An Invitation to Cognitive Science
Vol-2
Osherson, D. (ed)
MIT Press, 1990

Chapter 5

Seeing, Believing, and Knowing

Fred Dretske

Epistemology is a branch of philosophy devoted to the study of knowledge and topics—such as truth, memory, perception—relating to knowledge. Epistemology is a philosopher's version of cognitive studies.

Truth is an important part of this study because a central conception of knowledge is knowledge *of the truth*. Though you can know that something isn't so—that, say, the cat *isn't* under the sofa—you can't know something —that the cat *is* under the sofa—that isn't so. To know the whereabouts of the cat requires one to be in possession of the truth about the cat's location. This being so, the idea of truth, as a necessary condition for knowledge, has figured prominently in philosophical discussions of cognition.

Memory and perception also occupy a prominent place in epistemology. Much of our knowledge (some would say *all* of our knowledge) is acquired

This chapter was prepared while I was a Fellow at the Center for Advanced Study in the Behavioral Sciences, Stanford, California. I am grateful for the financial support of the Center, the National Endowment for the Humanities grant FC-20060-85, and the Andrew W. Mellon Foundation.

by perceptual means: we come to know where the cat is by seeing it on the sofa. We might also hear, smell, and feel the cat. These are some of the ways we have of finding out, ways of coming to know, the content and character of our world. The general term for such ways of finding out, ways of coming to know, is *perception*. *Memory* is the name we give to the ways we have for retaining (through time) the acquired knowledge. Powerful mechanisms for acquiring knowledge (keen eyesight, for example) are of little value to animals that cannot remember, if even for a few seconds, anything they learn. A large storage capacity, on the other hand, is wasted on systems with no way of getting information to be stored.

As earlier chapters reveal, there has been a dramatic increase in our scientific understanding of *how* we know some of the things we know. Nevertheless, despite this progress, certain classical philosophical problems, problems concerning the nature, scope, and limits of visual cognition, remain unanswered—or, better, remain without answers that command widespread assent. As we learn more about the way things actually work, these problems tend to be expressed in somewhat different ways. In the past forty years, for instance, computer terminology, a terminology that is embodied in information-processing models of perception and cognition, has become popular. Nomenclature aside, though, the problems are still the old problems, the ones philosophers have pondered and debated for centuries. John Locke, the famous seventeenth-century philosopher, would have little trouble understanding the issues discussed here. Indeed, he had well-developed views on most of these topics.

Although some find it frustrating, this continuing lack of agreement about the right answers to certain puzzling questions—the so-called philosophical questions—is not unexpected. Problems tend to be classified *as* philosophical when they elude established methods, including scientific methods, of solution. But this is no reason to belittle the problems or to despair of their eventual solution. Solutions may lie in finding better methods. This chapter is an attempt to survey some of the more intractable of these problems, to indicate options for dealing with them, and to introduce, when it seems useful, appropriate distinctions and clarifications.

5.1 Seeing Objects and Seeing Facts

When cognitive scientists speak of visual perception, it seems reasonable to suppose that they are referring to something that we normally describe using the verb *to see*. Seeing the cat on the sofa is to visually perceive the cat on the sofa.

To avoid misunderstanding, though, one should ask, at the very beginning, whether visual perception (or seeing) is to be reserved for objects, facts, or something else. After all, we normally speak of seeing objects (like cats and sofas), the properties of objects (the color of a cat, the size of the sofa), events (the cat's jumping onto the sofa), states of affairs (the cat's being on the sofa), and facts (*that* the cat is on the sofa). If these are all to be counted as instances of visual perception, as they appear to be in ordinary language, then care must be taken in a scientific study of visual perception to specify *what* is being perceived: an object, a property, an event, a state of affairs, or a fact. For it is not at all clear that the *same* processes and mechanisms are, or need be, involved in the perception of these different things. Quite the contrary.

Consider, for example, a small child glancing at the sofa and mistaking a sleeping cat for an old sweater. Does the child see an object? Yes, of course. Besides the sofa there is an object, the black cat on the sofa, that the child mistakenly believes to be a black sweater. Though the child does not recognize the cat (*as* a cat), she must, in some sense, see the cat in order to mistake it for a sweater. Nevertheless, though the child sees a black cat on the sofa, sees an object fitting this description, she does not realize that this is a correct description of what she sees. She thereby fails to see the corresponding fact: *that* there is a black cat on the sofa. She sees an object (the black cat on the sofa) but not the fact (that there is a black cat on the sofa) corresponding to it. Shall we say, then, that the child *perceives* the black cat on the sofa? The answer to this question will obviously depend on whether one is thinking of objects (black cats) or facts (that they are black cats).

We can, of course, merely stipulate that visual perception is a way of seeing objects that involves, in some essential way, a knowledge of the object. So when a child—or, indeed, any other kind of animal (an unsuspecting mouse, for instance)—sees a cat on the sofa without realizing what it is, without learning or coming to know that it is a cat, then this way of seeing the cat will not count as *perceiving* the cat. To perceive a cat is, according to this way of using words, to come to know, by visual means, by the use of one's eyes, that it is a cat. Perception is restricted to seeing facts—to seeing that a cat is a cat.

We are free to use words as we please. There is nothing to prevent our restricting visual perception to visual cognition, to a coming-to-know-by-visual-means. It would seem that this particular restriction is, in fact, rather widespread in cognitive psychology. Interested, as they are, in what subjects learn in their perceptual encounters with objects, cognitive psychologists tend to focus on a subject's recognition or identification of objects, ways of seeing (hearing, smelling) things that require some knowledge of what is seen (heard, smelled). So, for instance, recognizing a geometric figure *as* a triangle requires the subject to realize, to come to know, upon seeing it, that it *is* a triangle. If he, upon seeing it, doesn't know what kind

of figure it is, doesn't at least distinguish it from other sorts of figures, then he doesn't *recognize* it—not, at least, *as* a triangle. Recognizing triangles is a way of seeing a fact—the fact, namely, that they are triangles.

