

CHAPTER SIXTEEN: DETACHING GOAL STATE REPRESENTATIONS

The animal that constructs intentional representations of temporal contingencies must have a use for these representations. That follows from the description of intentional signs in Chapter Six. The obvious use of these representations is preparation for the future. The animal uses representations of the future as representations of current affordances, as guides to current action, enabling it to begin now to position itself to make use of upcoming events and upcoming situations, or to take steps now to avoid them. Depending, perhaps, on its current needs, or its current state of potentiation, representations of future events serve for it as pushmi sides of P-P signs which direct activities such as ducking an approaching ball or blow, salivating to be ready for swallowing, preparing to execute a sharp turn coming up just ahead when running, and so forth. That is, perception of what is distal in time operates exactly as does perception of what is distal in space. Just as the animal is guided here by a perception of what is there, the animal is guided now by a perception of what will be later.

I have spelled this out rather carefully for the following reason. In the case of an animal that predicts and represents the result of its own purposive action, it is easy to slip into thinking that in representing the future it is guided, not just by the future it represents, but toward the future it represents. That is, it is easy to confuse anticipating future events for which it itself is purposively responsible with using representations of those future events as guiding goals. Let me first illustrate the difference between these two, and then attempt to analyze the relation between having a purpose that one also expects the fulfillment of, and having that fulfillment as a goal guiding one's behavior.

I once had a cat, named Sam, who learned (as cats will) to push open the screen door to let himself out. But the door had a very strong spring on it, and would close just in time to catch the end of Sam's tail. Sam tried to avoid this by running faster, but the faster he ran, the more he instinctively lowered his tail to streamline himself, and the further behind him the end of his tail would be. The irony was, of course, that if he had just walked slowly through the door with his tail straight up, it would never have been pinched at all. Now the lowering of his tail was instinctive and biologically purposeful. And he predicted the effect of this purposeful motion and tried to avoid it. But representing the possible future effects of his activity in lowering his tail so as to prepare for or ward off these effects clearly was not the same as controlling the purposive activity of his tail itself in accordance with a representation of its possible effects.

Similarly, for an animal whose behavior was controlled entirely by perceptions of affordances, every facet of activity would be controlled by perceptions of what the brute facts are or are going to be. That some of the brute facts about what is going to be result from the animal's own dispositions to follow certain affordances would not change this matter. Adjusting to the consequences of the manner in which it is already disposed to act is something such an animal might learn, over time, to do. And learning not to do, in future, what it is currently disposed to do is also something that, over time, it might learn. But the goals of conditioning mechanisms are not projected ahead. They govern changes in behavior only after the fact. For such an animal, then, perceptions of the future effects of its activities could not serve as guides to control those same activities currently, but could at best control concurrent compensating activities. That an animal anticipates the outcome of what it is doing does not imply that this anticipation is what

controls or guides what it is doing. Something more must be added if we are to understand what it is for an animal to project a goal and adjust its activity, on the spot, in order to meet that goal, rather than merely predicting the outcome of its previously learned dispositions to follow affordances.

But if it is my purpose to do A, and I anticipate that what I am actually doing will not lead to doing A, isn't it obvious that this will cause me to alter what I am doing? First, recall the protective eye blink reflex (Chapter One). When the eye doctor is trying to put drops in your eye, you anticipate your purposive blink; it does not follow that you can control it. In this case the reason, or at least a reason, seems clear. The purpose of the eye blink reflex is not a represented purpose. It is not on the same level as your conscious represented purpose of keeping the eye open, and so will not be canceled out by this represented purpose. Similarly, recall the greylag goose that doesn't notice that its egg-retrieval routine has failed in the middle. Again, a reasonable assumption is that this happens because the purpose of the goose's activity is not a represented purpose. It is the purpose of an automatic response. Can we also explain, in this manner, why the hamsters didn't notice that the purpose of their food gathering and storing activities was not being accomplished as they stumbled over one another, retrieving the same crackers again and again from one another's corners (Chapter Thirteen)?

That the hamsters didn't know what they were doing might be harder to imagine, because their activity was clearly being guided by various perceived affordances in the environment. They had to perceive the crackers as things to be taken home and stored, and they had to perceive the paths to their corners as paths to be followed, and so forth. However, the crucial question is whether perceiving an affordance involves representing the end state to which the affordance leads. I have spoken of the perception of an affordance --the harboring of a P-P representation-- as a perception at once of what the case is and of what to do about it. But, first, a representation that directly guides one's doing, one's motions, is not the same as a representation of where one will end up as a result of that motion. The movement of my hand and arm may be directly guided by the turns of the bannister on the way down the old fashioned staircase without involving any representation of the place the bannister ends. Similarly, a path can guide my walking without my knowing where it, hence I, will end up. Second, notice that if the representation of an affordance does somehow represent the goal state at the end of the path it offers, it surely can't represent it in the same representational system in which it represents the current situation. If the very same sign says both what the current state is and what future state is directed to be, it must be saying these things in different languages.

