Ruth Garrett Millikan 12/12/2002

University of Connecticut

On Reading Signs; Some Differences between Us and The Others

On Reading Signs

Ruth Garrett Millikan

Philosophy Department

University of Connecticut

Storrs CT 06269-2054

860-429-1284

ruth.millikan@uconn.edu

Number of words: 7713

On Reading Signs; Some Differences between Us and The Others

If there are certain kinds of signs that an animal cannot learn to interpret, that might be for any of a number of reasons. It might be, first, because the animal cannot discriminate the signs from one another. For example, although human babies learn to discriminate human speech sounds according to the phonological structures of their native languages very easily, it may be that few if any other animals are capable of fully grasping the phonological structures of human languages. If an animal cannot learn to interpret certain signs it might be, second, because the decoding is too difficult for it. It could be, for example, that some animals are incapable of decoding signs that exhibit syntactic embedding, or signs that are spread out over time as opposed to over space. Problems of these various kinds might be solved by using another sign system, say, gestures rather than noises, or visual icons laid out in spatial order, or by separating out embedded propositions and presenting each separately. But a more interesting reason that an animal might be incapable of understanding a sign would be that it lacked mental representations of the necessary kind. It might be incapable of representing mentally what the sign conveys. When discussing what signs animals can understand or might learn to understand, one question it may be important to have in mind concerns what kinds of <u>mental</u> representations these animals are likely to possess. To this end, a fairly explicit theory of mental representation, and of its various types, would be needed. In this essay I am going, very quickly, to sketch a general theory of mental representation and of the most basic varieties of mental representation. I will suggest some ways in which the most sophisticated kinds of mental representations that humans use seem to differ from those used by other animals.

Mental representations are a species of what I will call "intentional signs." Intentional signs must be distinguished, first, from "natural signs." "Natural signs," in general philosophical usage and in the usage of pragmatics, are signs that are not designed to be used as signs, hence are not conventional (e.g., Augustine 1986, c. 427; Ockham 1495, p. 19) and not voluntary (e.g., Ockham 1495; Kant 1798). Because they are not designed for use as signs, it makes no sense to attribute truth or falsehood to natural signs. Smoke means fire only when it has actually been caused by fire. Black clouds mean rain only when they actually produce rain. Red spots mean measles only when caused by measles.

Natural signs themselves cannot be deceitful or wrong, though it is of course possible for an interpreter to make mistakes when trying to read them. In my usage, natural signs contrast with "intentional signs," which, following Franz Brentano's technical usage of the term "intentionality," are signs that can be false or that may sometimes signify nothing real. By intentional signs I mean those that have been "designed," in accordance with human or animal purposes, or by learning mechanisms, _or by natural selection_, to be interpreted according to predetermined (semantical) rules to which targeted interpreters are cooperatively adjusted. Thus it is possible for intentional signs to be false or misleading. To remind us that this usage of "intentional" is technical, and to be sure that the intentional, in this sense, does not become mixed in our minds with the very different notion expressed by the philosopher's terms "intension," "intensional," and "intensionality," I will sometimes capitalize the "T," thus: "intenTion," "intenTional," "intenTionality." One kind of intenTional sign is a mental representation, as will become clearer below.

In my terminology, "mental representation," does not imply consciousness. I am not going to talk about what is before or within an animal's conscious mind. Mental representations have to do with the mechanics of behavior control, how this control is accomplished, presumably, neurologically. A place to start is with Gallistel's usage of the term "representation" which, as he says, is derived from the mathematical sense: "The brain is said to represent an aspect of the environment when there is a functioning isomorphism between some aspect of the environment and a brain process that adapts the animal's behavior to it," and later, "The exploitation of the correspondence to solve problems in the one domain using operations belonging to the other establishes a functional isomorphism: an isomorphism in which the capacity of one system to represent another is put to use" (1990 p. 15-16). Gallistel seems to have in mind the classic 20th Century view that putting a representation to use involves calculation, but this image is unnecessarily restrictive. Think, instead, of tracing a line with your finger, the visual representation of the line guiding the motion of your hand to conform to the contours of the line. This seems, anyway, to involve something more like translation (in the physicists's sense) than calculation.[1]

