is glued to a beige wall, a person sees the paper only if he differentiates it from the (adjacent portions of) the wall. A person eating a carrot–orange soup can be said to taste the orange only if the orange flavor is differentiated by him from the other component flavor. A person listening to an orchestral passage can be said to hear the violas only if he differentiates the sound of the violas from the sounds of the other instrumental sections.

There are two problems with this approach. The first, which I shall not try to solve completely, is what 'differentiates' means. The best effort Dretske makes is to say (for the case of vision) that *S* differentiates *O* from its immediate environment just in case *O* '*looks different*' to him than the objects in the immediate environment. What does 'looks different' mean? Does it mean that *O* appears to have some property that the other objects in the immediate environment do not appear to have? This requirement is not strong enough. There is a property which the piece of paper appears to have that the adjacent portions of the wall don't appear to have, viz., occupying its particular location. Yet *S* does not see the piece of paper.

The second problem is also suggested by an example of Dretske's. If we select a reasonably large portion of a perfectly homogeneous wall that you see, you would also be said to see that *portion*, though it isn't differentiated from adjacent portions. So the differentiation condition, even if properly explicated, is not universal.

I would approach the two problems as follows. A potential perception-attributor divides the environment into *natural units*, that have certain discontinuities or inhomogeneities. A piece of paper is a distinct unit from a wall, but the several portions of a wall are not distinct units.[7] With respect to a natural unit *O*, *S* is not said to perceive *O* (where *O* is a sub-unit of the aggregate perceptual stimulus) unless something in *S*'s percept represents its distinctness, that is, unless something in the percept is a clue to there being such a distinct natural unit present. (I say a 'clue' because it would be too much to require that *S believe* there is such a natural unit.) A representation of distinctness might be a representation of an edge or boundary, or of a distinctive flavor. Mere difference in positional representation, though, does not constitute such a representation, because it isn't a clue to a distinct natural unit.

The above condition only applies to the perception of natural units. A portion of a wall is said to be perceived even when it is not differentiated from its environment because it is not a distinct natural unit.

With respect to natural units, perceptual verbs are, in Goodman's terminology, *non-dissective* (Goodman, 1951). That is, where O is a proper part of O^* and a natural unit, the fact that O^* is M-perceived by S (at t) does not entail that O is M-perceived by S (at t).

The non-dissectiveness of perceptual verbs partly explains the oddity of saying that S sees a mass of molecules (when, for example, he sees a table). In interpreting 'S sees a mass of molecules' there is a temptation to read 'mass of molecules' distributively, so that the sentence implies, 'S sees at least one molecule'. But this sentence is false. Given the non-dissectiveness of 'see', S may see something made up of molecules without seeing any molecule. Since a molecule is a 'natural unit', he would have to differentiate one molecule from adjacent objects. But no one can do this with the naked eye. Of course, one can interpret 'mass of molecules' non-distributively. Such a reading is mandatory when the foregoing sentence is rewritten as: 'S sees something that is in fact a mass of molecules'. This sentence, however, is neither odd nor false.

VIII

Let me close with a difficulty for the *causal* component of my theory. Consider the following example (which I owe to Jaegwon Kim). You are looking at the night-time sky through a telescope, and detect a region of total blackness against a background of diffuse light. This portion of your phenomenal field, it turns out, is attributable to a black hole. Wouldn't we say that you *see* the black hole? But there is no light-transmission from the black hole to your eyes; nor is any other energy transmitted from it to your eyes. We seem to have a counterexample to the causal theory of perception.

Is it really a counterexample? The best approach is to distinguish two senses of 'cause'. (Like other philosophers, I try not to multiply senses; but sometimes it is natural.) One sense of 'cause' is a *counterfactual* sense, the second a *physical mechanism* sense. The counterfactual sense of 'cause' is roughly the sense captured by David Lewis' analysis (Lewis, 1973), or a somewhat stronger analysis along those lines. In this sense, it is sufficient for C to cause E (where C and E are events) that E wouldn't have occurred if C hadn't occurred. The counterfactual sense of 'cause' is satisfied in the black hole case. Your percept wouldn't be the same (you wouldn't have *that* percept) if the black hole weren't where it is. If the causal theory of perception requires causation only in the *counterfactual* sense, the black hole is not a counterexample. It is a counterexample, though, if the theory involves the physical mechanism sense of 'cause'. Such a sense would say that an event (or object) causes another event only if there is some transfer of energy or force from the first event to the second. This requirement is violated in the black hole case.

Now the theory of the several modalities (which belongs to the theory of perception) seems to need the physical mechanism sense of 'cause'. Modalities are partly distinguished by *kinds* of causal mechanisms. If *no* physical mechanism is at work, how can we say there is *vision, hearing*, or what have you? We could dismiss this as a case where environmental and counterfactual components of the modality-concept take precedence over the causal-mechanism component. But I think another solution is preferable.

Even in the black hole case, the quality of your percept is partly attributable to light transmission. There is no light transmission from the black hole, but light transmission from other regions partly explains the salient blackness in the middle of your visual field. The salient blackness results from a *difference* or *contrast* in light-transmission: *some* light from adjacent regions and *no* light from the black hole. Since the black hole is one of the objects responsible for this contrast (in the counterfactual sense of 'responsible'), it satisfies all that should properly be required by way of a causal constraint on objects of vision.