We are indeed free to use words as we please. But this proposed restriction of visual perception to the perception of facts, to recognition, to a way of seeing things that requires a knowledge of the thing seen, has unfortunate consequences. For we now have no natural way of describing the child who mistakes the cat for a sweater. Since the child does not know it is a cat, she does not, on this way of using words, *perceive* the cat. What, then, is the relation that exists between the child and the cat? The child is not blind. Light rays, reflected from the cat, are entering the child's eyes and, in some perfectly normal way, causing within her a visual experience that would be quite different if the cat were not there. This being so, it seems most natural to say, from a commonsense standpoint, that the child sees the cat but does not realize that this is what she sees. If, because of the way we have decided to use the word *perception*, this does not count as perceiving the cat, it must surely count as *seeing* the cat. Using the word *perception* in this restricted way, then, would not let us count, as visual perception, a person's seeing a cat in perfectly normal circumstances.

It seems preferable, therefore, to distinguish between seeing objects and seeing facts, not (as above) by artificially reserving the word *perception* for one way of seeing, the way of seeing that requires knowledge of the thing seen (that is, seeing facts), but rather by distinguishing two forms of perception, two ways of seeing. We are then free to speak of seeing a black cat without necessarily realizing (knowing or believing) that it is a black cat (or, indeed, an animal at all) as, say, *sense perception* (of a black cat), and another, recognitional, way of seeing the cat as, say, *cognitive perception* (that it is a black cat). This brings our use of the term *visual perception* (including as it now does both cognitive and sense perception) into closer harmony with the ordinary verb *to see* and at the same time allows us to preserve the important distinction between seeing a cat on the sofa and seeing what it is that is on the sofa.

Given this way of using words, we are then free to describe the efforts of cognitive scientists as investigating the processes underlying these forms of perception, examining their differences and commonalities. Perhaps it will turn out, for instance, that processes described as early vision are merely the processes involved in sense perception, the seeing of objects, and later vision comprises whatever additional conceptual or cognitive processes are essential to the perception of facts (cognitive perception) relating to these objects. Perhaps, also, debates about whether perceptual processes are top-down or bottom-up, about the inferential or constructive character of perceptual processes, about whether these processes are massively parallel or sequential, and about their modularity are all debates that

can be given sharper focus by distinguishing between the *kind* of perception the debate is a debate about. Discussions or perceptual learning and development will also benefit by a close observance of the difference between cognitive and sensory forms of perception.

For these reasons we will adopt in this chapter the device of speaking of sensory and cognitive perception. The first is a way of seeing (or perceiving) cats (or triangles) that does not require (though it may in fact be accompanied by) knowledge that it is a cat (or a triangle) that is seen. This is what we have been calling *object perception*. *Cognitive perception* of a cat (or triangle) will be reserved for that way of seeing the cat (triangle) that necessarily involves a coming to know, a cognition (in fact, a recognition), that it is a cat (a triangle). If one, as we ordinarily describe things, sees a cat (a triangle) and recognizes it only as an animal (a figure) of some sort, fails (for whatever reason) to know or realize that it is a cat (a triangle), then one has sensory, but not cognitive, perception of a cat (triangle). Cognitively one perceives only an animal (figure) of some (unspecified) sort. I leave open the question (but see question 5.1) of whether it is possible to have sensory perception of an object without *any* cognitive perception of it—whether, for instance, one might see a cat without recognizing it as anything whatsoever (not even as an animal of some sort).[1]

5.2 Perceptual Objects

Many, perhaps most, of our cognitive perceptions, the facts we come to know by visual means, are mediated in some way. Our visual knowledge of A depends on, and derives from, our visual knowledge of B. We see that we need gas (come to know, by visual means, that we need gas) by seeing that our fuel gauge registers "empty." We see one fact (that our gas tank is nearly empty) by seeing another fact (that our gauge registers "empty"). We see *by the newspapers* that there has been a tragic plane crash, *by the*

1. The topic of *seeing as*—at one time a fashionable topic in the philosophy of perception—is a hybrid form of perception, a way of seeing that goes beyond sensory perception (requiring a fairly specific cognitive or judgmental attitude or tendency on the part of the perceiver) but falling short of full cognitive perception (knowledge not being required). One sees a stick as a snake. The stick obviously does not have to *be* a snake for one to see it as a snake. Hence, this cannot be cognitive perception, at least not cognitive perception of a snake (for this would require one to recognize it as a snake, something one cannot do of something, like a stick, that is not a snake). Nonetheless, one sees (sensory perception) the stick and takes or judges it to be a snake. The knowledge required of cognitive perception (knowing that the X is an X) is replaced by some variant of belief: one believes, or is inclined to believe, or would believe if one did not know better, of the object (it may or may not be an X) that it is an X.

tracks that the animal went this way, *by her frown* that she is displeased, and *by the thermometer* that the patient has a fever.

Given this dependence of some visually known facts on other visually known facts, the question naturally arises whether some facts are basic in the sense of being known directly and without this kind of dependence on other visually known facts. If my knowledge of the plane crash derives, or is somehow inferred, from my knowledge of what is printed in the newspapers, if my knowledge of what other people are thinking and feeling is somehow inferred from what I can see of their observable behavior and expression, are the latter pieces of knowledge themselves derived from some more fundamental, even more basic, kind of knowledge—possibly a knowledge of how the light (reflected from a newspaper page or a person's face) is structured, how this light is affecting my eyes, or how my brain is reacting to all these external events? Might it turn out, as some philosophers have argued, that all our knowledge of external, objective, facts—that there was a plane crash, that the newspaper reports a plane crash, that Susan is displeased, that she is frowning—derive, ultimately, from our knowledge of subjective facts, facts about the current state of our own mind (how things look)?

