OF SENSORY SYSTEMS AND THE "ABOUTNESS" OF MENTAL STATES*

Our thoughts, we believe, are "about" things. When I look at a picture of the Eiffel Tower, I am thinking about an object, a particular structure that exists in Paris; when I remember the fragrance of gardenias, I recall a property that certain flowers have; and when I wonder how George Smith is getting along, I have in mind a particular person, a philosopher I met at Tufts University. Somehow or other, my thoughts are linked to properties and objects. The question, of course, is 'How'? How are mental events tied to the objects they represent? What is "aboutness" and how does it come into being?

I want to discuss recent naturalistic theories of "aboutness," naturalistic theories of, roughly, how certain psychological or neural states come to represent, stand for, or be about properties and objects. The theories at issue are, among others, the theories of Patricia and Paul Churchland, Daniel C. Dennett, Fred Dretske, Jerry Fodor, David Papineau, Denis Stampe, and Kim Sterelny.[1] What I wish to show is that the naturalists' project, as commonly conceived, rests upon an intuitive and seemingly banal view of what the senses do—

* This paper began as the first chapter of my doctoral dissertation, "On Piranhas, Narcissism, and Mental Representation: An Essay on Intentionality and Naturalism," Ann Arbor, Michigan (1989), and gone through a number of incarnations as both papers and talks. The comments of many people have been very helpful—but I owe special thanks to Joseph Malpeli, Daniel C. Dennett, Kim Sterelny, Jaegwon Kim, Martin Hahn, the members of the Embedded Computation Project at Xerox PARC, and the members of The Spatial Representation Project at King's College, Cambridge. My largest debt, however, is to Brian C. Smith, whose views on intentionality have profoundly influenced my own philosophical views. This paper was written, in part, with the financial support of the Center for the Study of Language and Information (under a grant from the Systems Development Foundation) and the Xerox Palo Alto Research Center.

[1] See, for example, Paul Churchland's *Scientific Realism and the Plasticity of Mind* (New York: Cambridge, 1979); a joint article by both Paul and Patricia Churchland, "Functionalism, Qualia and Intentionality," *Philosophical Topics*, I (Spring 1981): 121-45; Patricia Churchland's *Neurophilosophy: Toward a Unified Science of the Mind-Brain* (Cambridge: MIT, 1986); Dennett's *The Intentional Stance* (Cambridge: MIT, 1987); Dretske's *Knowledge and the Flow of Information* (Cambridge: MIT, 1981) and *Explaining Behavior: Reasons in a World of Causes* (Cambridge: MIT, 1988); Fodor's *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge: MIT, 1987) and *A Theory of Content and Other Essays* (Cambridge: MIT, 1990); Papineau's "Representation and Explanation," *Philosophy of Science*, LI, 4 (December 1984): 550-72; Stampe's "Show and Tell," in B. Freeman et alia, eds., *Forms of Representation* (Amsterdam: North-Holland, 1975): 221–45; and Sterelny's *The Representational Theory of Mind* (Cambridge: MIT, 1990).

0022-362X/96/9307/337–72

that the senses function to inform the brain of what is going on "out there," in the external world and in one's own body—and that this "banal" view about sensory function is false or, more moderately, unlikely to be true in the strong guise that the naturalists' project requires. Hence I suspect that the naturalists' project, at least in its present form, is unlikely to work.

As one would imagine, much of this paper will be concerned with how I think the senses do function, in order to show exactly how the naturalists' project goes astray. The purpose of this paper, however, is not merely to cast empirical aspersions upon a philosophical project, to lob stones into the naturalists' camp, so to speak. Like most naturalists, I, too, believe that our thoughts/neural states are about the world—about the scent of gardenias, about the Eiffel Tower, about the tickles and itches in one's feet—and that this fact of "aboutness" is fundamentally a *biological* fact about persons. Broadly construed, I share the naturalists' goal. It is the trodden path that has me worried. So the purpose of this paper is, first, to dissuade the average naturalist from the current route and, more importantly, to suggest a different path that a naturalistic theory of aboutness might take.

I. THE NATURALIST CAMP AND THE TRADITIONAL VIEW OF THE SENSES

What, then, is the naturalists' project? Part of the problem with characterizing current naturalistic theories of representation is simply that, as a "camp," they comprise quite a diverse lot. That is, there is little agreement among the naturalists about what the project is—what exactly requires explanation—much less how the naturalistic theory should go. Take, for example, three of the more well-known naturalist programs, three among the many. First, there are those theorists, such as Fodor, who are committed realists, philosophers who take themselves to be explaining the content and directedness of our ordinary folk-psychological ascriptions—hence to be solving the traditional problem of the intentionality of mental states. On the opposite side of the clearing, one finds the eliminativists, with the Churchlands as their leading force. They deny that there are folk-psychological states, or more mutedly, they worry that the taxonomy of folk psychology will fail to survive the advances of the neurosciences. Eliminativists, then, are not concerned with the intentionality of folk-psychological states per se (for there likely are none) but with the problem of how *neural states* or *computational states* come to represent. Finally, a third distinct group, following Dennett, adopts "the intentional stance." On this view, the ascription of intentional states is a matter of holistic interpretation. When we ascribe psychological states to an individual, we are not describing her inner representational events. We are simply attempting to ascribe that set of intentional states that will have the greatest predictive power for the behavior of the individual concerned. Qua elements of a predictive strategy, intentional states are "real" but they are not the sorts of "things" that stand in need of a naturalistic explanation. Still, Dennett holds, there are representations of some sort in the brain and it is the "aboutness" of these states, whatever they might be, which will need a biological explanation. Hence the middle ground occupied by the intentional stance: a commitment to the taxonomy of folk psychology but a denial that folk-psychological states are the correct subject of a naturalistic theory. Clearly, then, the naturalists are divided by some deep philosophical differences: there is no consensus among them about what property, exactly, requires explanation—whether this is the property of intentionality, of psychological content more broadly construed, or of representational content, writ large. Nor is there agreement about what sorts of states have "that" property—psychological states as delineated by ordinary ascriptions, psychological states as taxonomized by scientific psychology, neural representations, or computational states. Follow the sounds discord, emanating from the forest, and there you will find the naturalists' camp.

Their disagreements about the nature of mind and mental representation not withstanding, all of the above naturalists share a common project. Each is trying to provide a naturalistic explanation of psychological representation and each is concerned (wholly or in part) with how, in the natural order of things, representational content comes into being—how a psychological state comes to be about a property or object. In an ecumenical spirit, then, I shall call "the" property at issue *aboutness* and take the naturalists' project as one of explaining aboutness.

More specifically, most naturalists start with the view that aboutness is a relational property: for a representation to be about a property or object is for it to bear a specific kind of relation—call it the *aboutness relation*—to an object or property in the world. It is the nature of the aboutness relation, between a mental/computational/neural state and its object of representation, that the theory should explain.

Second, all of these theories are naturalistic in that they are motivated by the general conviction that our explanation of human psychology ought to be reconciled with our scientific understanding of the physical world. The goal is to explain this property of directedness as part and parcel of the physical world, as a natural phenomenon. In more concrete terms, this belief takes the form of a restriction on what will count as acceptable theoretic terms: the theory is limited to the terms of a natural science, a theoretic language devoid of semantic predicates.

Third, the naturalist usually begins by selecting a natural relation, one that will be used to explain the aboutness relation. Assume, for example, that a person is standing gazing at the Eiffel Tower. Those theorists who are attempting to explain the content of ordinary propositional attitudes might say of this case that the subject's perception is caused by the Eiffel Tower or carries information about it—they might posit an informational or causal relation between the subject's mental state and the Eiffel Tower. An eliminativist, on the other hand, who denies the categories of folk psychology writ large, might explain the event by positing a natural relation between a neural state and certain world events—perhaps it is the function of that type of neural state to indicate the Eiffel Tower or, more likely, to indicate towers in general. Naturalists agree, in other words, that, if a given mental/neural state is about *x*, there must be some natural relation (or relations) between that state of the individual and certain events in the world which ultimately explains this aboutness relation. What is contentious for the naturalist are two further issues: which natural relation ought to be chosen for this explanatory role and how, exactly, the posited natural relation(s) will be called upon to play that part. Will the natural relation be that of causation, of information, of indication, or will the natural relation be cashed out in terms of biological function? Is the intertheoretic relation, between the psychological realm and neurological/computational/biological realm, one of identity, of reduction, of supervenience, or perhaps something else? Here the views vary widely. What is generally agreed, however, is that the aboutness relation will be ultimately explained by a natural relation.

Fourth, virtually all naturalistic theorists agree on an important methodological point, namely, that a naturalistic theory should start with the static perceptual case. If we want to understand how our mental states come to be about the world, we should begin with those occasions on which our connection to the world is most clear cut—when, say, the subject sits, with his eyes open, staring at the world before him and has visual experiences of the forest beyond him. There is general agreement, in other words, that, if there are any simple cases of aboutness to be found, it is static perceptual events which will furnish them. To put this another way, we all recognize that the relations of our ordinary mental states to their objects can be very complex. We have thoughts about things that are presently inaccessible to us by perception; we can ponder the nature of objects that are fictitious or abstract; we can think about things with which we are only vaguely acquainted, or which we are unable to identify, or about which we may have false beliefs. I can think about the Eiffel Tower even if I have never seen it; I can wonder how a certain George Smith is doing even though there are a number of philosophers named 'George Smith'; I can ponder the nature of causality or irrational numbers. These, we agree, are the difficult cases, the ones we ought best leave aside, at least until there is a story to tell about the simple case—when the object of perception is physically present to the subject and when he has some thoughts or representations that must surely be about that object.[2] It is the simple perceptual case that will, in the end, ground the aboutness of complex representational states.