Also, Gallistel probably has in mind a more restricted notion of a functioning isomorphism than I intend. Just as, strictly speaking, <u>nothing</u> is a quantity and <u>zero</u> is a number, strictly speaking, a sign system that maps times onto themselves and/or maps places onto themselves, as in simple signaling systems, exploits isomorphisms. The isomorphisms are what allow these simple systems to exhibit productivity. For example, a warning cry to conspecifics that tells when a predator has been sighted is a member of a potentially infinite set of such signals, each telling of a predator at a different time and/or in a different place. The set of possible signals is isomorphic to its corresponding set of possible signifieds. It is not, of course, that any intrinsic property of the cry is isomorphic to any intrinsic property of the predator. Similarly, no intrinsic property of the dot on the map that indicates the village of Storrs is isomorphic to any intrinsic property of Storrs. More accurately, any such isomorphisms are inoperative, nonfunctioning, within this system of representation. Every isomorphism other than the isomorphism that maps a domain onto itself by the identity function involves some arbitrary correspondences.

Gallistel is clear that what makes a brain state into a representation of some aspect of the environment is not just an isomorphism but an isomorphism that is used to adapt behaviors to this aspect. A mental representation is of <u>whatever it is designed to be used as a representation of</u>. This is the same as to say that mental representations are intenTional signs or intenTional representations. Perceptual or cognitive systems that produce intentional representations have been selected for producing representations to be used by targeted interpreters. If the representations they produced were never used as representations, these mechanisms could not have been selected for producing representations according to rules to which targeted interpreters are adjusted. They might be very efficient at producing natural signs, but natural signs are not intenTional representations.

Gallistel's description covers only representations of what is the case -- that is, only "indicative" representations. It does not cover representations of what is to be done -- "imperative representations." For example, it does not cover explicit intentions or goal representations. Paraphrasing Gallistel, I will say that the brain represents something to be done by the organism when there is a specific kind of isomorphism between a brain process and some aspect of the environment (or of the organism-environment relation) that this process functions to produce. The use or purpose of the brain process is to guide behavior so as to produce what it represents. Again, notice that being a mental representation depends on there being uses for it. Brain-environment correlations and covariations that are mere side effects of proper functioning --and there undoubtedly are many-- do not count.

Mental representations, then, can be used either to reflect states of affairs or to produce them. That representations can face either of these directions is not news. Classic statements are in (Anscomb 1957) and (Searle 1983.) What has not been generally recognized is that many representations face both ways at once. The principle is easiest to grasp in the case of simple external representations used for communication between nonhuman conspecifics. Does the dance of the honey bee tell where the nectar is or does it tell worker bees where to go? Clearly it does both. The genes for producing and responding to these dances have been selected because they result in dances that map nectar locations and also because they result in worker bees being guided to those locations. Similarly, alarm calls of the various species do not just represent present danger but are signs directing conspecifics to run or to take cover. If beavers did not dive in response to the danger splashes of their conspecifics, the disposition to splash when sensing danger surely would not have been selected for. These calls and signals are intenTional signs or representations that are at once descriptive and directive.

What then occurs in the head of a bee who understands a fellow bee's dance? Does the bee come to believe there is nectar at location L, desire to collect nectar, know that to collect nectar at L requires going to L, hence desire to go to L, hence, no other desires being stronger at the moment, decide to go to L, and proceed accordingly? Surely not. To posit anything more complicated than, as it were, a literal translation of the dance into bee mentalese is surely superfluous. The comprehending bee merely acquires an inner representation that is at the same time a picture, as it were, of the location of nectar (relative to its hive) and that guides the bee's direction of flight. The very same representation tells in one breath both what is the case and what to do about it. I call representations having this sort of double aspect "pushmi-pullyu" representations (or "PPRs") after Hugh Lofting's charming two-headed Janus-faced creature by that name. (For more details see Millikan 1996.)

J. J. and E. J. Gibson claimed that the direct objects of perception are affordances (Gibson 1966; Gibson 1977). That is, what an animal directly perceives is places to climb up on, things to sit on, places to hide, things to eat or to run from and so forth. One way to understand this is that the natural signs in the ambient energy read via perception are translated directly into mental representations that face two directions at once. They tell what is located in what regions nearby, and at the same time guide appropriate responses to this information. Contemporary Gibsonians postulate "perception-action" cycles whereby structures of ambient energy impinging on the organism and carrying information about the distal environment, added to information about the current configuration of the organism's body, are directly translated into structured action that takes account of both these factors, directly producing behaviors that will be productive given these factors. But by "directly" they don't, of course, mean without mediation by the nervous system. So this is tantamount to postulating basic perceptual representations as being PPRs.