This is a question about cognitive perception, about the structure of our knowledge. Are there some facts we know that are fundamental—*foundational*, as philosophers like to put it—in the sense that all other things we know are derived from them? Is our knowledge of the way the world *is* derived from, and ultimately dependent upon, our knowledge of the way the world *appears*?

The answer to this question depends on the answer to a somewhat different question, a question about sense perception. What objects do we see? Do we see cats, sofas, newspapers, and people? If not, then it would seem that our knowledge of these things (the fact, for instance, that the newspapers say there was a plane crash and the fact that Susan is frowning) must derive from our factual knowledge about other things (whatever objects we do see). My knowledge of the plane crash derives from my knowledge of the newspapers because I did not *see* the plane crash. I see only the newspaper. Hence, whatever facts I learn about the plane crash, including the fact that there was a plane crash, must derive from facts I learn about the newspaper.

What facts we see, and which of these facts are fundamental, therefore depends on what objects we see. If you don't see the gas tank, then your visual knowledge of the gas tank, that it still contains gas, must derive from your visual knowledge about whatever objects you do see—in this case, typically, facts about your fuel gauge. You see that you have some gas left *by* seeing what your gauge registers, and this dependence among cognitive perceptions (your knowledge of the gauge being primary) derives from a fact about sensory perception, from the fact, namely, that you see the gauge but not the tank.

Even when we speak of perceiving one object *by*, *through*, or *in* perceiving another—in the way we speak, for instance, of seeing the game *on* TV or seeing someone *in* a movie (or photograph)—our knowledge of the game or person will be secondary relative to our knowledge of the electronic or photographic image. Insofar as we regard the image appearing on our television or movie screen as the primary, or real, *object* of perception, we regard facts about these images as cognitively primary. Facts about the people and events being represented are secondary. For instance, we learn (see) that a player kicked a field goal by observing the behavior of the electronically produced images of the player, the ball, and the goalposts appearing on our television screen.

Hence, a question about the structure of cognitive perception—whether in fact there is a fundamental level of visual knowledge, and if so, whether this is knowledge of objective or subjective facts—awaits the answer to a prior question: What is the structure of sense perception? What objects do we see? The answer to this question will constrain, if not determine, the answer to the questions about cognitive perception. If we do not see physical objects, if we are (in sense perception) always aware of mental images (representations) of external objects (as some philosophers and psychologists seem to believe), then our knowledge of objective reality (if, indeed, we have such knowledge) will necessarily be derivative from and secondary to our knowledge of our own mental states.

Discussions of these issues are often clouded by failure to appreciate the difference between cognitive perception and sense perception. It is sometimes argued, for instance, that we do not perceive ordinary physical objects because, for whatever reason (the reason is usually skeptical in character), we do not know, or cannot be absolutely certain, that there *are* physical objects. For all we know, all experience, even the experience we take to be *of* a real external world, may be illusory. It could all be a dream. This argument, though it has a distinguished history, is a fairly obvious conflation of cognitive perception and sense perception. One does not have to know, let alone know for certain (whatever that might mean), that there are physical objects in order to see (sense perception) physical objects. Such knowledge is only required for cognitive perception. Just as the child described above saw a cat on the sofa without knowing what it was, it may turn out that we see ordinary physical objects (including cats and sofas) every moment of our waking life without ever being able to know (if the philosophical skeptic is right) that this is what we are seeing. Questions about what objects we see are quite different from questions about what facts we know.

Failure to keep the distinction between sense and cognitive perception clearly in mind also tempts students into mistakenly supposing that if our knowledge of physical objects is somehow derivative from the way they *appear* to us, from the way they *look*, then what we really perceive when (as we ordinarily say) we see a cat is an internal mental image of the cat. We see (as it were) the *look* or *appearance* of the cat. Such an inference would be fallacious because even if our cognitive perceptions rest on subjective foundations (on the way things look to us), our sense perceptions need not rest on similar foundations. We may know that there are physical objects by the way they appear to us (so that cognitive perception has, in this sense, a subjective basis), but our sense perception of objects is itself direct and unmediated. In other words, we may come to know (see) it is a cat (a fact) by the way it appears, but what we see (the object) is the cat itself, not its appearance.

Aside from these possible confusions, though, there are a variety of positions that have been, and continue to be, taken on the nature of both cognitive and sense perception—on what facts and objects are most immediately and directly seen. Though these theories, in both their classical and their modern form, are often hard to classify because of their failure to be clear about whether it is cognitive perception or sense perception they are talking about, they can be roughly characterized as follows.

Direct (Naive) Realism (sometimes said to be the view of the person-on-the-street) holds (1) that there is a real physical world, objects and facts whose existence is independent of our perception of them (this makes the view a form of physical *realism*) and (2) that under normal conditions we are, in a direct and unmediated way, perceptually aware of these objects and facts (hence, *direct* and therefore, according to its detractors, *naive* realism). In other words, what we are directly aware of in sense perception is, unlike a headache or an afterimage, something physical that continues to exist when we are no longer aware of it.

Representative Realism (also called the *Causal Theory of Perception*) shares with Direct Realism (and common sense) the first of these two doctrines. It disagrees, though, about the second. According to Representative Realism, our perception of physical objects is indirect, mediated by a more direct apprehension of something mental, some internal representation (hence the name *representative* realism) of external physical reality. These mental representations have been given various names: sensations, ideas, impressions, percepts, sense-data, experiences, and so on. But the idea is almost always the same. Just as we see what is happening on the playing field *by* seeing what is happening on our television screen (so that our knowledge of the game, when viewed on television, is indirect), so knowledge of even the most obvious physical fact—the fact, say, that there is a table (or, indeed, a television set) in front of us—is itself indirect. We see that there is a table

in front of us *by* seeing, or somehow being aware of, its internal, mental representation. When we are watching a game on television, then, our knowledge of the actual game is *doubly* indirect: we know about a game occurring 1,000 miles away by knowing what is happening on a television screen a few feet away, and we come to know what is happening a few feet away by becoming aware of what is happening (presumably *no* distance away) in our own minds. In the last analysis, then, all our knowledge of objective (physical) fact rests on a knowledge of subjective (mental) fact because the only objects perceived (directly) are mental objects—the way things appear.