The four broad assumptions, then, that naturalistic theories of aboutness usually espouse are these: that to give a theory of aboutness is to explain a relation between a representational state and some property or object; that this relation is to be explained in the terms of the natural sciences without recourse to semantic predicates; that this psychological relation will be ultimately explained by some relation from the natural sciences; and that the simple perceptual case is in some sense "basic," that it constitutes the most likely starting point for such a theory. This agreement, I think, is a surprising fact, given the diverse explanatory goals of naturalistic theories. If the goals of the authors are so heterogeneous, why do their theories have essentially the same broad theoretic form?

Enter the "traditional view" of the senses, our intuitive understanding of sensory function. To present the traditional view in its

[2] Of course, simplicity is not the only reason why the static perceptual case has seemed such an obvious starting point. Nor need this assumption hinge entirely on any view of the senses. For philosophers with an empiricist bent, who believe that the content of *all* of our mental/computational states is somehow wrought from the content of perceptual states, perception is not merely the most simple case of mental directedness. Perceptual states are also the most *basic*. If it is our perceptions that somehow provide the raw materials for our more complex intentional states (about unperceived, fictive, or abstract properties), then perceptual states *ground* all other intentional relations. For an empiricist, this methodological point of entry is dictated by his philosophical program. Even a rationalist, however, will find the study of the simple perceptual case a perfectly obvious starting point. All one need believe is that, in the absence of all sensory input, there would not be (could not be) any thought at all, that sensory input is a necessary condition for intentional states. On this view, perceptual states are also "basic," although not in the strong sense endorsed by the empiricist, given above.

strongest guise, think for a moment about the solipsistic plight of the brain. There floats the brain, the "I," encased and protected by the skull. It is the control center of the body, the mover of limbs, the planner of actions, the origin of all the body's thoughts and feelings. This central function notwithstanding, the brain resides in a kind of intracranial isolation. It has only mediated access, through its sensory and motor attachments, to its central concerns, to the body and the external world. But for a series of outgoing and incoming "wires," the brain is alone with its thoughts. It is because of the brain's solipsistic existence, then, that the senses have a clear role to play. They are the brain's window on the world. The senses show the brain, otherwise blind, how things stand, "out there," both in the external world and in its own distal body.

To flesh out this picture, consider one example of external perception, peripheral thermoreception, the system that reacts to surface skin temperature. Think about the sensations that we have of temperature—the bitter icy cold of a January wind, the pleasant coolness of a chilled drink in hand, the warmth of the spring sun, the singeing pain of a hot iron. From this internal evidence, one might guess that our thermal sensations are the products of skin receptors that are finely tuned to surface temperature. Like miniature biological thermometers, the receptors record the temperature of their immediate surround, its ups and downs. If, say, the skin temperature is eighty-nine degrees Fahrenheit, then the receptors must send a signal, some firing pattern, that "means" eighty-nine degrees Fahrenheit; if the skin is cooler, say eight-three degrees, the receptors must give an appropriate response (perhaps they fire less rapidly). The receptors, we think, must react with a unique signal, one that correlates with a particular temperature state. Thus, the brain gains information about what it is like "outside."

Although we realize that human thermoreception probably does not function exactly as described above (for we all know about the various temperature illusions), the example captures a certain feature of the traditional view: sensory systems must be veridical in some sense of the word. If the senses are the brain's window on the world, then any system worth its salt (and functioning correctly) ought to provide an accurate account of just how things are: the brain must be able to tell, from the signals it receives, how things stand in the world.

First, it would not do to have a sensory system that is fickle or unreliable in its reports. No one would want a thermoreceptive system that suddenly started to send a signal, usually reserved for a skin temperature of eighty-nine degrees Fahrenheit, at arbitrary intervals; nor should the system be such that, when the skin *is* eighty-nine degrees Fahrenheit, the thermoreceptors fail to register that fact, forget to send the signal at all. What good, qua informant of the brain, would such a system be? In nonmetaphorical terms, this aspect of sensory veridicality is usually expressed as that of *constant correlation*: if a signal is to be informative ("tell the truth"), it must be produced when and only when a particular stimulus (or stimulus set) is present.

Second, if the brain is to have an accurate understanding of how the world lies, then constant correlation is not enough. The relevant *structure* of the external events or properties must also be preserved by the sensory signals: the representational relations among the sensory signals must mirror the relevant relations in the sensed domain. In the case of thermoreception, the brain must be able to discern a one-dimensional relation between temperatures—whether this temperature is greater or less than some other one. So a thermoreceptive system must be, well, thermometer-like: the relations among the temperature states in the world ought to be at least roughly discernible from (if not strictly isomorphic to) the relations among the individual thermoreceptive signals. Constant correlation alone is not very helpful, especially if the question before the brain is whether or not to pull a hand away from the warming fire (that is, is the hand getting too warm now?).[3]

Third, a veridical sensory system is also a *servile* system (or if you prefer, a system that acts as the brain's loyal retainer). By 'servile', I do not mean a system that reacts in only strictly prescribed (law-like) ways to the world's sensory impingements. Much of sensory processing, we now believe, involves active or "top down" processing: the brain uses stored information, default assumptions, or hypothesis generation and testing in order to solve the computational problem at hand ('What is that object'?) or to increase its efficiency. Quite obviously, this kind of activity does not challenge our intuition of veridicality: such sensory systems are still trying to represent the world, in the face of limited information, as accurately as possible. What would not jibe with our notions of a proper (veridical) sensory system is one which actively embroidered upon the nature of the impinging sensory stimuli or which simply made things up. What we expect from sensory systems, in other words, is that they are the brain's

[3] For a well-worked out example of the notion of isomorphic representations, see C.R. Gallistel's *The Organization of Learning* (Cambridge: MIT, 1990).

"ontological drones." In the service of the brain, they toil tirelessly to report the "what, when, and where" of the world's events. They do not interject their own opinions into their reports; they do not slyly skew the information to reflect their own interests or prejudices. Their job is to state the facts. To say that a veridical sensory system must be 'servile', then, is to say that it represents the world as accurately as possible, without embroidery or fiction, given the information available. (I take servility to be a slightly different notion than that of reliability as given above, though perhaps it is merely a subtype. By analogy, the problem with the town gossip is not that, for any event in town, she might fail to have something to say about it. What we doubt is the *accuracy* of what she says. One worries that it will incorporate her own likes and dislikes, prejudices, and interests.)

On the traditional picture, then, the senses, using a system of signals that captures the structure of a domain of external properties, tell the brain, without exaggeration or omission, "what is where." It is this view of the senses which dovetails with the philosophical problem at hand. The naturalists' project, stated in its most general form, is to explain the psychological relation of aboutness in some way that fits neatly with our scientific picture of the world. On the traditional view of the senses, there is also a relation of aboutness, a relation between any state of a veridical sensory system and its object. Indeed, fulfilling this relation is the raison d'être of the senses—that without which the brain would not know how things stood, either in its own body or in the external world beyond it. So what the traditional view of the senses provides, from the natural sciences, is exactly the sort of relation a theory of aboutness requires—hence an obvious starting place for any naturalistic theory of aboutness. Not all intentional states, the naturalist freely admits, are sensory states; nor need the subject have ever had sensory contact with a given object of thought; nor indeed are the intentional objects of such states necessarily physical objects or properties, the sort of entities with which the subject could have sensory contact. Still, the path to choose is clear, the naturalist will contend. Indeed, at first glance, it is the only path visible at all—the only path out from the solipsistic existence of the brain.

II. NARCISSISTIC SENSORY SYSTEMS

The problem with the traditional theory of sensory processing, I think, is that it is not universally or even generally true. In some cases, sensory mechanisms do behave as one intuitively expects, in accordance with the traditional view. More often than not, however, sensory systems fail to be "veridical" in the sense given above. Indeed, if one had to pick a single predicate to describe all sensory systems, it is that each and every sensory system is, well, "narcissistic."

As a first pass at explaining this metaphor, think of the usual case, the human narcissist, a person whose worldview is informed by exactly one question: "But how does this all relate to ME?" In a classic story, a narcissist goes to his therapist for his regular appointment. At the door he is met by another therapist in the same practice: there will not be a session that day, she informs him. The narcissist's therapist has been in a boating accident—she is alive, but in critical condition in the hospital. It is gently explained that there is a good chance that the therapist will survive and recover, lead a reasonably normal life, perhaps even return to her practice. At this, the narcissist looks stricken, and wails: "But why do these things always have to happen to ME?!?"

Needless to say, a narcissist's world picture, being informed by but one question, is strangely askew. In this regard, it is important to realize that the problem with the narcissistic worldview is not just that his range of interests is idiosyncratic, that self-interest directs his attention toward a limited portion of the world, a small part of what other people see. Rather, by asking the narcissistic question, the *form* of the answer is compromised: it always has a self-entered glow. In this case, it is obvious that the causal factors leading up to the accident are entirely independent of the narcissist's existence. But it does not seem this way to him. The narcissist cannot see himself and his relation to the world in an objective light, as one life among others. Hence his understanding of the emotions and actions of other people, and of events in the world in general, will necessarily incorporate his own particular interests. For the most part, it is not possible for the narcissist to stand back and remove himself from the picture—and that, of course, is exactly the property which gives the narcissist away. I do not mean to imply by this that the narcissist never gets things right, never sees the world for how it is. Sometimes the question 'How does all this relate to me'? is just the right question. 'What does that person want from me now'? is the appropriate thing to wonder in a dark alley. Hence it yields a legitimate—"veridical"—answer. It is nonetheless true that the narcissist's question shapes his worldview at all times, and this is so regardless of whether the question is "appropriate" or not.