A virtue of this 'contrastive' approach to the causal requirement is that it accords with the way perceptual systems work in general: they respond to *differences* in light intensity or *differences* in pressure, rather than absolute levels of intensity or pressure. (Cf. Lindsay and Norman, 1972, p. 111, and Bower, 1974.) An example of this is the phenomenon of habituation. Sensory systems in general become habituated to recurrent patterns of stimulation, so that these patterns no longer reach the level of consciousness. Only changes or differences in stimulation are registered at the perceptual level. We stop hearing the ticking of a clock or the hum of an air-conditioner, but notice the *cessation* of such a pattern. Cornsweet (1970) tells an illustrative story. At 11:46 in the evening, an electrical failure disabled Big Ben's chimes so that they could not sound. At 12:00 sharp, the man who lived next door leapt out of bed and shouted: 'What was that?' It is tempting to say that there was *something* this man *heard*: the non-ringing of the chimes. Of course, this non-ringing did not cause the man's auditory percept via sound-wave transmission. But the non-ringing was causally responsible for his auditory percept because it (among other things) was causally responsible (in the counterfactual sense) for the *difference* between *past* sound-wave patterns on his ear-drums and the *current* sound-wave pattern (i.e., *no* impingement) on his ear-drums.

The University of Michigan

NOTES

* This paper was written while the author was a John Simon Guggenheim Fellow and a Fellow of the Center for Advanced Study in the Behavioral Sciences. I am indebted to Holly S. Goldman for a number of valuable criticisms and suggestions. I have also received helpful comments from several Center Fellows, including Keith S. Donnellan.

[1] Perhaps I should say, here and elsewhere, '*non-trivial* true belief'. This would exclude beliefs such as '*O* is an object'.

[2] It is not clear how we pick out exactly relation *R*, rather than some slightly more or slightly less inclusive relation. Clearly, considerations of simplicity and naturalness are involved; but I shall not speculate on the details of these considerations.

[3] Or, one might treat the rod-mediated case like mirror-vision. The standard environmental relation is not satisfied, but the counterfactual and causal relations are satisfied.

[4] It should not be concluded, of course, that *all* objects causally relevant to an auditory or olfactory percept are said to be heard or smelled. As in vision, some causes are not perceptual objects. An auditory percept is causally dependent on the acoustic properties of the environment. Whether a sound is heard at all, and how loud or lush it sounds, depends, for example, on the curvature of the ceiling and on the absorbent properties of the furnishings. (Designers of whisper chambers and orchestra halls take such things into account.) But people are not said to *hear* the ceiling or the furniture. Again, this is explicable in terms of true belief acquisition. We seldom form percept-generated beliefs about the acoustic properties of the environment, such as the ceiling or the furniture.

[5] This discussion echoes themes I have stressed in Goldman (1976). In several respects the present paper is a companion piece to that one.

[6] In these passages the term 'information' does not usually mean 'true belief', but has a connotation derived from information theory. See the introduction to Cornsweet (1970).

[7] Which units are 'natural' units is a problem broached (in somewhat different terms) by Roderick Firth, in Firth (1967). I shall not attempt to resolve this problem.

REFERENCES

Armstrong, D. M., *Perception and the Physical World*, London, 1961.

Bornstein, M. H., W. Kessen, and S. Weiskopf, 'The Categories of Hue in Infancy', *Science* **191**, (1976).

Bower, T. G. R., 'The Evolution of Sensory Systems', in R. B. MacLeod and H. L. Pick, Jr., eds., *Essays in Honor of James J. Gibson*, Ithaca, 1974.

Cornsweet, T. N., *Visual Perception*, New York, 1970.

Dretske, F. I., *Seeing and Knowing*, Chicago, 1969.

Eimas, P. D., E. R. Siqueland, P. Jusczyk, and J. Vigorito, 'Speech Perception in Infants', *Science* **171**, (1971).

Firth, R., 'The Men Themselves; or the Role of Causation in Our Concept of Seeing', in H. N. Castañeda, ed., *Intentionality, Minds, and Perception*, Detroit, 1967.

Gibson, E. J., *Principles of Perceptual Learning and Development*, New York, 1969.

Gibson, J. J., *The Senses Considered as Perceptual Systems*, Boston, 1966.

Goldman, A. I., 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* **73**, (1976).

Goodman, N., *The Structure of Appearance*, Cambridge, 1951.

Grice, H. P., 'The Causal Theory of Perception', *Proceedings of the Aristotelian Society*, Supplementary Volume XXXV, 1961.

Halle, M. and K. Stevens, 'Speech Recognition: A Model and a Program for Research', *IRE Transactions on Information Theory*, 1962.

Hirst, R. J., *The Problems of Perception*, London, 1959.

Hubel, D. H. and T. N. Wiesel, 'Receptive Fields of Single Neurones in the Cat's Striate Cortex', *Journal of Physiology* **148**, (1962).

Kohler, I., 'The Formation and Transformation of the Visual World', *Psychological Issues* **3**, (1964).

Lettvin, J. Y., H. R. Maturana, W. S. McCulloch, and W. H. Pitts, 'What the Frog's Eye Tells the Frog's Brain', *Proceedings of the IRE* **47**, (1959).

Lewis, D. K., 'Causation', *The Journal of Philosophy* **70**, (1973).

Lindsay, P. H. and D. A. Norman, *Human Information Processing*, New York, 1972.

Neisser, U., *Cognitive Psychology*, New York, 1967.

Norman, D. A., *Memory and Attention*, New York, 1969.

Paivio, A., *Imagery and Verbal Processes*, New York, 1971.

Pitcher, G., *A Theory of Perception*, Princeton, 1971.

Pylyshyn, Z. W., 'What the Mind's Eye Tells the Mind's Brain: A Critique of Mental Imagery', *Psychological Bulletin* **80** (1973).

Shepard, R. N. and J. Metzler, 'Mental Rotation of Three-Dimensional Objects', *Science* **171**, (1971).

Sperling, G., 'The Information Available in Brief Visual Presentations', *Psychological Monographs* **74** (1960).

Treisman, A. M., 'Verbal Cues, Language and Meaning in Selective Attention', *American Journal of Psychology* **77** (1964).