Going beyond these forms of realism are various forms of *idealism* (sometimes called *phenomenalism*), theories that deny an objective physical reality altogether. Everything that exists depends for its existence (like a headache or an afterimage) on someone's awareness of it; hence, everything is in the nature of a mental entity like an idea (hence, *ideal*ism). Since these extreme views have few, if any, serious advocates within the philosophical (not to mention cognitive science) community today, we will leave them without further comment.

As indicated earlier, one might be a Direct Realist on sense perception but an Indirect Realist on cognitive perception. The objects we see are physical objects, but we know about them via their effect on us (the way they appear to us) in sense perception. The problem with this mixed position is the problem of saying just how one might come to know how objects look—which, according to some theorists, is a knowledge of how, in sense perception, we internally represent them—without thereby becoming aware of, and hence perceiving, the internal representations themselves (thereby becoming an Indirect Realist on sense perception also). To put it crudely, how can one *know* how things look without perceiving, or somehow being aware of, their look?

The debate between Direct and Indirect Realists becomes very technical at this point. Indirect Realists maintain that we are directly aware of mental objects—images—in hallucinations and dreams. Aside from the *cause* of the experience, though, there is no reason to distinguish between these illusory experiences and our ordinary veridical perception of (physical) objects. In both cases we are directly aware of the internal mental representation. When we speak, as we commonly do, of seeing an ordinary object (like a cat), we are, if we speak truly, being caused to experience some catlike image by a real cat (a real cat that we do not directly perceive). When we hallucinate or dream of a cat, there is no such external cause—hence, we speak of these experiences as illusory. In all cases, though, it is the image that we directly experience. Only the cause of the experience is different. Direct Realists try to counter this, and related, arguments by insisting that although sensory perception of real objects requires the

having (and thereby the existence) of internal representations, and though such representations in fact determine the way these objects look or appear to us, there is no reason to suppose we perceive these representations themselves. We perceive a cat by (internally) representing a cat, not by perceiving an internal representation of a cat.

5.3 Perceptual Processes

The debate about the objects of perception is related to a debate (not always clearly distinguished from it) about the kind of processes underlying perception. Do perceptual processes, those culminating in our seeing something, exhibit the qualities of reason and intelligence? Do they, despite being unconscious, have an inferential or computational character, moving from premise to conclusion (deductive reasoning), or from data to explanatory hypothesis (inductive reasoning), in something like the way human agents consciously solve problems? When I see a cat on the sofa, or that there is a cat on the sofa, does my visual system do something similar to what clever detectives do when they infer, on the basis of certain clues and signs, that a certain state of affairs not directly apprehended *must be* the case?

We can, of course, metaphorically describe the operations of anything, even the simplest machine, in thoughtlike, semicognitive terms. We are especially fond of doing this with computers. We say they know, that they remember, recognize, infer, and conclude. If one counts arithmetical operations as forms of computation, even dime store calculators perform (or are *described* as performing) impressive feats of reasoning—multiplying, taking square roots, and calculating percentages in fractions of a second. We even speak of such comparatively humble devices as thermostats and electric eyes in quasi-perceptual terms—as, for example, "sensing" a drop in room temperature or the approach of a person and responding by turning the furnace on or opening a door. The question, then, is not whether we *can* speak this way, not even whether it is sometimes useful to talk this way (to adopt what Dennett (1987) calls the *intentional stance*), but whether this is anything more than a metaphorical crutch—a figure of speech that conceals or masks our ignorance about underlying causal processes and mechanisms. Do visual systems ever literally solve problems, infer that something is so, formulate (on the basis of sensory input) hypotheses about the distant source of stimulation in the way that rational agents do this at the conscious level?

Hermann von Helmholtz, the great nineteenth-century physiologist, thought so, and many investigators today (see, for example, Gregory 1974a, 1978; Rock 1977, 1983; Ullman 1980) are inclined to agree. At least

they view the processing of visual information as a form of problem solving and hence as a form of reasoning that, though unconscious, exhibits enough of the essential properties of fully rational thought and judgment to make it, in a fairly literal sense, an instance of problem solving itself. The light reaching the receptors (sometimes called the *proximal stimulus*) carries information—fragmentary and impoverished (and thereby ambiguous) information to be sure, but information nonetheless—about distant situations (the *distal stimuli*). The visual system's function is to take these data and to construct, as best it can, a reasonable conjecture (hypothesis, judgment) about the distal source of this stimulation. It begins with premises describing receptor activity, data concerning the distribution and intensity of energy reaching the receptor surface, and is charged with the task of arriving at useful conclusions about the distal source of this stimulation. The conclusion it reaches (for instance, it must be a cat out there causing this pattern of retinal activity) constitutes the subject's perception of a cat. If the visual system reaches a different conclusion—that, for instance, it is probably an old black sweater—the subject sees an old black sweater instead of a fluffy black cat. If the perceptual system can't make up its mind, or keeps changing its mind (it's a cat; no, on second thought, it's probably a sweater; no, that can't be right, it's probably a cat), the subject sees first a cat, then a sweater, then a cat again. Though such flip-flopping seldom occurs when we are looking at real cats (because, in normal circumstances, light from real cats is generally richer in information—hence, less ambiguous—about the kind of object that has structured the light), it sometimes happens with specially constructed figures viewed under restricted (say, monocular) conditions—Necker cubes, for instance. Since so much emphasis is placed on the visual system's efforts at constructing a reasonable interpretation or hypothesis (about the distal stimulus) from information reaching the receptor surfaces, this approach to perceptual processing is often described as a *Constructivist* or *Computational* approach to visual perception.