Turn again to a simple sensory system, that of thermoreception. As I said above, our first guess about thermoreceptors is that they act like small thermometers, signaling the brain about the location and the temperature state of the skin. In fact, however, the nature of thermoreception is quite different (the following draws largely upon from the work of H. Hensel[4]). First, the apparently continuous temperature gradient that we feel is not the result of the continuous response of a single thermomechanism. Instead, our sensations are the result of the action of four different types of receptors: two thermoreceptors, "warm spots" and "cold spots," and two pain receptors (nociceptors) that fire only in the extreme conditions of very high or very low temperature. At very high and very low temperatures, we feel only pain, sensations that are qualitatively indistinguishable from one another. In the middle zone, we rely upon one or the other kind of thermoreceptors for our sensations of warmth and cold. (The temperature at which warm spots begin to respond is roughly the same temperature at which cold spots leave off.) Second, there are far more cold receptors than warmth receptors, the exact ratio differing from location to location. So, for example, on the face, the nose has a ratio of 8:1 cold to warm spots, the cheeks and the chin are somewhat more sensitive to warmth with a ratio of 4:1, and the lips (entirely counterintuitively) are sensitive to only cold, with almost no warm spots whatsoever. One result of this variability is that different parts of the body are more sensitive to heat or cold than others. Because conscious sensation is the result of cumulative neural response, the more receptors there are, the more cumulative neural activity. This is a fact that will strike you as immediately plausible if you imagine wading into a cold lake. As a matter of fact, some steps *are* harder to take than others. Another result of this variability is that the "temperature of neutrality"—the "comfort zone" as thermostat makers call it—differs from one area of the body to another. Whether a temperature feels comfortable is as much a function of location on the body as it is of temperature.

What do the individual receptors do? Each kind of receptor, warm spots and cold spots, have both a static and a dynamic function (figure 1). The receptor's response to a constant temperature is its static function, represented by a curve that plots the rate at which the neuron fires against the stimulus temperature. For both the warm spots and cold spots this response is nonlinear; the static functions for the two receptors also differ, one from the other. The warm receptor responds over a narrow range of temperatures with a steep rise in firing rate at the high end; then the firing abruptly halts at a maximum temperature. The cold receptor has a less intuitive response pattern. It has a wider window of response, a gentle curve with a maximum response at

[4] *Thermal Sensations and Thermoreceptors in Man* (Springfield, IL: Thomas, 1982).

(a) Warm "Spots"

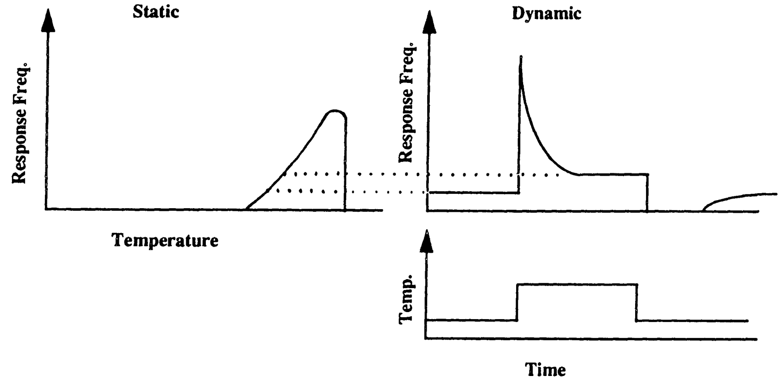

(b) Cold "Spots"

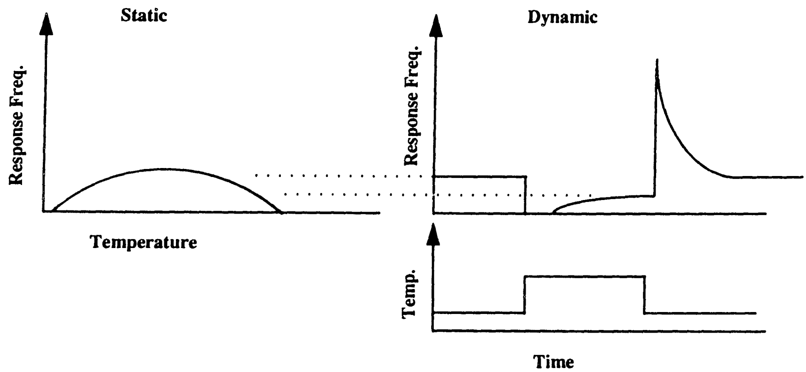

Figures 1(a) and 1(b). The static and dynamic functions of the warm and cold receptors. To illustrate dynamic function, both receptors were subjected to a sudden temperature increase and then a sudden temperature decrease. The warm receptor shows a dynamic response to temperature increase alone; the cold receptor shows a dynamic response to temperature decrease alone. (Adapted from Hensel, *Thermal Sensations and Thermoreceptors in Man.*)

the midpoint, tapering off thereafter. The static functions of neither the warm spots nor the cold spots are thermometer-like, with a certain set increase in firing rate per degree of temperature change.

Both kinds of thermoreceptors also have a dynamic function, a response to temperature change. When the temperature of a warm spot is increased, a burst of activity occurs; then the firing rate gradually slows and settles into a new higher base rate, a rate determined by the static function. For example, after shoveling a snowy walk, putting your hands under only tepid water initially feels very warm. The dense liquid causes sudden energy transfer, a sudden increase in temperature, and the dynamic function's burst of neural activity. When the temperature of your hands stabilizes, the neural activity also decreases and the water feels as it normally would, cool. What is important here is that the size of the initial activity burst is *variable*: it depends upon the starting temperature. So, for example, if you transfer your hand from tepid to warm water—say, a forty degree Fahrenheit change in temperature—there will be a sudden burst of activity in the warm receptors. But if you then transfer your hand from warm to hot water, again a difference of forty degrees Fahrenheit, the dynamic burst will be much greater. The higher the starting temperature, the larger the initial burst of the warm receptor. Thus, the transition from warm to hot water is felt more keenly than the change from tepid to warm water. Note also that when a warm spot is cooled there is no neural burst at all, only a gradual decrease in firing rate until the new lower base rate (determined by the static function) is reached. These same principles apply to cold spots as well, except in reverse. A *cold* stimulus evokes a dramatic over-compensatory response; a *warm* stimulus causes a gradual decrease in activity; and the *lower* the initial skin temperature, the greater the initial burst of activity in response to a cold stimulus. (These opposing characteristics are, in fact, the defining functional qualities of warm and cold receptors.)

As one can see, the response properties of the thermoreceptors are complex. There is no single thermometer-like receptor or even two "abutting" thermometer-like receptors; the warm spots and cold spots exhibit two different nonlinear (static) functions; the transducers exhibit adaptation and behave in a context-dependent fashion (the dynamic response depends upon the current base-line skin temperature). What is more, our temperature sensations do not result directly from the functional properties of individual receptors but are determined by the characteristics of the neural population as a whole. Where there are more cold receptors, the area seems colder; if a location has few warm spots, heat is hardly noticed at all.

All of this seems somewhat strange on the traditional view of sensory processing, of thermoreception as a system that disinterestedly records temperature facts. Just how inept could this system be? Viewed as narcissistic, however, the system makes perfect sense.

What the organism is worried about, in the best of narcissistic traditions, is its own comfort. The system is not asking, 'What is it like out there'?—a question about the objective temperature states of the body's skin. Rather, it is *doing* something—informing the brain about the presence of any relevant thermal events. Relevant, of course, to itself.

In this light, reconsider the old illusion created by placing one hand in cold water, the other in hot, and then, and after a few minutes, placing both hands simultaneously in some tepid water. Stupid sensors. They tell you that the tepid water is two different temperatures. But the sensors seem less dull-witted if you think of them as telling you how a stimulus is affecting your skin—that one hand is suddenly cooling while the other is rapidly warming. Skin damage can occur from extremes of temperature (being burnt or frozen) but also from rapid temperature changes alone, even if the changes occur within the temperature range for healthy skin. So the information provided by the dynamic function—whether there has been a change in temperature, in which direction and how rapidly that change is occurring—is crucial to the organism for avoiding skin damage. It allows you to pull away your hand before it is burnt, to remove your hand when the change in temperature will be fatal to the skin.

From this perspective, one can also see why it makes good sense to have a dynamic function that is context dependent—why the burst rate depends upon both the amount of change and the starting temperature. If the skin temperature is already high, any increase is likely to cross the upper limits of skin tolerance. So the extreme bursts of the dynamic function, produced by stimuli at the upper end of warm-spot function, warn that tissue damage is about to occur. (If the temperature sensitive nociceptors fire, then damage has already been done.) In other words, the system produces strong signals whenever potentially dangerous temperature changes occur. From this perspective, the utility of the dynamic function is clear.