Now there is a sense in which the perception of an affordance is always a directive representation of a goal state. But usually it is an extremely inarticulate representation, directing no more than when the goal state is to be achieved, namely, at some more or less definite interval after the time of the representation. This is what it directs because this is what it is its purpose to achieve and this is what it maps onto, time mapping onto time (Chapters Four and Thirteen). Other significant variables of the directive aspect of the P-P perception direct motions or paths, not end states. Consider, for example, how a guided missile tracks a target. The information that guides it concerns the angle between its direction of motion and the target. It responds by continuing to correct its direction of motion until it matches the direction of the target, so if all goes well it ends up where the target is. But it employs no representation of itself ending up at the

target. To suppose that it does would be to confuse a representation of the target's relation to its direction of motion with a representation of hitting the target. It would be to confuse a representation of an enabling relation to the target with a representation of what guidance by that relation will enable.

Similarly, consider how your motion is guided when reaching to pick up an object. What directly guides your motion is a perception of the current relation, the distance and direction, of the object relative to your hand (and of your hand to the rest of your body, for leverage and balance). Your motion is a complicated function of that distance and direction. The function is such that the motion of the hand (and compensating motions of the body) will result in your hand arriving where the object is. Nothing has been said yet about a representation of your hand arriving or having arrived where the object is. That would be another kind of representation. Following an affordance that results in picking up the object is being guided by your current relation to the object so that your movement is a function of that relation. The purpose of following that affordance is to achieve a state in which your hand is on the object. But merely in perceiving and following the affordance, that end state is not represented, or not represented articulately. A "goal," in the sense that a goal may be perceived or represented in thought in the perception of an affordance, is just a place, such as the inside of a net or a basket. The goal may be represented, but the goal state corresponding to the affordance, say, the ball's being in the basket, is not part of what mere perception of the affordance represents. It is not part of what the P-P representation represents.

This brings into sharper focus a central question I have proposed to consider in Part IV. How does it happen that the descriptive and the directive sides of primitive P-P representations eventually separate and become independent? It now becomes clear that the answer cannot be entirely straightforward. It cannot be just that the P-P representation breaks in two, the two aspects or sides coming apart so that each stands alone. First, a projected goal, one that guides the animal in planning and in knowing when it has reached its goal, has to be a representation of a goal state. And a projected goal state that will allow the animal to know whether or when that goal has been reached needs to be represented in the same language as the animal's descriptive representations. It needs to be represented in pushmi language. The goal state that is projected during an explicitly purposeful act is, in general, arrival at a state of affairs that will afford something further. It will afford arrival at the perception of another affordance, in the final instance, arrival at perception of a B-affordance (Chapter Thirteen). But for arrival of an affordance to be recognized as such, the language in which this arrival is explicitly projected as a goal needs to match the language in which it is later to be perceived as a fact. It has to match the descriptive side of the P-P representation that it anticipates.

A point that may superficially appear to be parallel is recognized by many researchers when discussing the roles either of efference copy or of reafference in guiding motor activity. For example, Jeannerod (1997) says, "In order to be useful for the comparison process, the reafferent signals must be compatible with the efferent ones. In other words, the two must be coded in the same 'language' for being mutually understandable and for the matching process to be possible" (1997, p. 178). However, this seems to presuppose that what is articulately represented in the

language in which motor activity is directed is the end state to be achieved rather than a path to be followed, which would surely lead to a problem of individuating or counting these end states. Jeannerod says, "Goal-directed behavior implies that the action should continue until the goal has been satisfied. The description of the motor representation must account for this property, that is, it must involve, not only mechanisms for steering and directing the action, but also mechanisms for monitoring and eventually correcting its course, and for checking its completion" (p. 173-4). This image collapses the representation of the changing relation of the organism to its target that dynamically guides its motor activity with a representation of the end product to be achieved. It confuses dynamic perception or prediction of a changing enabling relation to a target with representation of a goal state. Representations of a variety of goal states and goal sub-states may be divided or individuated during an activity, but surely representation of changing relations to a target during short movements is dynamic. Indeed, with practice, movements directly under the control of dynamic input are extended to have longer and longer duration. The experienced tennis player sees the ball coming and hits it to a targeted place. Typically, no representations of any sub-goal states intervene from start to finish. That, I take it, is a part of the Gibsonian vision that we want to preserve.