It seems clear that many primitive animal behaviors, even our own most primitive behaviors, are controlled in this way. This becomes transparently clear if we remain strict in our mathematical reading of "isomorphism" in the definition of representation, recognizing time and place as significant variables in representations. Examples are everywhere. The neural signal that triggers your protective eye blink reflex is technically a PPR. It represents that something is approaching the eye too closely right here right now and gives the instruction to close the eye right here right now. It does this by mapping time and place of the approach of the object onto time and place of the neural signal, and in turn onto time and place of the blink. Similarly, the neural signals that mediate between those environmental signals that are "behavior releasers" and the behaviors called "fixed action patterns" (Lorenz and Tinbergen 1939; Tinbergen 1951; McFarland 1981, p. 1990 ff, Gould 1982) that are thereby released in many animals are PPRs. Very simple internal mechanisms that control tropistic behaviors in primitive animals employ PPRs. And if the Gibsonians are at least partly right, many more flexible behaviors such as grasping, chasing, climbing, and so forth, may fit this pattern as well. One possibility is that the simplest animals, at the level of insects, for example, may be governed almost entirely by a set of perception-action cycles arranged in a hierarchy that determines which shall take precedence over which, depending on need, or when more than one currently relevant affordance is perceived. Some animals may be pure pushmi-pullyu animals.

Notice that to serve as an unmediated guide to immediate action, the indicative face of a PPR has to represent the relation of the affording situation or object to the perceiving animal. An animal's action has to be initiated from the animal's own location. So in order to act, the animal has to take account of how the things to be acted on are related to itself, not just how they are related to one another. In the simplest cases, the relevant relation may consist, merely, in the affording situation occurring in roughly the same location and at the same time as the animal's perception and consequent action. More typically, it will include a more specific relation to an affording object, such as a spatial relation, or a size relative to the animal's size, or a weight relative to the animal's weight or strength, and so forth. That the indicative faces of PPRs have to show relevant relations of the affording situation or object to the animal does not make what the PPRs represent in any way "subjective," however. PPRs must give objective information about perceiver-world relations. That the PPR must represent a relation or relations between the affording situation or object and the animal itself does not imply, however, that the PPR expresses a self concept or contains an independent element or aspect that refers to the animal itself. Reference to the acting animal itself is not an articulated part of a PPR. To see this, we must look again at what is involved in the functioning or "significant" isomorphism between a representation and what it represents.

A representation is, as such, a member of a representational <u>system</u> defined by an isomorphism between the domain of the signs and the domain of the signifieds. There will be certain significant mathematical transformations of any sign in the system that will yield other signs in the system, these transformations corresponding in a regular way to transformations of the states of affairs that would be signified. That is what defines the functioning or significant isomorphism and makes a sign system productive. This isomorphism will be defined by certain entirely definite relations among the signs in the system that correspond to definite relations among the correlative signifieds. What the signs say explicitly or <u>articulately</u> will be only what these relations show. For example, the commonest kind of bee dance contrasts significantly with other bee dances only along three dimensions. (See, for example, Gould and Gould 1988). One dimension shows direction of nectar location relative to the hive and the sun. Another dimension shows the rough distance of the nectar from the hive. A third dimension says <u>when</u> this is so, namely, at roughly the same time as the dance. There is no way to transform a bee dance so that it talks instead about peanut butter rather than nectar, or about the moon rather than the sun, or about the tall oak tree rather than the hive, or about something that was the case last week. Exactly similarly, there no way to transform the dance so that it tells not that the watching worker bees should fly off in a certain direction but that just Susy bee should or Sally bee should or, of course, that the wasps should. Nor can we suppose that the mental representational system into which the bee translates a bee dance would allow it to think, alternatively, about the relation of nectar to the tall oak tree or of peanut butter to the moon last week, or that Susy bee, rather than she herself should fly off in a certain direction. What use would a worker bee have for any such representations? Or consider the beaver splash. It's articulation is even more restricted than that of the bee dance, having only two dimensions of contrast: time, and place. Nor can I see much motivation for supposing that the beaver understands the splash by translating it into a mental representation in a system allowing also representations of danger next week or of peace and quiet last week, or of what bears or other beavers should now do. For the same reason, the PPR that tells the relation of the perceiving animal to the affording object and directs the animal's action toward that object does not need to represent the animal itself explicitly. There need not be transformations of it that would represent, instead, the relation of objects other than the perceiving animal to the affording object.