Since Constructivists regard sensory stimulation, even in the best of viewing conditions, as inherently ambiguous (there are always a variety of distal arrangements that could have produced that pattern of proximal stimulation), they view perceptual processing as primarily a matter of *adding* information to the stimulus (or *supplementing* the information available in the stimulus) to reach a perceptual outcome: seeing a cat. Since the proximal stimulation does not unequivocally specify the distant object as a cat, and since we nonetheless (under optimal viewing conditions) see a cat (the visual system reaches this conclusion), the perceptual system must exploit some other source of information to reach this judgment—adding or supplementing (via some inductive inference) the information contained in the stimulus.

There has been a vigorous challenge to this (more or less) orthodox position in the last forty years. Gibson, in a series of influential books (1950, 1966, 1979) and articles (1960, 1972) has argued that the stimulus, *properly understood*, contains all the information needed to specify the distal state of affairs. If the proximal stimulus is understood, not as a static distribution of energy occurring on the receptor surfaces *at a time*, but as the total dynamic pattern of stimulation reaching a mobile observer *over time*, there is no need for inference, reasoning, and problem solving. There is sufficient information in the stimulus (thus broadly conceived) to specify (unambiguously determine) the character of the distal object. Why reason about what is out there when the stimulus tells you what is out there? Why suppose, as Constructivists do, that perceptual systems are smart detectives when all they really have to be (given reliable informants—that is, information-rich stimuli) is good listeners, good extractors of the information in the signals reaching the receptors? Since this approach tends to eliminate all intervening cognitive (indeed, all intervening psychological) mechanisms from the processes resulting in our perception of objects, it is often referred to as a *Direct Theory* of perceptual processing.

Relevant to the question of whether perceptual systems are more like good detectives doing their best with ambiguous data (Constructivism) or more like good listeners faithfully registering stimulus information (Direct Theory) is what Fodor (1983) describes as the *modularity* of information-processing systems. A system is (comparatively) modular when it is (comparatively) insulated from information available to other parts of the total system. If I am told (and thereby know) that it is a cat on the sofa, for instance, does this, *can* this, affect my visual perception of the cat? If not, my visual system exhibits modularity with respect to this kind of information (information available to the central processor from auditory sources). If this collateral information is capable of affecting what I see, then the visual system (understood as that subsystem responsible for my seeing what I see) is not modular in relation to this kind of information.

If the visual system is modular, its operation (and therefore presumably what the subject perceives) is unaffected by what other information may be available to other parts of the system (or what the subject may know as a result of information received from these other parts). Modular systems are therefore described as *stimulus-driven* (the processing is bottom-up rather than top-down): it is the stimulus itself (information at the bottom, as it were), not the system's (possibly variable) hypotheses about that stimulus (information available at the top) that guides the processing of incoming signals and thereby determines what the subject perceives. Modular systems are therefore most naturally thought of in the second of the two ways described above—as good extractors of preexistent information, information that is already in the stimuli, not as good detectives or problem

solvers about the best interpretation of informationally ambiguous stimuli. There is no point in supposing that a process of reasoning is occurring in modular systems when the process, being modular, is not allowed to use information (other than what is in the stimulus itself) to generate perceptual conclusions. Modular systems are not intelligent. They don't have to be. They have no problems to solve. They just do what the stimulus tells them to do.

It is by no means obvious that these two approaches to the analysis of perceptual processes are incompatible. It may turn out, for example, that although the stimulus, properly understood, *is* rich in information about distal objects, rich enough (let us suppose) to unambiguously determine what distal objects produced it, it nonetheless requires inferential (reason-like) processes to decode the signal, to extract this information from the stimulus. Fingerprints, being unique to their bearers, may unambiguously determine or specify (in an information-theoretic sense) who held the gun. It nevertheless takes a good deal of problem solving, after one has discovered the incriminating prints, to figure out who held the gun. One has to know which people go with which prints, and this may take memory, inference, and prior learning (the sort of cognitive work that organizations like the FBI invest into the creation of a fingerprint file). As Ullman (1980, 380–381) puts it, the role of processing may not be to create information, but to extract it, integrate it, make it explicit and usable.

There are, then, a variety of ways of expressing questions about the nature of those processes underlying our perception of the world. But these questions should not be confused, as they often are, with questions about the objects of perception, the questions discussed in section 5.2. Gibson's views have been described (by both Gibson and others) as a theory of *direct* perception. This can be misleading. It certainly is confusing. The sense (if any) in which this theory is direct is much different from the sense in which Direct Realism is direct. Direct Realism is a theory about the *objects* of perception, about *what* we see. The kind of direct realism we are now talking about, the kind associated with Gibson's work, is a theory about the processes underlying perception, about *how* we see what we see. One can be a Direct Realist about the objects of perception, holding that we directly apprehend physical objects (not sensations or mental intermediaries), and be a Constructivist about the processes underlying our (direct) perception of these objects. One can suppose that intelligence, some kind of thought-like process, is involved in the construction of internal representations without supposing that one thereby sees (or in any way perceives) the representations so constructed. One can, in other words, be a Direct Realist about the objects of perception and an Indirect Realist, a Constructivist, about the processes underlying this direct relationship.