As another example, think again about a cold mountain lake—this time, about putting your head into a cold mountain lake. Most of us have had the experience of going for a hike and then, after a strenuous day, coming upon an inviting-looking lake. First you put your hands in the water to test the temperature; then, after the initial shock of is over (after the dynamic burst of the cold spots has subsided), you decide to risk a swim. Note, however, that making your brain very cold is not a good idea. Heat loss through the scalp is one of the many ways that human life comes to an end. So, from an evolutionary perspective, it is not surprising that the scalp has an abundance of cold receptors—that when you dive in, the water feels colder to your head than it did to your hands earlier. Given how dangerous heat loss can be, it makes sense to receive a sharp warning signal whenever the scalp is cooled, before you loose the sense to react rationally. More generally, if one looks at the complexity of the human body and at the large range of behaviors we are capable of performing, it is clear that the importance of thermal stimuli can vary, depending upon where the stimuli occur. You need to keep the top of your head warm; you need to make sure that your nose is not frozen in weather where your cheeks will still be fine; you need to be able to sense both cold and warm objects with your hands. Hence the evolutionary "solution" of a variable ratio of cold to warm spots. Through a difference in the number of cold and warm spots, the cumulative signals are tailor-made to reflect the relevance of temperature stimuli for a particular place on the body.

Looking back at thermoreceptive function, then, one realizes that this system is not merely inept, a defective indicator of surface temperature. Rather, the system as a whole constitutes one solution to man's various thermal needs—that he be warned when thermal damage is occurring or before it is likely to occur, when temperature changes are likely to have specific consequences, and so on. These are the thermal problems which have been "posed" over time to the evolving organism and which have resulted in the idiosyncratic functions described above. In other words, this system, qua system, reflects a constellation of interacting facts about one species—the environmental conditions encountered by Homo Sapiens (for example, the thermal properties of the environment), the developed physiology of the species (for example, the mortality conditions of the skin, the details of internal and external thermoregulation), and the range of behaviors of which the species is capable.[5]

Let me now make explicit how the above characterization of thermoreception fails to fit the traditional view. Recall from the ways in which, according to the traditional view, sensory systems are veridical: each signal must correlate with some property (or range of properties) in the world; the structure of the relevant relations among the external properties must be preserved in a systematic encoding of those relations; and a sensory system must be servile in that it

[5] To make this point more tangible, try to imagine the thermoreceptive system of some other creature—say, the penguin.

must be seen to be reconstructing, without fiction or embellishment, the properties, objects, and events of the world external to the brain. In this case, the property of interest is, obviously enough, skin temperature at a specific part of the body. So the question that one would expect the thermoreceptive system to answer is: 'What is my skin temperature at *x*'? But this is not what it does.

First, particular thermal sensations do not necessarily correlate with any particular temperature or any particular temperature change. Because thermal sensations are a function of the firing rates of a neural population and because the absolute number and ratio of the two different receptors differ from one part of the body to another, exactly the same skin temperature can give rise to a variety of sensations. This is one reason. For another, the elliptical static response properties of cold receptors ensures ambiguity in its signals. As the skin becomes colder, a cold receptor fires more and more rapidly until the receptor's maximal response level is reached; then, as the temperature continues to drop, the firing rate starts to decrease. So a single cold receptor will fire in exactly the same way for two different stimuli, a temperature at the low end of its response range and a temperature at the high end. Necessarily, its signals (alone) are ambiguous. Nor do thermal sensations reflect temperature change. Here, it is the context dependency of the dynamic function of both warm spots and cold spots which presents the problem. The felt change in temperature for a specific temperature change will depend upon the starting temperature of the skin. If the temperature of a warm spot is increased at the bottom of its response range, the dynamic burst will be small; if it is warmed at the top of its response range, the burst will be very large. So neither absolute temperature nor temperature change is recorded by thermal sensations.

Second, for almost the same reasons, thermal sensations as a whole do not reflect the structure of thermal stimuli, whether stimulus T1 is greater than, less than, or equal to stimulus T2. The water, as we wade into it, initially lessens the skin temperature of each body part about equally; but it certainly does not feel that way. Some parts feel much colder. Similarly, cold receptors, with their elliptical static response function, represent disparate pairs of temperature stimuli as being equal—very low ones and relatively high (tepid) ones. Even warm receptors, with their nonlinear static responses (which abruptly stop at a certain maximal rate), do not quite get the picture right. Lower temperatures do elicit lower firing rates but the differences in temperature between the warmer and hotter stimuli are not uniformly recorded. Linear temperature increases are encoded as nonlinear changes.

Third, the thermoreceptive system embroiders its account of the temperature states of the world. Unlike many other sensory systems, it is probably unfair to say that this particular system actually manufactures fictions in the course of its ordinary function, but it does appear prone to chronic exaggeration. At the lower and upper limits of response for the cold spots and warms spots, respectively, a small temperature change elicits a hysterical response. Moreover, given the differing ratios of cold spots to warm spots, no body part simply reports its surface temperature. Each body part exaggerates its own state in accordance with of its own interests and sensitivities. So the thermoreceptive system is hardly the ontological drone that the traditional view had imagined, the tireless reporter of the "what, where, and when" of surface temperature.

It is time to step away from the details of human thermoreception in order to make some more general points about sensory function. In imagining the function of any sensory system, we slide very easily from the question 'To what are the receptors responding'?—a question about the nature of the proximal stimulus that will evoke a receptor response (for example, mechanical deformation or light in a certain range of wavelengths)—to the question 'What do the signals of the system detect'? In the grip of the traditional view, we treat these as one and the same question. But despite the entirely rule-governed nature of sensory systems, the question 'What is the system detecting'? may not be apt. Rather, the appropriate question, for any sensory system, is 'What is the system doing'? and by this we mean doing for the organism. This point requires some explanation.

Each and every sensory system, no matter how sophisticated or simple, is tied to a set (sometimes a very large set) of behavioral tasks. No matter what else the senses do, in the end, they must inform movement or action. In the case of behaviorally simple creatures, there may be a very short route from the response of the sensory receptors to the output of the motor system. In the slug, no doubt, a dynamic burst in a population of receptors sensitive to sodium chloride will trigger an immediate aversive reaction, movement away from the stimulus. Here, the utility of matching the information recorded by the senses to the needs of the motor system is clear. There is no need to represent the world "the way it is," for simple lives (behavioral repertoires) require limited information. Indeed, evolution will favor sensory solutions that package the information in efficient and quickly accessible formats, in ways that match the particular physical form of the motor system, its motor tasks, and hence informational requirements. For every evolved system, there will be a symbiotic relationship between the information gathering of the sensory system and the informational needs of the motor system—and the elegant solutions that evolution eventually selects need not involve any straightforward (to our eyes) "veridical" encoding of sensory information.

In our own case, of course, we think before we act. We are not simple stimulus-response organisms, with our behaviors triggered by simple sensory signals. Moreover, our perceptual thoughts are genuinely intentional, tied as they are to the objects and properties of the external world. I sit and stare at my empty coffee cup; I wonder whether I would like another cup, or at least, want one enough to get up and make one; I ask my husband whether he would like a cup of coffee. (And so on—philosophical drama at its finest.) These two facts, that our responses to the world are not stimulus driven and that our intervening thoughts are about stable objects and their properties, may lure us into thinking that somehow human sensory systems (and those of other behaviorally complex creatures) are of a different, "veridical" sort. Our senses could not be narcissistic, or at least, not for the most part. (Let us dismiss thermoreception as some crude and embarrassing vestige of our primitive past.)

Neither of these two facts, however, obviates the need to direct motor behavior in a timely and efficient fashion by means of sensory information. Ruminate about coffee all you like, but should a full cup appear, you will have to reach out your hand and grasp the cup. From your perspective, of course, the point of view of the conscious subject, it may seem to you as if you look out at the world, contemplate the cup of coffee on the table, and then decide to pick it up—as if your unitary perception of the coffee cup qua cup itself unmediatedly directs your hand. After all, if the result of all those visual processes is a perception of a cup, surely that (unified) perception then informs your actions. No matter. What directs your movements, from the moment you begin to lean forward in anticipation of reaching with your arm to the moment the coffee cup makes contact with your lips, is a host of information from numerous different sources. There will be *visual information* about the egocentric position of the cup relative to your body, about its position relative to your reaching hand, about the shape of the handle relative to your grasp, about the cup's rotation (shape) relative to a horizontal plane as you pick it up (do not spill it now), and about the cup's speed of movement; there will be *proprioceptive information* about the position of your upper body as you lean forward, about the angles of your arm joints as you stretch toward the cup, about the weight of the cup relative to the firmness of your grip, about the fantastically minute adjustments of your fingers, hand, and arm muscles as you balance the cup of liquid in an upright position, about the position of the cup and your hand relative to your lips (after all, you do not have to stare cross-eyed to get the cup there); there will be *tactile information* about the pressure of the cup in your hand, the pressure of the cup on your lips, the shape of the cup in your hand. No doubt this is but a small part of what is actually involved in the "simple" activity of picking up a cup of coffee. Even our simplest actions, then, involve numerous sources and types of information (here, visual, proprioceptive, and haptic information) and, within a single system such as vision, specialized information (about shape, position using a variety of reference frames, rotation, movement, and so on) which requires diverse representational schemes. Thus, if one looks at the neuroanatomy of mammalian vision, most of the physically distinct sites, each associated with a specific informational problem, have connections with one or more motor site(s) and all visual areas have connections with subcortical sites which themselves have lines to motor areas. Diverse and complex information subserves simple motor movements. We, as perceivers of cups and saucers, then, tend to mistake our conscious perspective, about coffee and cups, for insight into how things work, how we actually manage to raise the cup. But this is not how it is. Although, from the first-person perspective, what you see is a cup sitting on the table, what actually guides your movement is a plethora of sensory signals. The upshot is that for us, as much as for other creatures, the symbiosis between sensory information and motor needs is equally strong. If the sensory information is to guide the requisite motor movements, it must be usable on-line, by numerous feedback systems, control loops, and cognitive "interrupts." All the sensory information must be encoded in a motor-friendly way. Even as intentional, conscious perceivers, we are equally in need of narcissistic sensory strategies, indeed, perhaps more so given the complexity of our behavior and bodies.[6]

