ASPECTS OF COGNITIVE REPRESENTATION

Fred Dretske
Department of Philosophy
University of Wisconsin, Madison

An organism's cognitive system is a control mechanism whose function it is to initiate, adjust and, if necessary, suppress behavior in the service of need and (if applicable) desire satisfaction. In order to perform this function the control mechanism must have direct and continuing access to intelligence about the circumstances in which activity is to be carried out. Hence, the cognitive system is that part of the executive mechanism whose contributions to control are themselves under the control of (or at least sensitive to) whatever information is (or has been) available about the theater in which operations are to be performed. It is an information-driven control system.

Our ordinary attributions of perception, knowledge and belief reflect this general picture of cognitive processes. To say what someone sees, knows or believes is to identify particular control states by means of their representational properties--in terms of the kind of information they are themselves under the control of. Such, at least, is the representational theory of cognitive processes as we find it (or as I find it) embodied in ordinary descriptions and explanations of animal and human behavior. What I want to do here is to explore this way of looking at cognitive systems in order to see how much weight it can bear. That is to say, is the general idea of a representational control system fertile enough to support the enormously rich and variegated attributions of semantic content characteristic of our ordinary descriptions of what we see, know and believe? Can it, furthermore, provide us with a way of understanding how such attributions figure, as they are ordinarily thought to do, in explanations of behavior?

By a representational system I shall mean any system whose function it is to indicate, by its various states, how things stand with respect to some other object, condition or magnitude. This, obviously, is not to require much of a representational system. A variety of simple devices qualify. A tachometer, the sort of instrument found on the dashboard of many automobiles, is a representational system according to this characterization. Its various states indicate something about the rate at which the engine is rotating. Hence, it represents the angular velocity of the crankshaft. A doorbell, in virtue of indicating the condition of the doorbutton (depressed or not), thereby represents the position of the button. And the firing of a neural cell, by indicating the presence and orientation of a certain energy gradient on the surface of a photoreceptor, represents the whereabouts and orientation of an "edge" in the optical input.

In speaking of a representational system I shall continue to speak of the information the system carries about the quantities, conditions and objects it represents. I intend nothing subversive in this way of speaking--nothing, I hope, that begs the questions which it is my project to explore. For by information I mean nothing more, and certainly nothing less, than what the particular states of a representational system indicate to be so.

Thus the tachometer's registration of "1,000 rpm" indicates, and thereby carries the information that, the engine is running at 1,000 revolutions per minute. The ringing doorbell carries the information that the doorbutton is being depressed (and therefore the information that someone is at the door) because, presumably, this is what it indicates. Representational systems and information processing systems are, on this way of thinking, two sides of the same coin. Information is what representational systems need in order to represent, and representation is what information processing systems do for the things about which they carry information.

In thinking of a cognitive system as an information-driven control mechanism, therefore, we are thinking of it in representational terms. This may not be all we have to say about a cognitive system in order to distinguish it from other control mechanisms (e.g., those associated with the autonomic nervous system), but it will do as a start. The project is to see how large an oak we can tease out of this tiny acorn, how far one can go in understanding perception, knowledge and belief with these semantically meager resources. If they prove too meager, it will at least tell us something about the special character of cognitive processes.

In thinking about a representational system, there are at least two questions one can ask about its representational capacity. And when the representational system is, in addition, a control system, there are at least three questions that should be asked. One can ask, first, what it is, what quantity, property, object, person or condition, the system is representing. A thermometer represents the temperature, a fuel gauge the amount of gasoline in the tank, a photograph the objects (persons, building, foliage) that the picture was taken of. Secondly, one can ask how what is represented is represented. What does the representation say about what it represents? That it is 95 (in the case of the thermometer)? That the gas tank is almost empty (fuel gauge)? That your niece has let her hair grow long (photograph)? The first question is a question about the reference or denotation of the representation. The second question is about the content of the representation. Topic and comment.

There are, in other words, pictures of black horses and what Nelson Goodman (Languages of Art, Hackett; Indianapolis, 1976, p. 29) has called black-horse pictures. Unless the picture of a black horse is a black-horse picture, it will not represent the black horse as a black horse. Imagine, for example, a picture of a black horse in which the horse is photographed at a great distance in bad light with the camera slightly out of focus. The horse appears as a blurry spot in the distance. This is a picture of a black horse but not what Goodman calls a black-horse picture. When invited to see pictures of your friend's black horse, you expect to see, not only pictures of a black horse, but black-horse pictures—pictures in which the denotation of the picture is identifiably a black horse, pictures in which the black horse is represented as a black horse.