Once again, though, controversy about the intelligence, or lack of it, of perceptual processes is often muddled by failure to be clear about exactly which processes are in question. It should be obvious that cognitive perception—our perception of facts, our seeing *that* (and hence coming to know that) there is a cat on the sofa—is the result of a process that is strongly influenced by higher-level cognitive factors. Cognitive perception is clearly not modular. A subject who does not already know what a cat is, or does not already know what they look like—a small child or an inexperienced animal, say—will be unable to *see* (recognize) *what* is on the sofa, unable to *see that* there is a cat there (to be carefully distinguished from an ability to see the cat there). For cognitive perception of the cat on the sofa, in contrast to sense perception of the cat, requires not only the appropriate concepts (for *cat* and *sofa*) but some intelligence in the application of these concepts to the objects being perceived (the cat and the sofa). The upshot of cognitive perception is some *known fact* (say, that there is a cat on the sofa) and such facts are not learned without the cooperation of the entire cognitive system. By changing a subject's cognitive set—changing what the subject knows or believes about the way things look, for instance—one easily changes what the subject learns, comes to know, hence perceives in a cognitive way, about the objects it sees (in a sensory way). Some form of Constructivism or Computationalism is therefore inevitable for seeing *facts*.

The real question is, or should be, whether that part of the visual system given over to *sense perception*, to seeing objects (like cats and sofas), is also intelligent. Does *it* exhibit some (any? all?) of the marks of reasoned judgment? Is *it* modular?

The answer to this question will depend on just what one takes to be involved in the perception of objects, in seeing, say, a cat on the sofa or a person in the room. If the upshot or outcome of cognitive perception is some known fact—that there is a cat on the sofa or a person in the room—what is the upshot or culmination of sense perception? When, at exactly what stage in the processing of incoming information, do we see the cat on the sofa and the person in the room? If recognizing the object *as* a cat or *as* a person is not necessary to the sensory perception of these objects (as it is to their cognitive perception), what is necessary? Since we can see a cat at a distance, in bad lighting, or in unusual conditions (circumstances in which it does not even look like a cat), we cannot suppose, following Gibson, that to see a cat is to have information in the stimulus that specifies the cat *as* a cat. For in such cases there may be no information in the stimulus about what it is we see. That does not prevent our seeing it.

It is true, but unilluminating, to be told that the sensory perception of an object occurs when the visual system constructs a sensory representation of the object. What we want to know is what kind of representation a sensory representation is. If cognitive perception of a cat occurs when the system constructs a cognitive representation of the cat, an internal judgment or belief that it is a cat (some kind of internal description of the cat *as a cat*), what is a sensory representation of the cat, the kind of internal representation whose occurrence constitutes a sensory perception of the cat? Is it something like what philosophers and psychologists used to call a *sensation*? Or is it more like what they (or some of them) now call a *percept*? Or, to use even more fashionable jargon, is it more like what Marr (1982) and his associates call a *2 ½-D sketch*?

Until these questions are answered, we can expect little progress on questions about the nature of perception itself. How can we tell whether sensory perception is best thought of in terms of a clever detective or a good listener if we cannot say, in any clear way, what final product, what kind of internal representation, this kind of perception is supposed to produce?

5.4 Perceptual Change

Do we learn to see things? Does prolonged experience of the world change what we perceive or the way we perceive it? Do people with radically different languages, radically different ways of describing their surroundings, see their surroundings differently? Do completely different world views—what Kuhn (1962), for instance, calls *incommensurable scientific theories*—generate differences in what people can observe and, hence, in the data on which their theoretical differences rest?

Such questions have fascinated philosophers and psychologists, linguists and anthropologists, for centuries. The answers to these questions are not easy. Nevertheless, some things seem reasonably clear—if not the final answers themselves, then at least the sorts of considerations that must inform the search for final answers.

The first point, a point that has been made repeatedly in this chapter, is that before rushing in with answers to any of these questions, one should first be very clear about the question. What kind of perception is the question a question about?

As a case in point, the question about whether we learn to see things has a reasonably straightforward answer if it is a question about cognitive perception, about the facts we come to know by visual means. The first time (as a very small child presumably) I saw a maple tree I probably didn't know what kind of tree it was. Having no experience or knowledge of maple trees, being ignorant of what maple trees looked like (or, indeed, of what maple trees were), I didn't recognize what I saw *as* a maple tree. I

didn't see what kind of tree it was. Now, however, I am quite expert in this kind of identification. I can look at maples, at least the more common varieties, and quickly recognize them as maples. I can see, by their general shape, their bark and leaf structure, *what kind of tree* they are, that they are maple trees. There has been a change, therefore, in my ability to cognitively perceive objects around me, a change that came about by experience, learning, and (in this case) diligent study and practice. Learning of this kind is a pervasive and familiar phenomenon.

But if the question about perceptual learning is a question about sensory perception, about the objects we see, about whether we learn to see maple trees themselves (and not just the fact that they are maple trees), the answer appears to be quite different. I did not learn to see maple trees. I could do that when I was a very young child—before I learned to recognize them. What I learned is how to identify the things I saw, things I therefore saw before I learned to identify them. Sensory perception of objects normally comes before the cognitive perception of these same objects. If it did not, there would be no way to learn what objects look like. How can you learn what objects look like if you cannot see them? Humans do not see things at the moment of birth, of course. Certain physiological changes must first occur before we can, for instance, focus on objects and coherently process information contained in light. But these maturational processes are not to be classified as *learning* in any ordinary sense. We no more learn to see solid objects than we learn to digest solid food.