When we examine a sensory system, then, we are looking at an evolved solution to a specific informational problem. Perhaps it is a

[6] For a very similar view of sensory processing, see Patricia Churchland, V.S. Ramahandran, and Terry Sejnowski, "A Critique of Pure Vision," in C. Koch and J.L. Davies, eds., *Large-Scale Neuronal Theories of the Brain* (Cambridge: MIT, 1994): 23–60.

system that is able to increase or decrease the sensitivity of its receptors in response to present needs (without recording the change in sensitivity), that applies filters for the creation of specific features, that adjusts an outgoing signal in order to ensure certain characteristics in the signal return—or that encodes the facts about skin temperature in some way designed to aid creature comfort. The relevant question to ask of such a system is 'What is it doing'? This is a question to which the answer might be 'it is measuring, with variable discrimination, the animal's tilt away from the vertical in order to maintain its upright posture', 'it is providing a visual signal that can be easily processed for movement information', 'it is monitoring the stretch of the flexor muscle in order to adjust the length of the tensor by an equal amount', or 'it is indicating an edible insect'—answers that may or may not make reference to processes of veridical perception.

III. DEFENDING THE TRADITIONAL VIEW: THREE OBJECTIONS AND REPLIES

In the face of the above kind of neurophysiological example (for human thermoreception is merely one among many, chosen for its familiarity), there are a number of responses one might give in order to vindicate the traditional view. I shall give three of the most common replies. The first response is, in effect (but not in intention), an a priori defense of the traditional view; the second response is a straightforward denial of the thesis that sensory systems are, at bottom, narcissistic; the third response, the most promising of the three, grants that sensory systems are narcissistic—but it denies any conflict with the traditional view.

A. The a priori defense. Given the above example of thermoreception, one might take away quite a different lesson, not a lesson about sensory processing per se, but a metaphysical view about the nature of *properties.* Dennett, for example, in talking about our conscious perception of secondary properties, takes the clearly self-interested nature of sensory systems to show that their normal function can *define* properties in the world. It does not matter that our sensations of, say, redness fail to correlate with any easily delineable set of surface reflectance properties, a single property that is sanctioned by the scientific image. There is redness in the world because we have a disposition to respond in a "discriminating" fashion to "red" stimuli.[7] That is, if a sen-

[7] In *Consciousness Explained* (New York: Little, Brown, 1991), p. 382, Dennett says: "What property does Otto judge something to have when he judges it to be pink? The property he calls 'pink'. And what property is that? It's hard to say but that should not embarrass us, because we can say why it is hard to say. The best we can do, practically, when asked what surface properties we detect with color vision, is to say, uninformatively, that we detect the properties we detect. If someone wants a more informative story about these properties there is a large and incompressible literature in biology, neuroscience, and psychophysics to consult."

sory system responds in one and the same way to a rather diverse set of properties and conditions, then that disorderly set simply *is* a property of the world. I am not sure that Dennett himself would wish to frame the view in this general form, but let us assume this strong view for the sake of argument.

Setting aside the merits of this view qua metaphysical thesis about properties, note that as an answer to the question 'Are sensory systems veridical'? it is empty. If properties are defined by the ordinary causes of sensory signals, then, by definition, all sensory systems are veridical: the system "captures" the structure of the domain of external properties just because it *defines* that domain and, of course, any signal reliably records, without exaggeration or omission, whatever stimuli ordinarily cause it. It is only against a firm prior ontology of properties, one that is specified and justified independently of actual sensory response, that questions about veridicality have any bite at all. Thus, on this reply, what seemed like a prototypically empirical question—'Are the senses veridical'?—is recast as a question of definition, the answer to which has the status of an analytic truth. This alone, of course, is not a reason to dismiss the response—perhaps the traditional view ought to be granted the status of analytic truth—but this seems to me unlikely. (Below, in answering the third objection, a more decisive objection to this response will emerge.)

B. The appeal to signal information. Whether or not a sensory state appears to correlate with a particular external property, this response begins, sensory states nonetheless *carry information* about their causes. Take the case of thermal sensations. Although any given thermal sensation might be caused by a large variety of temperature states, a given thermal sensation still carries information about its cause. It just does not wear this information on its sleeve. After all, thermoreceptors act in orderly ways—the response functions of both warm spots and cold spots are shown in figure 1 and there is some function that sums the responses of individual receptors. So starting at the beginning of the process, with the initial state of the thermoreceptive system (the current static response rate of the receptors), plus the neural population characteristics (the ratio of warm spots to cold spots and their numbers) and stimulus temperature of the skin, one could predict the sensations of the subject. Conversely, starting with the resultant thermal sensation (or population response), plus the initial state of the system and the neural population characteristics, one could compute the value of the stimulus temperature. In other words, it is possible, using the appropriate calculations, to deduce skin temperature from the thermal sensation under standard conditions—and this is just what it means to say that a signal (thermal sensation) carries information about a source (skin temperature). Hence, in this sense, an informational sense, thermal sensations do reliably indicate temperature states.

There is certainly some truth to this response. Given the causal regularities that govern sensory systems (for there is nothing supernatural about them), the sensations or signals they give rise to will often carry information about their causes. That is, whenever there is a computable function that describes the input-output relation, the response of a sensory system will carry information about its causes. In the case of thermal sensations, sometimes our sensations carry information about skin temperature and temperature change, sometimes not. This is because the static response of the cold receptors does not define a computable function, hence there is no algorithm that could determine, given the population response of the cold-spots (a sensation of cold), the stimulus temperature.[8] Indeed, a bell-shaped response curve for sensory receptors is more the norm than the exception, but let us put aside this fact for the moment. The question that concerns us here is whether, given that information about the stimulus is often carried in the sensory signal, this will be of any practical use in constructing a theory of aboutness.

The problem here is that it makes little sense to identify the contents of sensory states with whatever information they carry. Return to the example of wading into the lake. There you stand in water up to your thighs, debating the relative advantages and disadvantages of total submersion. (Would not a bit of splashing and swishing be just as effective, you wonder?) Eventually, you wade in, gasping and sputtering at various intervals. Now, on the informational view, all this silliness is for naught. The dynamic bursts of the various cold-spot populations, no matter where they may lie, all signal exactly the same thing—a single temperature for the water or, if you like, a single temperature change. This is the information carried objectively by each population of cold spots—information that you already had when you waded in up to your thighs. Of course, *you* may behave a little oddly, as you immerse yourself, but there is nothing odd about the behavior of your thermoreceptive system. It reliably indicates a single change in skin temperature, no matter what parts are presently getting wet.

This, of course, is a parody of the information response, but it makes a serious point. Information that is carried by, but not en-

[8] My thanks here to José Antonio Diez Calzada for clarifying this point.

coded in, a signal is information that is available only *in theory*. To say that the information is *present* is to say only that there exists a computable function that, if used, would yield the correct result—in this case, the actual temperature change. It is present, as it were, from the point of view of the universe. But no creature has ever acted upon information that is available only in principle. Thus, to posit a theory of aboutness based upon signal information alone is to a posit a theory that, by hypothesis, makes no connection with an organism's action, behavior, or thought. I take it that this is not what the objector had in mind.

Rather, when an objector points out that even a narcissistic response contains information about its causes, the hope is that somewhere down the line, this information is *extracted*. Despite our conscious narcissistic sensations or the narcissistic responses of sensory systems, somewhere higher up in each sensory system the appropriate calculations are made and the true state of the world is inferred. (Thus, for example, in case of human color vision, the ambiguity inherent in the bell-shaped response curve of each cone—ambiguity between intensity and frequency information—is partially resolved by the overlapping range of responses of the three cone types. Through comparison, ambiguity is resolved.) How else could a sensory system be of use? This suggestion, however, amounts to little more than an expression of one's faith in the traditional view. Empirically, there is little reason to think that all sensory systems carry within them the means to "decode" their own responses.

Once again, in the case of thermoreception, there is no anatomical or physiological evidence to suggest that the thermoreceptive system has the additional information needed to compute stimulus temperature—that it has knowledge of the number of receptors at each location, or the nature of the static and dynamic functions of the receptors, or that it keeps track of the system's initial state. Nor is there any further psychological or behavioral evidence that our thermoreceptive systems have this information or make these calculations. Given a thermal sensation alone, we have no capacity consciously to "see through" the narcissistic thermal sensation to the objective skin temperature (or temperature change) nor, behaviorally, do we act as if (unbeknownst to us) our brains make these kinds of calculations. This information is simply not available to us through sensory means. (Of course, there are many other ways that one might use thermal sensations to infer temperature. After reading Dr. Spock, you will know what temperature of milk correlates with a certain warm feeling on your wrist. Knowing this, the sensation can be used as an indicator of a particular milk temperature. But to make this inference requires exactly those cognitive capacities which the naturalistic theory hopes to explain—intentional thoughts about Dr. Spock, baby books, wrists, sensations qua sensations, and so on. The thermoreceptive system alone does not provide this information.) The point, here, is a general one: when a sensory system uses a narcissistic strategy to encode information, there need not be any counteracting system that has the task of decoding the output state. If the very point of a narcissistic encoding is to match the incoming information to an organism's behavioral needs, then for most part, there is no reason to re-encode the sensory information into a veridical format.