Similarly, in the simplest cases, there is no reason to suppose that the affording situation or object represented by a PPR is explicitly represented. Consider any behavior triggered by an environmental releaser, for example, the feeding behavior of the song bird triggered by the sight of the red inside of the open beak of its young. This behavior is mediated by a mental PPR whose indicative content is that *a hungry baby of mine is right here and at this time needy and ready to receive food* and whose imperative content is the directive *at this time drop food into this baby's mouth* Xor something of that sort. But none of that complicated content is articulated, of course. This PPR need not contrast, for example, with any PPR that says anything about any non babies-of-mine or about actions to take other than dropping food. Indeed, the bird may have, as we might say, "no idea" what it is doing, as we would conceive what it is doing. Similarly, it would surprise me if the beaver tail splash gave rise to anything we would consider "thought" at all in the beavers that hear it. But it doesn't follow that the splash is not representational.

The very simplest of inner representations, then, seem to have the following three characteristics, none of which seem characterize the kinds of inner representations humans typically communicate using sentences. First, these representations tell in one undifferentiated breath both what the case is and what to do about it. Second, they represent the relation of the representing animal itself to whatever else they also represent. Third, they tend to be highly inarticulate, the representational systems in which they occur being devoted to highly specific tasks, so that very few contrasts in possible content are needed or possible.

The tendency for systems of inner representation to be devoted to highly specific tasks in most animals is evident from studies of animal learning. On this point, it may be sufficient to quote the Princeton ethologist James Gould on what he terms the "rigidly programmed plasticity" (Gould 1982, p. 268) characteristic of most animals:

> ...learning is adaptively programmed so that specific context, recognized by an animal's neural circuitry on the basis of one or more specific cues, trigger specific learning programs. The programs themselves are constrained to a particular critical period, ... and to a particular subset of

> possible cues. Nothing is left to chance, yet all the behavioral flexibility which learning makes possible is preserved. (J. Gould p. 272.)
>
> Learning, even in higher vertebrates, seems less a general quality of intelligence and more a specific, goal-oriented tool of instinct. Bouts of learning such as food avoidance conditioning, imprinting, song learning, and so on, are specialized so as to focus on specific cues --releasers-- during well-defined critical periods in particular contexts. Releasers trigger and direct the learning, and in general the learned material is thereafter used to replace the releaser in directing behavior. As a result animals know what in their busy and confusing world to learn and when, and what to do with the information once it has been acquired. Most learning, then, is as innate and preordained as the most rigid piece of instinctive behavior. (p. 276).

In this preordained way, many animals learn either by trial and error or from conspecifics what to eat and what not to eat; some learn from others which local species are their predators; the European red squirrel laboriously learns how to open, specifically, hazel nuts; the oyster catcher laboriously learns to open oysters; and the chimp laboriously learns to open nuts by using a rock and an anvil. Speaking generally, what animals are capable of learning, hence, it is reasonable to suppose, what they are capable of developing representational systems to support, tends to be closely tied to specific skills or at least specific ends found to be useful in the past history of the animal's species. Count this as a fourth typical characteristic of many inner representations in other animals.

In contrast to these reflections on typical inner representations in other animals, we humans are capable, first, of having many beliefs that we know of no practical uses for. And we can have many explicit desires and goals that we don't know how to implement because we lack the relevant information. Clearly, indicative and imperative mental representations can occur quite independently in us, obliging us to use practical inference to join them together again in the production of action. We definitely are not pure pushmi-pullyu animals. Second, we are capable of having beliefs about things and affairs that are very distant from us and about things whose spatial and temporal relations to us we have no knowledge of at all. Our indicative representations do not, in general, represent relations of situations and objects to us. Third, on the assumption that beliefs and desires can be at least as articulate as sentences used to express them, they must have considerable inner articulation, allowing contrasts in at least subject and verb phrase, and often in direct and indirect object, prepositional phrases and so forth. Fourth, we seem to be capable of learning many skills and of learning about many kinds of affairs that neither we ourselves nor our species has previously had any use for, and developing the necessary representations accordingly. We seem to be capable of harboring many representations that are not dedicated to any particular practical purpose, but that remain quite uncommitted.