This is not to say that some changes in our perception, our sensory perception, of objects may not occur after prolonged experience. Perhaps objects start looking different after they become familiar or after we know certain things about them. Does a familiar face—the face of a loved one, say—look different after it has become familiar from the way it looked the first time you saw it? Do coins look larger to poor children than they do to rich children? Do lines in an optical illusion that look to be of different lengths start looking the same after you learn (by measuring them) that they are the same length? These questions are questions about the way things look, about the character of our visual experience, about something we earlier dubbed (without really knowing or explaining what it was) the *sensory representation* of objects. They are not questions about the way we cognitively represent objects, about our perceptual beliefs or judgments. Changes and differences in cognitive representations are an obvious and familiar fact of life. That such changes exist is not worth arguing about (though the changes themselves are certainly worth studying). Changes, if any, in our sensory representations are not so obvious. Quite the contrary. To document such changes one has to be very clear about what sensory representations are and what constitutes a change in them. To answer questions about whether we learn to see in this sense, then, requires a clear,

at least a much better, understanding of the nature of sensory representation, of what kind of internal response to an external object constitutes our seeing the object.

Similar remarks can be made about various forms of perceptual *relativity*. Is perception relative? Well, cognitive perception is certainly relative to many things—everything, in fact, capable of influencing what one comes to believe. If not having a word for X or a theory about X means I cannot come to have certain beliefs about X, then not having a word (or a theory) for X will prevent me from cognitively perceiving X. Without an appropriate language for talking about oxygen, without some knowledge (however crude) of chemical theory, I can hardly be expected to see *when* oxidation is occurring (see *that* it is occurring) even when it happens under my nose. I just will not recognize it—certainly not *as* oxidation. So the cognitive perception of oxidation is relative to those factors—factors like possession of the right concepts and knowledge of the appropriate scientific theories—that are essential to a knowledge that oxidation is occurring. For the same reason, people who have badly mistaken astronomical views will not be able to see what others see when a lunar eclipse occurs—that the moon is moving into the earth's shadow. They will not see that a lunar eclipse is occurring because, with mistaken views about what is happening (they think the gods are showing displeasure by extinguishing the moon), they will not learn what everyone else learns when they see the same thing: that the earth is casting a shadow on the moon.

But though cognitive perception is obviously relative in this way, there is no reason to think—in fact, there is a lot of reason *not* to think—that sensory perception is similarly relative. Though the astronomically ignorant may not see that an eclipse is occurring, they certainly see *the eclipse* (= the earth's shadow moving across the face of the moon). That, in fact, is what frightens them. And though the chemically ignorant can hardly be expected to see that oxidation is occurring, they can, given normal eyesight, witness the oxidation, the blazing fire, as well as everyone else.

To say that perception is relative to a certain factor is to say that our perception of things depends on that factor. Change that factor (enough) and we change what is perceived or, possibly, whether anything at all is perceived. To suggest, then, that sense perception is *not* relative to a variety of factors affecting our perception of facts is a way of suggesting that sense perception is comparatively modular. It is *not* sensitive to the cognitive influences (a subject's language, conceptual scheme, or scientific world view) that determine one's perception of facts. The issue of perceptual relativity, and more generally of perceptual change and learning, then, is merely another way of approaching questions raised in earlier sections of this chapter about perceptual processes in general. It seems, therefore, that the answers to a variety of questions, both philosophical (raised in this

chapter) and scientific (addressed in earlier chapters), depends on a deeper understanding of perceptual processes and the different outcomes, sensory versus cognitive, that they support. Achieving deeper understanding of this sort will require the combined efforts of investigators from many fields.

Suggestions for Further Reading

For more detailed treatments of the distinction between the perception of objects and the perception of facts, and for a defense of the idea that sense perception does not require cognitive perception, that seeing is not (or at least not necessarily) believing, see Dretske 1969, 1978, 1981, Sibley 1971, and Warnock 1955. For opposing viewpoints (defending the idea that all perception involves, if not knowledge, then a kind of judgment or belief), read Armstrong 1961, Hamlyn 1957, Heil 1983, and Pitcher 1971.

For discussion of the issues surrounding the controversy between Direct (Naive) and Representative Realists about the objects of perception (whether we directly see physical objects or some mental surrogate), see Dretske 1969, 1981, Goldman 1977, Sanford 1976, and Chisholm 1957 for direct theories and Jackson 1977, Ayer 1956, 1962, and Price 1932 for indirect theories. A (by now) classic article on causal theories of perception is Grice 1961.

Concerning the question of whether perceptual processes are constructive or not, an exchange that brings out most of the issues can be found in Ullman 1980, an article in the Constructivist vein, and the comments on it (many of which defend a direct theory). For vigorous exposition and defense of Constructivism, see Gregory 1974a,b, 1978, Rock 1977, 1983, and Fodor and Pylyshyn 1981. Works generally supportive of a direct theory of perceptual processing (and therefore sympathetic to many of Gibson's ideas) include Turvey 1977, Mace 1977, Michaels and Carello 1981, and many of the contributions to Shaw and Bransford 1977 and Macleod and Pick 1974. For further discussion, including evaluations of the empirical status of these two approaches, see Hayes-Roth 1977, Johansson, von Hofsten, and Jansson 1980, and Epstein 1973.

For perceptual learning, change, and development, consult the references in chapter 4. For perceptual relativity, see Churchland 1979, 1988, chapter 5 of Dretske 1969, Hanson 1958, Kuhn 1962, Brown 1987, and Shapere 1982.

Questions

5.1 Is sensory perception possible without cognitive perception? Can one see an object—like a cat—without thereby coming to know *something* (not necessarily that it is a cat) about it? If not, does this mean that some kind of conceptual ability (whatever is needed to know) is necessary for vision—the ability to see things?

Do animals see the same things we do? Do they have beliefs? If so, do they have the same kind of beliefs we have? Does every animal with eyes (and therefore, presumably, vision—the ability to see things) have thoughts?

5.2 Is it possible to see facts while seeing no objects—to have cognitive perception without sense perception? What is the best way to describe what happens when one detects a change in overall illumination (that the lights went out, say) with one's eyes closed? Is this a case of seeing a fact (that the lights went out) without seeing an object?