That there is a level on which representations of intended goal states and perceptual representations speak the same language has recently been argued for from experimental evidence by Hommel et al (2001). Their claim is that there is a level of common coding for perception and for plans of action, and that this coding is of distal events. In the "Authors' Response" section of the paper, they are explicit that by "action planning" they do not mean motor coding, and that the level of perception they have in mind is not the sort produced by the dorsal system. Thus their claim, as I understand it, is compatible with the general point I have been trying to make here. An animal that not only looks ahead, predicting the results of its actions, but plans ahead, adjusting the predicted results of its actions to represented goal states, and that knows when and whether it has reached those goal states, must code its action plans in the same representational system in which it codes perception of the results of its actions. It has to represent goal states in such a way that it can recognize when they have been achieved. Moreover, since this kind of planning ahead is a stage by stage sort of process, it has to be undergirded by more dynamical action processes, presumably by direct guidance in accordance with perceptions of affordances, governed most directly by representations of enabling relations supplied perhaps mainly by the dorsal system. This also fits, for example, with the suggestion from Chapter Fourteen that the achievement of objective representation may allow the animal to analyze its various activities into distinct achievement stages, understanding its activities as composed of series of transformations of one objective situation into another objective situation, hence as a series of completion stages in an objective process.

Compare the way in which language forms having the same satisfaction conditions but different linguistic functions or purposes may employ what are basically the same semantic mapping functions. Thus "Close the door," "The door will be closed" and "The door is closed" may all have the same satisfaction conditions, though the latter two have to be said at different points in time for this to be so. And they all express their satisfaction conditions in the same basic representational system. There are systematic ways to change an imperative mood sentence

into an indicative mood sentence with the same satisfaction conditions, and systematic ways to change a future tense sentence said at one time into a present tense sentence said at a later time while retaining the same truth conditions. Compare also the Fregean idea that one can have different "propositional attitudes" toward the same proposition (though distinctions among the tenses have not generally been understood this way). But I don't want to suggest too close an analogy with sentences or with Fregean propositional attitudes either. Unlike sentences, mental representations with the same satisfaction conditions toward which different mental attitudes were taken need not, presumably, display different syntax but merely need to have different uses or functions. And as for propositional attitudes, we have been dealing with perceptual representations and with representations continuous with or joined to perceptual representations through recall, not with conceptual thought. No position on the relation of perceptual representation to conceptual representation, or even on whether there is such a distinction, or whether one fades into the other and so forth, is intended.

The idea to be explored is that beyond pushmi-pullyu representation, a common system of mental representation may develop in which projected goal states, objectively represented future states, and objectively represented present states can all be expressed. ("Objectively represented" means that the representing organism is not implicitly represented, hence if represented at all is explicitly represented, represented as an object. See Chapter Thirteen.) We suppose that the pushmi side of an inner P-P sign of an objectively represented affair is already expressed in the required format, for example, the pushmi side of your objective representation that coffee is now in the cup for drinking. Recall that the pushmi sides of some inner P-P signs represent coming events rather than present ones. Sometimes these coming events are not represented objectively, for example, it is doubtful that the nutcracker who responds to signs of winter coming by caching nuts represents winter coming as an objective future occurrence. The nutcracker's perception is a representation of winter coming in that it will serve its own or proper function in a Normal way only if winter is indeed coming, but it is unlikely to be an articulate representation of that, for example, it is unlikely that any transformation of that representation says that summer or spring is coming or, as is more relevant here, that winter is here. We are now postulating a new kind of representational capability, such that coming events and present events are coded in a common representational system, the difference between a representation of something future and a representation of the same thing as present lying in the use rather than the coding of these representations. What this makes possible, in the first instance, is that the representing organism may learn to represent certain coming events, on the basis of present local signs, by a trial and error or corrective process in which its predictions are remembered and compared with later outcomes. It can do some learning about what in the present is a sign of what in the future without having to use the criterion that is the success or failure of immediate practical activities that depend on these predictions.

Besides present and future states and events, we postulate that projected goal states are also represented in this common representational system. Thus the organism is able to know when or whether it has reached these states. This allows cessation at appropriate times of activities designed to lead toward projected goals. It thus prevents the sphex effect (Chapter Thirteen). Does it also allow learning what kinds of motor activities do and do not accomplish

certain goals?