Now I think there is no question but that we humans also use many forms of representation, in perception, below the perceptual level, and so forth, that are PPRs; that show relations of ourselves to affording objects; that are inarticulate or, like food aversions, that are learned according to built in triggers. And I think there is no question that many animals harbor inner representations that are not just PPRs; that show relations of objects to one another rather than merely to the animal; or that are somewhat articulate, and so forth. My point is merely that in exploring the question what kinds of signs a particular animal might be capable of learning to interpret, we should explicitly take into consideration whatever we can discover about the kinds of inner representations the animal is capable of employing. For it seems clear that to comprehend a sign with a certain force, content, and articulation, the animal must be able to match it with an inner representation with similar force, content, and articulation.

What then are the steps from, beginning at the bottom rung, the sort of inarticulate pushmi-pullyu comprehension the bee has and the dim sort of pushmi-pullyu comprehension that mediates responses to behavior releasers, to articulate, well-differentiated and uncommitted human beliefs and desires at the top? Well, of course, I don't really know, but here are a few speculations.

One thing that apparently occurred with the evolutionary development of the forebrain is that much incoming perceptual information became divided into two somewhat independent channels, a dorsal channel that yields representations, for example, of direction, distance, angle, location and size of objects <u>relative to the perceiving organism</u>, and a ventral channel that yields representations of objective or nonrelative shape, size, color, texture and so forth, used for determining what object or objective kind of object is being perceived. (For a review, see, for example, Norman 2003?). The capacity to represent objective non-observer-relative properties of objects as distinguished from the effects these properties are currently having on the perceiver requires the development of what are called "perceptual constancies," such as the ability to recognize the same size at various distances, the same shape at various angles, the same originating sound through various kinds of interfering noise, the same color under various lighting conditions, and so forth.[2] This is not the ability to make anything like subject-predicate judgments, of course, but merely to represent observer-independent properties, certain configurations of which are then recognized as indicating certain objects or kinds of objects. Consider, for example, a connectionist net that has learned to recognize seven faces each from any of various angles and at various distances, but which, if given any new face to learn, has as hard a time as it did with the first (indeed harder, because of interference). Suppose instead it had somehow learned to recognize same-shape-again quite generally. If it could do that, the next face might be learned in one trial.

Representations exist and show significant articulation only in so far as they are used as representations, and in so far as the contrasts corresponding to these articulations matter for these purposes. The other side of the division between the two kinds of representations must be the development of two kinds of uses for these representations. On the one hand, the animal develops general skills in navigating among and manipulating objects-in-general, skills that might be applied to <u>any</u> object whose shape, size, and orientation, distance and so forth, relative to the organism, are perceived. On the other hand the animal develops the capacity to recognize various specific objects and specific kinds of objects, each from a variety of distances and perspectives, and through a variety of intervening media and through different sensory modalities. These are objects and kinds of objects suitable to certain purposes, such as chasing, fleeing from, eating, nest building and so forth, but that require to be navigated among or manipulated in order to be used. Thus the animal perceives via the ventral system which kinds of objects to run from, which to approach, which to pick up, which to eat, which to climb up on, and so forth, while it perceives via the dorsal system, the relations to itself of these objects, which relations must be taken into account to guide its motions with respect to the objects. How would this development affect the four aspects of mental representation mentioned above?

First, representations that result from the achievement of perceptual constancies --representations of objective shape, color, size and so forth-- would seem to be intrinsically uncommitted representations. There is unlikely to be anything relevant to an animal's immediate activities that follows from the presence, for example, of objective sizes or colors or shapes simply as such. Representation of these properties, say, in early vision, have no one particular use but any of many possible uses, depending on what kind of situation or object in the environment they help to identify. An indicative representation that is not dedicated to any particular use but has many uses is still a representation only because it has uses, but it is not a PPR. At least it is moving away from being a PPR. If it has an open ended set of uses, as in the case of an animal that can learn to identify many new kinds of objects for use by first representing their properties, it certainly is not a PPR, but has a purely indicative character.