C. The appeal to detector cells and biologically salient properties. Of the three objections, the third is the most promising as well as the most obvious: Granted that sensory systems utilize narcissistic encodings, why not say that the function of a sensory system is to detect *narcissistic properties*, properties that are defined relative to an organism's interests? The shocking cold of the lake that you feel on your scalp, for example, is a property, one that *that* sensation (population response) detects, namely, the property of being-too-cold-for-my-head. The same will hold for the properties detected by similarly narcissistic mechanisms. To put this another way, when the question 'What temperature states do thermal sensations indicate'? was asked, the answer was restricted to the temperature readings of some scientific scale, degrees Fahrenheit or Celsius or Kelvin. On this division of the world, the answer was 'no': one thermal sensation does not indicate any single skin temperature (or temperature change). But if we look to the neuroethological or neurophysiological literature on sensory function, we find descriptions that make use of a variety of biologically salient properties, properties that are not described by the predicates of physics and chemistry. We find, that is, descriptions of "detectors" in simple organisms, some of which respond selectively to "legitimate" properties, such as magnetic north, being a complex sugar or a certain amino acid, or having a certain wavelength, others of which respond to "messy" properties, such as vertical symmetry, small flying insects, movement in the left of the visual field, or being a poisonous substance. Moreover, referring to these messy properties is essential to characterizing the function of these detectors; they are an ineliminable part of the neurobiological description. Thus, if we recognize "narcissistic" properties, along with "legitimate" and "messy" properties, as biologically salient, we can recast sensory systems in a way that conforms to the traditional view—as detectors or reliable indictors of external properties.

What makes this a particularly good objection, of course, is that much of it is right. I take it as incontrovertible that there are neurons that act as detectors, and that these detectors are tuned to properties such as predator and prey. Biologically salient properties are exactly those properties which neuroethology and neurophysiology appeal to in describing sensory function and those without which sensory function qua function could not be characterized. Moreover, prima facie, there is no reason to disbar narcissistic properties from the club, if neurobiological description demands them. Admitting all of this, however, does not give the critic what he needs.

Notice, first, just how strong the suggestion is qua neurobiological thesis. The claim is that each and every sensory system functions to detect properties, be they narcissistic properties (defined relative to organism's needs), biologically salient "messy" properties (for example, the property of vertical symmetry), or "legitimate" properties (those recognized by the other physical sciences, say, the property of containing NaCl). Call this the *detection thesis*. This is an extremely strong universal claim about sensory function and hence about the form of all our future biological explanations of the senses. Moreover, it is a universal claim made in advance of the lion's share of empirical research on, and biological theorizing about, a vast topic—a claim, the only present empirical basis for which could be a few notorious examples from the comparative neurobiological literature (for example, the famous fly detectors of the frog). It is a testimony to the strong intuitive pull of the traditional view, I think, that the prematurity and all-encompassing nature of the detection thesis seems to have escaped most of its backers.

Note that, even if one were to take all of the present literature on sensory processing and think up narcissistic or messy properties for those systems to detect, this would not provide any evidence for the universal claim. There are, after all, many devices, natural and human-made, for which one could claim a rough function from input to output states, for example, for vacuum cleaners and stereo amplifiers, as well as lungs, livers, and intestines. With a bit of imagination and a good sense of humor, one could think up some "messy" or "narcissistic" properties for these devices to detect. What the objector is claiming, however, is not merely that we could conjure up properties for each system to detect. He is making the stronger claim that neurobiologists will usefully describe the systems as such—that characterizing sensory systems as detecting properties will *always* provide us with some further insight into how or why the mechanisms function as they do. In other words, the detection thesis is committed to the view that our best biological explanations will characterize sensory systems as detecting. As I said above, this is a very broad empirical claim, especially when one realizes the infant state of neurobiological research on the senses. On what grounds might someone believe it to be true?

Above, in giving the example of the human thermoreceptive system, I argued that the neurobiological research does not support the traditional view nor, for that reason, the detection thesis—that we must ask the broader question 'What is the system doing'? not merely 'What is the system detecting'? In the literature of that nascent science, one finds, along with references to the now-famous detector cells, a variety of other descriptions. There, the neurobiologist begins his research by trying to pinpoint the informational problem(s) that the system must solve (a good alternative to poking blindly about) and then looks for mechanisms that solve them. Here, because the problems and answers are often framed in standard engineering terms, function is explained using standard engineering concepts. Thus, one reads about mechanisms that "turn up the gain," "act as a resistor," "apply a cut-off filter," "use a step-function," "shift a spectral sensitivity," and so on. Surely, the objector says (or must be saying), we could profitably recast this "engineering talk" in terms of detection and indeed surely we shall, once the full explanations are in hand. I think not. Let me give one short example of a sensory process for which this sort of "recasting," far from illuminating function, would positively obscure it.

Consider our proprioceptive system, the sense through which we know where our limbs are in space (without looking, that is).[9] In this system, the receptors are not external transducers, ones affected by events beyond or on the skin.[10] Rather, our sense of proprioception relies upon muscle spindles, internal receptors that are stimulated by the mechanical stretching of the muscles at our joints. In this system, the "engineering" dilemma that arises is that the range of activity to be encoded exceeds the ability of the neurons to respond. The el-

[9] For an extremely interesting account of what it is like to live *without* proprioception, that is, with only visual knowledge of body position, see Jonathan Cole's *Pride and a Daily Marathon* (Cambridge: MIT, 1995).

[10] The following information on stretch receptors comes from T. Carew and C. Ghez, "Muscles and Muscle Receptors," in Eric Kandel and James Schwartz, eds., *Principles of Neural Science* (New York: Elsevier, 1985, 2nd ed.), pp. 443–56, here pp. 454 ff.

bow joint, for example, has a wide range of angular extension; as you extend your arm from its fully bent to the straight position, the flexor and tensor muscles are shortened and lengthened considerably. Sensory neurons, however, have a fairly restricted range of response—there is a limit on just how quickly a neuron can fire (that is, four times a second). One has, then, what looks like a dilemma: if the muscle spindles are set to respond over the full range of joint movement, they will give only coarse-grained information about muscle stretch; but if spindles are adjusted to provide fine-grained information about muscle extension, they will respond over only a limited range of joint movement. What to do? The neural solution to this dilemma is a certain kind of feedback process: central control alters the sensitivity of the receptors, the response of the muscle spindles to change in muscle length. As the limb extends, and as the muscle becomes longer and longer, central control "lowers the gain" on the spindle response. This allows the spindle to maintain a continuing, fine-grained response throughout the entire range of muscle movement—although it does so at the "expense" of determining muscle length or extension.

This kind of problem, encountered by the proprioceptive system, is a very common one in sensory processing. Quite often the range of stimulation over which the system must respond far exceeds the ability of a sensory neuron to signal those events. (For example, the visual system of the cat, it is thought, responds over an illumination range of about sixteen log units!) Sometimes the solution to the problem is to build a variety of receptors, in order to cover the full range of stimuli; just as often, the solution is to use a mechanism that adjusts sensitivity. What is important to realize, here, is that there need not be any further device that records the "position" of the gain mechanism. If the purpose of a stretch receptor is to keep the flexor and tensor muscles balanced (as one stretches, the other contracts), the absolute muscle length is not what is important (at least not for this task). One could say, of course, that the stretch receptors detect muscle length if one wanted to give a rough characterization of what they do (after all, they do react to mechanical stimulation—the muscle stretching). But this way of talking quickly becomes misleading—say, in the face of motor disfunction—if one wants to explain why a person has tensor rigidity (or spasticity). For a careful characterization, one that will explain both function and misfunction, talk of detection only confuses the matter. Thus, while the proximal stimulus is the stretching of the muscle, the function of the muscle spindles is characterized more broadly, in terms of what the system is doing.

As I said above, the third objection depends upon a very strong and broad empirical conjecture about sensory function. Quite obviously, given that neurobiology is just getting off the ground and that, as a result, there are almost no complete accounts of any single sensory system, this question, about the correct principle of sensory processing, remains open. Still, the objector's claim is very strong, and (my guess is) not well grounded in any neurobiological evidence. Is there any other science, then, apart from neurobiology, that might prove the detection thesis?

A different defense of the detection thesis is based upon general evolutionary considerations. After all, the defense goes, organisms simply could not have survived if their senses did not function as reliable detectors of salient properties. If animals are going to find prey, avoid predators, and attract mates, then their sensory systems must detect these properties of the world. As Patricia Churchland[11] once said, in the course of defending an early theory of representational content called *correlational content*: "I take it as obvious that if there were no systematic relations between the external world and states of the brain, the animals could not survive or if they did it would be a miracle in the strict sense of the word. An owl would not know where the mouse is and he would not intercept it as it runs" (*op. cit.*, p. 260). Stating this view in an even stronger form, Dretske says:

> Without such internal indicators, an organism has no way to negotiate its way through its environment, no way to avoid predators, find food, locate mates and do the things it has to do to survive and propagate. *This, indeed, is what perception is all about* (italics mine). An animal's senses...are merely the diverse ways nature has devised for making what happens inside an animal depend, in some indicator-relevant way, on what happens outside. If the firing of a particular neuron in a female cricket's brain did not indicate the distinctive chirp of a conspecific male, there would be nothing to guide the female in its efforts to find a mate.[12]

Moreover, because evolutionary theory explains just those factors which affect a species' survival or demise, evolutionary biology will make reference to exactly those environmental properties to which the creature reacts (or fails to react). To illustrate, a typical explanation in evolutionary biology might go something like this: "at the beginning of the egg-laying season, the parasites choose with care their

[11] "Replies to Comments" from "Symposium on Patricia Smith Churchland's Neurophilosophy," *Inquiry*, XXIX (June 1986): 139-273.