5.3 Is seeing an event (the cat's jumping on the sofa) and a state of affairs (the cat's being on the sofa) more like seeing objects (the cat on the sofa) or more like seeing facts (that the cat is on the sofa)? What is required to see the properties of objects—say, the size or color of a cat? Does one see the color of a black cat when one sees another (different) object of the same color (say, a black ball)?

5.4 Are objects and facts seen in dreams and hallucinations? Or does one merely dream (or hallucinate) that one is seeing an object or a fact? (Is this difference, the difference between seeing an object in one's dream and dreaming one sees an object, a *real* difference?) Are these colored things (of which one is aware in dreams and hallucinations) in the mind? Are there round red things *in the brain* when one dreams of something red and round?

5.5 If a star explodes and disappears when the light from it is still on its way to earth, does one nonetheless still see the star when the light reaches the earth many years later? If so, does this mean that one can see things that do not exist (any longer)? If not, what (if anything) does one see when the light from the star enters one's eyes and gives rise to an "experience" of a twinkling spot of light?

5.6 Are experts in a given field—auto mechanics (on cars), cooks (on food), and tailors (on fabrics), for instance—able to see things that the nonexpert, the layperson, cannot see? How is one to interpret auto mechanics' claims that they can *hear* things that laypeople can't hear—that, for example, a car's valves need adjusting or that it needs a tune-up? What kind of perception is this?

References

Armstrong, D. M. (1961). *Perception and the physical world*. London: Routledge and Kegan Paul.

Ayer, A. J. (1956). *The problems of knowledge*. London: Penguin Books.

Ayer, A. J. (1962). *The foundations of empirical knowledge*. London: Macmillan.

Brown, H. (1987). *Observation and objectivity*. Oxford: Oxford University Press.

Chisholm, R. (1957). *Perceiving: A philosophical study*. Ithaca, NY: Cornell University Press.

Churchland, P. (1979). *Scientific realism and the plasticity of mind*. Cambridge: Cambridge University Press.

Churchland, P. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science* 55, 167–187.

Dennett, D. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

Dretske, F. (1969). *Seeing and knowing*. Chicago: University of Chicago Press.

Dretske, F. (1978). The role of the percept in visual cognition. In C. Wade Savage, ed., *Perception and cognition*. (Minnesota Studies in the Philosophy of Science 9.) Minneapolis, MN: University of Minnesota Press.

Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

Epstein, W. (1973). The process of "taking-into-account" in visual perception. *Perception* 2, 267–285.

Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Fodor, J., and Z. Pylyshyn (1981). How direct is visual perception?: Some reflections on Gibson's "ecological approach." *Cognition* 9, 139–196.

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

Gibson, J. J. (1960). The concept of the stimulus in psychology. *American Psychologist* 15, 694–703.

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

Gibson, J. J. (1972). A theory of direct visual perception. In J. R. Royce and W. W. Rozeboom, eds., *The psychology of knowing*. New York: Gordon and Breach.

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

Goldman, A. (1977). Perceptual objects. *Synthese* 35, 257–284.

Gregory, R. (1974a). Choosing a paradigm for perception. In E. C. Carterette and M. P. Friedman, eds., *Handbook of perception*. Vol. 1: *Historical and philosophical roots of perception*. New York: Academic Press.

Gregory, R. (1974b). Perceptions as hypotheses. In S. C. Brown, ed., *Philosophy of psychology*. London: Macmillan.
Gregory, R. (1978). *Eye and brain: The psychology of seeing*. 3rd ed. New York: McGraw-Hill.
Grice, P. (1961) The causal theory of perception. *Proceedings of the Aristotelian Society, Supplementary Volume* 35.
Hamlyn, D. W. (1957). *The psychology of perception*. London: Routledge and Kegan Paul.
Hanson, N. R. (1958). *Patterns of discovery*. London: Cambridge University Press.
Hayes-Roth, F. (1977). Critique of Turvey's "Contrasting orientations to the theory of visual information processing." *Psychological Review* 84, 531–535.
Heil, J. (1983). *Perception and cognition*. Berkeley, CA: University of California Press.
Jackson, F. (1977). *Perception*. Cambridge: Cambridge University Press.
Johansson, G., C. von Hofsten, and G. Jansson (1980). Event perception. *Annual Review of Psychology* 31, 27–63.
Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
Mace, W. M. (1977). James Gibson's strategy for perceiving: Ask not what's in your head, but what your head is inside of. In Shaw and Bransford 1977.
Macleod, R. B., and H. L. Pick, Jr. eds. (1974). *Perception: Essays in honor of James J. Gibson*. Ithaca, NY: Cornell University Press.
Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
Michaels, C. F., and C. Carello (1981). *Direct perception*. Englewood Cliffs, NJ: Prentice-Hall.
Pitcher, G. (1971). *A theory of perception*. Princeton, NJ: Princeton University Press.
Price, H. H. (1932). *Perception*. London: Methuen.
Rock, I. (1977). In defense of unconscious inference. In W. Epstein, ed., *Stability and constancy in visual perception*. New York: Wiley.
Rock, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
Sanford, D. (1976). The primary objects of perception. *Mind* 85, 189–208.
Shapere, D. (1982). The concept of observation in science and philosophy. *Philosophy of Science* 49, 485–525.
Shaw, R., and J. Bransford, eds. (1977). *Perceiving, acting and knowing*. Hillsdale, NJ: L. Erlbaum Associates.
Sibley, F. N. (1971). Analyzing seeing. In F. N. Sibley, ed., *Perception*. London: Methuen.
Turvey, M. T. (1977). Contrasting orientations to the theory of visual information processing. *Psychological Review* 84, 67–88.
Ullman, S. (1980). Against direct perception. *Behavioral and Brain Sciences* 3, 373–415.
Warnock, J. (1955). Seeing. *Aristotelian Society Proceedings* 55.

Action