Similarly, the wolf's internal representation of the sick caribou may or may not be a sick-fleeing-caribou representation. But it certainly is a representation of a sick fleeing caribou. How it represents the animal is, to some degree, a matter of speculation, but unless it has some means of representing defenseless caribou, a way of commenting on these creatures which is, for practical purposes, extensionally equivalent to being a

defenseless caribou, its relentless and unerring pursuit of these particular animals is inexplicable. It would be like trying to explain the behavior of a thermostat in controlling the furnace if it had no means of representing the room temperature as above or below the desired setting. There has to be something in there that "tells" the thermostat what it needs to know for it to carry out its function. The same is true of the wolf.

I rehearse these familiar facts about representations only to emphasize that in thinking about a cognitive system as a representational control mechanism, the same questions can be asked about our cognitive states: what do they represent and how do they represent it? What are they getting information about and what information are they getting? But I said that there were three questions that can be asked about a representational system when it functions as a control mechanism. We can ask, not only about its topic (what is represented) and the comments it makes about that topic (the way it is represented), but about which of these comments, if any, has a control function. Which elements of the representation play a causal role in the determination of behavior? A representation may be ever so rich in the comments it makes about what it represents, but if none of this information is, or can be, used to control and direct movements, it is causally inert, hence, functionally irrelevant. Therefore the representation, qua representation, plays no role in the system's cognitive economy. A black-horse picture of a black-horse can be used to paper over a hole in my wall. In this case, the representationally significant aspects of the picture are irrelevant to the way it is functioning (concealing the hole). I could as well have used a picture of a white cow or a picture of nothing at all. And if our internal representations are to qualify as cognitive, they must (potentially at least) make some contribution to the way that mechanism functions in controlling and directing behavior. They must do so, furthermore, by means of their representationally significant properties. Only by so doing will the classification of cognitive states in terms of their content figure in the explanation of the behavior they produce.

Our ordinary attributions of sensory and cognitive states reflect the kind of distinctions just discussed. And this, to my mind at least, supports the idea that our familiar, folk psychological picture of cognition is a picture of a kind of representational control mechanism. We say, for example, that Clyde can see a black horse (in the distance) without (for various reasons having to do either with the great distance, the camouflage, the lighting, or the fact that he doesn't have his glasses on) its looking like a black horse, without its presenting a black-horse appearance. In describing what Clyde sees, we are describing what his sensory representation is a representation of. We are ignoring the kind of comment his perceptual system is making about that topic in order to specify the topic itself about which a comment is being made.

Thinking of the cognitive system as a representational mechanism, a mechanism no more complex than a simple gauge, gives us, therefore, the resources for understanding, not only the propositional attitudes, those attitudes (like knowledge and belief) that take (in their verbal expression) sentential clauses as complement of the verb, but also those attitudes (like seeing and hearing) that take concrete nominals as objects of the verb. S sees a bush, mistakes it for an animal crouching beside the path, and flees

in panic. The description of what S sees is an expression of what his
internal representation is a representation of--in this case a bush. What
he sees it as, what he takes it to be, or (under optimal conditions) what he
can see that it is, is an expression of the way he represents the bush. In
this case he has a crouching-animal representation of a bush--the analogue,
I submit, of a white-cow picture of a black horse.

There is then, as I see it, no real question about the validity of our
ordinary descriptive apparatus for assigning perceptual-cognitive states to
organisms. For in describing a creature as seeing, knowing or believing
something, we may be (and, I would urge, certainly are) doing no more than
what we are already doing with such simple representational devices as
gauges, instruments, and detectors. We are saying what and how things are
being represented. The only real question about the representational model
(aside from the causal efficacy of these representations--a point I will get
to later) is whether our ordinary descriptions of what a creature sees,
knows and believes are semantically too rich to be supported by the actual
physical representational resources of the organism. Are our ordinary
assignments of reference and content to an organism's internal
representations compatible with--and, if so, are they realizable in, the
actual neural machinery available for generating these representations?
There may be no real dispute in assigning the reference my gas tank and the
content half full to my gauge's representations; for the actual physical
construction of the device, and the laws governing its operation, clearly
reveal that this is what the device indicates and what it indicates it
about. But in saying that the wolf saw a sick caribou on the edge of the
herd, recognized it as sick and, therefore, as easy prey and, because of
this, pursued the animal while ignoring the thousands of healthy ones
nearby, are we assigning reference and content that exceed the
representational resources of the animal we are describing? If not, what
about our descriptions of Jimmy as seeing his uncle and recognizing him as
the man who promised to fix his bicycle?

INVITED ADDRESS: THE MEANING OF VISION-RELATED TERMS TO A BLIND CHILD

Barbara Landau, Columbia University

Lila Gleitman, University of Pennsylvania