[12] *Explaining Behavior: Reasons in a World of Causes*, p. 62.

egg-laying sites, rejecting less than optimal places; when the days grew shorter, signaling the end of the season, the parasites became far less discriminating, choosing exactly those kinds of sites which were rejected earlier." In the standard literature of behavioral ecology and evolutionary biology, even parasites are seen to recognize "optimal and suboptimal egg-laying sites," to detect the "shortening of the days" and "the end of the breeding season." What allows us to advance the detection thesis, then, is the simple fact that, in order to survive, all animals must react to salient properties of the environment—and evolutionary biologists cannot explain the factors affecting survival without characterizing the organism's reactions as such.

Again, there is much truth to the view. Of course, it is true that, as a whole, an animal's behavior must be directed toward certain salient objects or properties of its environment. Objects (and their properties) are important to the survival of all creatures. But from this fact alone, one cannot infer that the system of sensory encoding, used to produce that behavior, uses a veridical encoding. That is, it does not follow from the fact that the owl's behavior is *directed toward* the mouse or that the brain states bear systematic relations to stimuli, that there are any states of the owl's sensory system that are *about* or serve to detect that property of being a mouse. What is required for survival of the owl, for example, is that it finds its way to the mouse in a timely fashion and that, once it gets there, the owl can grab the mouse. To do so, the owl's visual system need not encode external space veridically, with anything like, say, a cartographer's topographic map, with the spatial relations of the world exactly mirrored by the spatial relations of the sensory map. Nor need the owl intercept the mouse by literally computing its trajectory based upon its present motion, nor need it use the deliverances of a mouse detector as part of the calculation. It is an open question what kinds of systems are actually used. In the case of the cricket, detector cells, for the chirp of the male conspecific, *are* used. But this is only one neural solution to one behavioral problem. Nothing about the directedness of an organism's *behavior* yields a firm conclusion about the directedness of the *internal states* of its sensory system.

Setting aside the question of whether there is any empirical evidence for the detection thesis, let me turn to a second response to it. Even if we granted the objection—even if we admitted that narcissistic properties will invariably figure in our neurobiological descriptions—this would not give the naturalist what he wants, namely, the beginnings of a theory of aboutness. The suggestion made is that states of sensory systems are about narcissistic properties, properties defined by relation to the subject's interests—*relational properties.* But what the naturalist wishes to explain, in the end, are representations that are about *objective properties* and events of the world. Whatever the neurophysiological facts about the thermoreceptive system, that is, we do not merely *feel* temperature sensations qua sensations in our skin. We have thoughts about the independent temperatures of objects. I put my hand out, grasp the coffee cup, and feel its warmth. I see it *as* a coffee cup and feel the cup *as* having a particular property, warmth. Similarly, when I put my hands into the dishwater, I feel *the water* as being warm, as having an objective property that is independent of my perceptions. Consider again the illusion produced by cooling one hand and warming the other, then placing both hands in a single bucket of tepid water. Here, the illusion, as you place both hands in the tepid water, is that *the water* is both hot and cold. That is what makes the exercise so surprising. There would be no illusion of a contradiction, however, if one assigned content as the objector suggests. Let one sensation indicate the property "precipitous drop in right-hand temperature" and the other sensation indicate "precipitous increase in left-hand temperature" and there is no explicit contradiction to be found. So the problem for the naturalist lies in the rather large gap between what our ordinary perceptions are about and the representational contents that the detection of narcissistic properties would assign, namely, mental states that are about "what is good for me," subject-dependent properties, as it were. To put this another way, the naturalist's hope was to find a relation between sensory states and external properties, a relation that would ground (in some undefined way) a theory of aboutness. But in order to save the traditional view, the objector introduces narcissistic properties and thereby gives up the link to the objective properties of the world. Prima facie, this is not a promising starting point for the naturalist theory. I shall have more to say about this later.

This same argument against the third objection, the detection thesis, also works against the first objection, the a priori defense. On the a priori view, recall, sensory mechanisms define properties of the world—large disjunctive sets of *objective properties.* Thus, on the a priori view, an individual temperature sensation defines a property, the disjunctive set of exactly those thermal stimuli which would evoke a particular sensation. The "hot" feeling will define one "property," a disjunctive set of thermal properties; the "cold" feeling will define another different disjunctive set. Does not this way of describing the properties "detected," as disjunctive sets of nonrelational properties, circumvent the problem? No. Even given this redescription of the properties, the a priori view fails to capture the content of ordinary thermal perceptions. When I put both hands into tepid water, the properties detected by the left hand and the right hand comprise two different disjunctive sets. *Ex hypothesi*, one temperature of water can bring about these two different properties in the course of normal function. Hence there is no contradiction between what my right hand says and what my left hand says (as there would be if there were no overlap between the disjuncts of both sets). In retrospect, this is exactly the problem one would expect given the form of the a priori defense: it simply takes a set of relational, narcissistic properties and redefines them in nonrelational terms. Defining the very same "properties" by means of a different device, however, does not alter the very nature of those properties. Thus, even though the a priori views picks out sets of objective properties with which to match our temperature perceptions, it, too, picks out the *wrong* sets.

IV. PHILOSOPHICAL IMPLICATIONS

So far, nothing that has been said casts any doubt upon our common-sense view about our perceptions qua intentional representations of the world. We have thoughts about the temperature of objects. We have thoughts about the Eiffel tower (real or fictive), about George Smith (the very one I met at Tufts), and the fragrance of gardenias (however vague or unreliable). However our sensory systems work, we are creatures who represent a world of objects, properties, and events. Granting all that has been said above, there is no reason, here, to question the existence of mental representations writ large. The philosophical problem of aboutness still stands.

On the other hand, if the above arguments are to believed, sensory systems do not seem to help us understand this fact. Let us take stock of the situation. First, sensory systems exemplify (usually quite elegant) solutions to specific processing problems. Sometimes reliable correlations are used (as in the case of the frog's proverbial fly detector) and sometimes they are not (as in the case of the thermoreceptors). The wiring fits the problem. Second, these kinds of nonrepresentational systems are characteristic of not only simple creatures—for example, the ones that swim around. These are the kinds of systems which connect *us* to the world. At our sensory and motor peripheries, our systems have narcissistic properties. Third, much of a simple organism's behavioral repertoire can be accounted for without the use of anything other than such narcissistic systems, without anything that looks like an internal representation of objects and properties. If one wants to build a simple organism that swims around and eats anchovies, it is not clear that *any* representational states must be used at all. For many classes of behavior, some of which are quite complex, "making it work" is all that counts. This is how sensory function looks from the point of view of the neurophysiologist.[13]

What lessons about representation can be wrought from these two viewpoints? First, insofar as the neurophysiologist's view conflicts with our firm intuitions about how a sensory system should work, we ought to be suspicious of our intuitions. These are, after all, intuitions about a quintessentially empirical question—'How do sensory systems work'?—a question to which a scientific answer is only now forming. Perhaps we have simply painted sensory systems in our own first-person image. Because our own conscious perceptions are genuinely representational (we have thoughts about the world, its objects, properties, and events), we have expected that both the sensory systems of simple creatures and the parts of our own complex systems will have this property as well, or at least some rough facsimile of it. We have assumed, without warrant, that all sensory systems (or parts thereof) are, in some sense or other, about the world as well.

More importantly, however, the distance between the neurophysiologist's view of sensory systems and our first-person perspective on conscious perception should raise a genuine puzzle (or rather, yet another genuine puzzle) about representation. We ought to wonder why and how we came to represent the world as we do, given the way in which our neural systems anchor us to the world, given how our sensory and motor systems function. The initial assumption of naturalistic theories has been that even simple perceptual systems must have states that have, at least in a limited sense, aboutness. If an organism did not know, at least roughly, what was out there, when, and where, how could it possibly survive? This is why even simple systems must represent the objects and properties of the world. But if the function of sensory systems is not to inform the brain (the ganglia?) about properties of the world per se, then genuinely representational systems are not merely the next bells and whistles added on to an established evolutionary trend. (First there was representation, and then there was *more.*) They are not mere refinements of or embellishments on an ongoing representational strategy. Our ability to represent the external world as containing objects, properties, and

[13] And to many computer scientists. For one defense of this view, see Rodney Brook's "Intelligence without Representations," *Artificial Intelligence*, XLVII (January 1991): 139-60.

events constitutes a distinct—different—capacity of an organism. What exactly is this capacity, and for what reasons did it come about? If an organism can get about, feed itself, and reproduce all without representational mechanisms, what neural/environmental "problem" was answered by the evolution of intentional states? How does this representational capacity work and, more specifically, how do our conscious intentional perceptions seem to form an apparently seamless union with our narcissistic sensory systems?