The assumption of classical ideomotor theories of perception, as Hommel et al (2001) put the matter, is that

[t]hough the learning refers to linkages between movements and their effects, the result of this learning needs to be organized in a way that allows to use the linkages the other way round, that is, go from intended effects to movements suited to realize them. ... If one takes for granted that the links between [representations of] movements and [representations of] effects can be used either way, a simple conceptual framework for the functional logic of voluntary action offers itself. This framework suggests that actions may be triggered and controlled by goal representations --that is, representations of events the system "knows" (on the basis of previous learning) to be produced by particular movements. (2.2.2)

Hommel et al's "goal representations" seem to be what I have been calling "goal state representations." But we are assuming here that motor activities are always directly guided by perceptions of affordances, where complete perceptions of affordances always include perception of enabling relations (Chapter Fourteen). Learning what kinds of motor activities do and don't accomplish certain goal states is thus the same as learning which affordances, which ways of being guided by perception, lead to which objective results. Actions will not be directly triggered by goal state representations unless suitable affordances, including enabling relations, are perceived. On this assumption, what will be learned by the animal is not what events will be produced by particular movements, but what affordances, if followed or acted on, will lead to what objective results. That, it seems likely, is part of what young mammals and birds are learning when they are playing. In the case of human infants, experimentation begins in the very first few hours of life as the infant moves its body parts in response to both internal and external stimulations, at first quite randomly, later with more and more purposive intent, to see and hear and feel what happens.

Hommel et al emphasize that the level at which "common coding" occurs concerns events at a distal level. For example, this coding may represent the effect that is a red light flashing on to your left rather than representing the hand motion (pushing a key down or lowering your finger) that causes the red light to flash. Their studies also suggest that when attention is shifted slightly inward to concentrate on a more proximal effect, say, on the key pressing rather than the light flashing, then the common coding level shifts also to the more proximal level. These conclusions are reached mostly from the results of reaction time experiments that suggest coding interferences and coding facilitations among actions of various kinds under different instructions. These interference and facilitation relations may shift within what is exactly the same experiment looked at in terms of actual outcomes produced by the experimental subjects, but in which the subjects are given different instructions designed to encourage them to focus their interest on more distal or less distal aspects of the outcome. For example, they are asked to be sure to press the left or the right key or, alternatively, to quickly make the light on the left or on the right come on. It appears the common coding may shift with attention shifts from representing the position of the hand or finger to representing the position of the intended light (Hommel 1993).

This suggests that motor-perceptual learning, learning how to be guided by perceptual

inputs so as to produce predicted outputs, may take place at many different levels of distality, depending on the interests of the learner. The infant learns how to be guided by perception so that an object is transferred (by reaching and grasping) from any of various positions into its own hand, or from one hand into the other, or so that an object across the room is transferred (by the infant's crawling) to being within reach. It learns how to perform simple manipulations such as turning things over, twisting, throwing, and so forth, such that the outcome is as anticipated. The tennis player learns how to be guided by perception so that the approaching ball is transferred to a designated position in the opposite court. Indeed, there is evidence that motor-perceptual learning that concerns distal events of this kind is more primitive than motor-perceptual learning that concerns merely motions of the organism's own body. As mentioned in Chapter Fourteen, very few animals can even learn to imitate bodily motions, and perhaps none but humans¹ do so naturally.

The animal that has learned its lessons well concerning the objective results of following certain kinds of affordances is in a position to predict the fulfillment of its immediately projected goal states in situations perceived as enabling. This clarifies the status of explicit intentions which, on the one hand, are projections of goal states, but on the other, are beliefs about the future. To firmly intend to do a thing consists, in part, in believing that one will do it. Otherwise one intends only to try. A confident intention, somewhat like a pushmi-pullyu representation, tells both what to do and what will have been done, so that further plans that depend on that settled future can now be made. Perhaps we could call a confident intention a "pullyu-pushmi" representation, though the symmetry is not perfect. Unlike the pushmi-pullyu representations, the two faces of a confident intention are both written in the same code. The very same aspect of the representation has two functions, one to produce action, the other to represent a future state of affairs.

A fascinating piece of evidence that the representation of goal states is not part of the mere representation of affordances in the case of humans is the difference between patients suffering severe damage to the prefrontal lobes, and those suffering with an "anarchic hand," caused by damage to the supplementary motor cortex. In both cases it appears that the behaviors of the patient or of the patient's hand simply follows the most obvious current affordances exhibited in the environment, whether appropriate or not. Thus the patient with severe prefrontal damage will pour water into a glass and drink it as many times as water is presented in a pitcher, regardless of whether she is thirsty. The hand of a patient with an anarchic hand may button and unbutton available buttons, grasp a pencil and scribble with it, or grasp a doorknob and turn it in entirely inappropriate situations. A difference between these two kinds of patients, however, is that the patient with prefrontal damage is in no way disturbed by her inappropriate behavior, while the patient with an anarchic hand is very disturbed by it, in some cases even tying the anarchic hand down to prevent its wanderings. This difference is interpreted by Frith et al (2000) to be the difference between patients that have intentions that conflict with the affordances they are following and those that simply don't have any intentions, intentions being a product of the prefrontal lobes. As mentioned in Chapter One, automatic behaviors are accepted as one's own so long as one doesn't disagree with them.

FOOTNOTES

-
1. Actually, it looks as though Dolphins may do this.
See (Herman forthcoming).