Let me try to set out this puzzle in more detail. On this view of things, aboutness constitutes something like an ontological "capacity," an ability to impose stability, order, and uniformity upon a conception of the world (and sometimes the world itself) on the basis of stimuli that do not themselves exhibit these properties. Walk around your kitchen and imagine the images that thereby are reflected on the back of your retinae (pretend, that is, that you are Bishop Berkeley making a cup of tea). The image of your refrigerator looms larger and then smaller as you walk toward it, then back away; despite uniform paint, its three sides have different spectral reflectances (given the location of the light source and each side's proximity to other colored surfaces); as you turn your head, the image shape changes dramatically; when you turn around and walk out the kitchen door, the image disappears completely. Add to this the fact that the images on both retinae are somewhat different, for each eye views the world from a slightly different vantage point. And on and on. That none of these changes in the stimuli matters to your perception of the world is a remarkable fact. You see the fridge as having a constant shape, with a uniform surface color, standing in a single place in the world. Despite the differences between this fridge and the hundreds of others you have seen, you regard it as one instance of a type, as a refrigerator; despite the dents and scratches it has gathered over the past years, you know it is the very same refrigerator that you purchased six years ago, the same one that sat in a warehouse several miles away. That you come to glean this stable ontology, of particulars that instantiate types, of particulars that occupy stable places in the world, is an astounding capacity. It requires that you (your brain) find stability despite stimulus change and uniformity, despite real difference, despite the dissimilarities between objects of the same type and the changes in objects over time. It does not matter whether, in the world, these stabilities and uniformities actually exist (or whether, in some cases, we merely impose stability via ontological categories). To conceive of types and tokens, places and objects as existing at all, given our sensory access to the world, is a fantastically difficult task. Call this the *ontological project.*

On the other hand, there is that small task, assigned to sensory systems, of getting the job done, of directing motor behavior. For the most part, this is not an ontological project, a task for which it is an advantage to see the world according to stable categories or as containing re-identifiable places and particulars. For one, because sensory systems encode information symbiotically with motor needs, the similarities and differences, uniformities and discontinuities that the senses "record" need not exist in the world. What matters is whether the system of encoding is effective, whether the encoded similarities and differences are useful.

The vestibular system affords an excellent example here. When the vestibular system, which provides our sense of balance, records the tilt of your head away from the vertical, it "partitions" this information in accordance with the requirements of keeping a human body upright. The physics of falling bodies being what they are, it is much easier to right oneself as the body begins to tilt or lean away from the vertical. So around the vertical point itself, the system makes extremely fine discriminations, thereby allowing the motor system to make the appropriate fine-grained postural adjustments. As the body tilts farther and farther away from the vertical, coarser discriminations (of angle of tilt) are made. So if one were to draw a polar map of human vestibular function, with the vertical at the center, the map would show concentric rings, with denser rings around the center, with wider and wider spaces between the circles, until the forty degree mark. (Why forty degrees? After forty degrees, the physics win. The vestibular system ceases to respond because the motor system cannot right the human body after that point.[14]) From your point of view, that of the conscious subject, very large changes in tilt, once you are already off balance, all count as the same, while the fine differences around the vertical which you do acknowledge do not mark any uniform fact about your position. The vestibular system's representations of tilt, its "concentric rings," mark imaginary divides that are useful for keeping you upright. In general, sensory systems both make and ignore distinctions in nature when it suits the organism's motor needs.

There is another reason, more difficult to explain, why sensory systems might be considered "pre-ontological," or unconcerned with a delineated world. Take the philosopher's favorite reptile, the frog. When the fly-detector cells react, they indicate (in normal circum-

[14] And why, one might ask, are cats so good landing on their feet? Because, through a neck reflex, they manage to keep their heads, and vestibular systems, turned upright thereby receiving vestibular information through almost the entire range of body angles. That, of course, and faster righting reflexes in general.

stances) the presence of a fly, but not the presence of any particular fly (call him "Herbert"). Just a fly, whichever one happens to present itself. For a frog, at least, it does not matter whether this fly is Herbert or Harold; it does not need to keep track of or identify Herbert *as a particular*, whatever that could mean ("You've flown your last flight, Herbert!"). Nor, for the purposes of eating the fly, does it matter that Herbert is in any particular place. The co-ordinates of the moving spot on the retina encode the fly's position relative to the frog and this is exactly the information which the tongue-swiper needs—the direction of swipe for the tongue relative to the frog. Nor, for the purposes of consuming the fly, does the frog need to know where, in the world, Herbert is located ("The last time I saw Herbert, he was sitting in the Savoy...."). All of these are familiar points from the philosophical literature. The same lesson, however, can be applied to our own case. When push comes to shove, so to speak, much of the information needed to make a movement is of the very same sort. Although you see the fly as a particular (and may even call him "Herbert"), reaching out to grab the fly will require information, not about Herbert qua particular fly (here in the lounge of the Savoy), but about his velocity relative to your arm motion, his position relative to your hand, and so on.[15] *That* information, without which you cannot catch the fly, is no different in kind than the information required by the frog. It is extremely precise information about a particular but not *qua* a particular. Call the narcissistic encoding of this type of information the *sensory-motor project.*

It is the gap between the needs of the sensory-motor project and the demands of the ontological project, I want to claim, that calls out for explanation. What were the behavioral/environmental conditions in virtue of which the development of an ontology—and hence things and properties to have thoughts about—came to have survival value to our predecessors?[16] Exactly what kinds of abilities or capacities are required in order to represent a stable ontology, of types and tokens, objects and places? And how exactly does the information

[15] For some interesting research on cortical visual maps that are "arm-centered," see Micheal Granziano, Gregory Yap, and Charles Gross's "Coding of Visual Space by Premotor Neurons," *Science*, CCLXVI (November 1994): 1054-57.

[16] To put this another way, there are certain circumstances when particulars (one's off-spring or mate) impinge upon the lives of simple creatures and for which having detector cells can be just what is needed. Detectors are reliable (indeed, they are often much more reliable than, say, the mechanisms of human facial recognition) and they are usually economical as well, requiring few neural resources (an ant, one must realize, does not have many neurons to spare). But while detectors may signal the presence of a particular, they do not represent them *as* particulars. So one can rephrase the above question as follows: What kinds of informational problems about particulars require, or are better served by, something other than detectors?

provided by our sensory systems co-exist with, form a whole with, the ontology imposed by a representational system? (Given that we *feel* thermal sensations as a function of a receptor population response, how does that fit with our conception of temperature as an objective property of objects?) These are the questions about representational directedness which immediately arise.

To grant that there is this sort of puzzle—even to ask the above questions about directedness without answering them—is to admit, pace Wilfrid Sellars, that in an important sense we do not really know what "aboutness" is. Certainly, at the outset, a vague realism about the directedness of mental/neural events is adopted: representations are "tied" to objects and properties and hence (there being no good reason to suppose otherwise) bear some kind of relation to them. But if we do not know exactly what it means to regard a particular as a particular, to see this thing as being of a certain type, this place as the same place, and so on—hence what kinds of capacities or abilities are involved in having representations that are about those things—then we do not know, in any substantive sense, in what that relationship consists. We only trust that it *is*.

To admit the above, I think, is to take on a different view about what kind of intertheoretic explanations are appropriate to intentional phenomena. On the naturalistic scheme, recall, the idea is to explain the relation of aboutness, in the first instance, by appeal to some other relation postulated by the natural sciences. Exactly how this natural relation figures in the larger theoretical picture varies from theory to theory, given the diversity of the authors' explanatory goals. Sometimes the natural relation is thought to confer a sort of "proto-intentionality" upon specific sensory states, states that are hypothesized to form a "ground" for the genuinely intentional states of folk theory (see, for example, Dretske's *Knowledge and the Flow of Information*); in other theories, the natural relation is simply identified with the "aboutness" relation of, say, a computational state (for example, the Churchland's notion of "calibrational content," developed in Paul Churchland's *Scientific Realism and the Plasticity of Mind*). Once one realizes that the demands of the sensory-motor project are, for the most part, distinct from the demands of the ontological project, however, one realizes that sensory systems need not be veridical reporters as portrayed by the traditional view. Hence sensory states need not bear the expected natural relation to external properties—the natural relation that was to mirror or ground the aboutness of mental or computational states. To ask the above questions is to admit, then, that the directedness of mental events constitutes a distinct set of representational abilities and capacities, which at this point, we can only roughly define. If this is so, then there is no longer any reason to think that a relational property at one level of theoretic explanation (the psychological/computational level) will have a clear mapping onto, or grounding in, any relational property at another (the neurophysiological or biological). The relation of aboutness need not be explained primarily in terms of some other natural relation at all.

Most importantly, questions of the above kind serve to shift the focus of theoretic attention away from static perceptual states. Recall the naturalist's hope that by closely scrutinizing simple sensory events and their relations to external causes, we would gain a toehold on the phenomenon of directedness—by understanding the most straightforward cases of mental directedness, we would have a route into more complex intentional phenomena. Because all of the senses, on the traditional view, are veridical, it will be the static perceptual case qua correlational state that will provide the essential key to aboutness. Once the assumptions of the traditional view are set aside, however, there is no assumption that sensory states will, in general, be about external properties of the world. More importantly, one can no longer assume that for those fully intentional perceptual states which indeed are about the world, it is the static perceptual case—where we sit with our eyes open staring at the object of perception—that will provide us with insight into that relation. Trace out the causal path between the object of perception, the stimulation of the receptors, and whatever neural events that thereafter eventuate and this alone will not *explain*, in the required sense, how genuine representation arises. The explanation of any particular perception can take place only against a background theory of the representational capacities at work. So, it is a theoretic understanding of those capacities qua capacities which will give the explanation of aboutness bite—and these are not capacities which we will understand simply by scrutinizing with care the static perceptual case. (This is not to say that the study of sensory systems is of no utility to the intentionality theorist—on the contrary—but that looking for correlations between sensory states and external properties is not what the study of perception will be all about.)

KATHLEEN AKINS

Simon Fraser University