On the other hand, second, there seems no reason to suppose, that the separation of representations of objective properties from representations of relations of these objects to the perceiving animal would result, just as such, in replacement of pushmi-pullyu representations with independent indicatives and imperatives. Rather, the immediate result would seem to be the replacement of inarticulate PPRs with articulated ones that explicitly represent what objects are where (or otherwise significantly related to the animal), and thus immediately guide the animal's activity. PPRs of this kind would represent the kind of affordances Gibson had in mind when he said that apples are perceived as affording eating, mailboxes as affording letter posting and so forth. Consider, for example, a cat frightened by an approaching dog. The dog affords (requires) escaping from, which can only be done if the direction of approach of the dog is part of what is represented in perception. The direction of the dog, combined with the direction of something perceived as affording cover, directs taking cover or hiding, rather, for example, than running (on some other occasion) to something perceived as climbable. Thus, although on a deep level the animal now harbors some purely indicative representations, there is no reason to suppose that it harbors any purely imperative representations.

Third, the articulate nature of the PPRs that result from the dorsal/ventral separation allows the decomposition of undifferentiated skills into sub-skills that may be learned or practiced within certain contexts and then recombined in new situations. The capacity to recognize a certain kind of affording object can be developed in some contexts but then reapplied in others. Likewise, the ability to manipulate or alter relations to objects can be developed or practiced in some contexts and then reapplied in others. Much playing in mammals seems to be devoted to developing such general skills.

Fourth, the capacity to recognize and represent objects articulately as differentiated from their relations to the perceiving subject might naturally be applied to the learning of new and different affordances connected with those very same objects. For example, if the dog is good at recognizing its master as an afforder (if approached in the right way) of food, this ability can be put to good use in learning to recognize and then approach its master in the right was to get let out the door.

But if the same object in the same relation to the animal affords the animal different things on different occasions, it begins to look as if a purely indicative representation of the object bearing a certain relation to the animal may be emerging: Master is in such and such spatial relation to

<u>me</u>. Here we should go slowly, however, for two reasons.

The first is that the completed representation of most affordances may be considered as involving perception, as well, of the animal's state of need or appetite. The more careful statement of the affordance the animal perceives will then be in terms of satisfaction of that need or appetite. The dog perceives a hunger satisfaction affordance or an exercise-need satisfaction affordance, and so forth. Lifting completely out of the domain in which pushmi-pullyu representations reign may not be so easily achieved.

The second reason is that, even though recognition of the same object or kind of object may in some cases be involved in the animal's recognition of more than one kind of affordance, limitations on what the animal is (as Gould put it above) preordained to learn may be very strict indeed. Thus the animal's perceptions of most situations and objects are likely to remain dedicated to quite specific kinds of tasks, the nature of which has been dictated by past history of the species or, to a some degree in more flexible animals, past history of the individual. Similarly, it is unlikely that an animal should learn to recognize any object or kind of object for which neither it nor its species has yet found any practical uses.

Notice, last, that the separation of ventral and dorsal channels for perception of objects and their relations to perceivers has no tendency to free inner representations from representing only objects as currently related to their perceivers. Thus far we have no account, for example, of how an animal might come to represent objects distal to it in time or space without also representing its current relation to those distal times or places, or without these relations being immediately germane to current action. A step in that direction is in fact taken, however, by many animals, indeed, perhaps even by some that are relatively simple, as follows.

Many animals apparently construct and use something like mental maps of the locales in which they live. Among these, perhaps, are the honey bees.[3] If you trap a honey bee and release it in a locale with which it is familiar but from which the hive cannot be seen, it will fly up a bit, circle around as if to identify its current location, then fly off in a beeline for home. A number of things are very interesting about this development. First, these maps are not just representations of the relations of other objects to the perceiver. Relations of that sort keep changing, so there would be no obvious point in recording them for future use. Rather, these maps apparently represent merely relations of other objects, of various places, to one another. A second interesting thing is that unlike perceptual representations, these maps are constructed gradually over time and stored away for future use. As such they appear to be purely factual representations, representations of what some part of the world is like, apart from any particular projects the organism currently has in progress. But perhaps the most interesting thing is that for any such representation to be used, it will have to be combined with or temporarily joined to a current perceptual representation, namely, one representing the animal's current location and orientation within the domain mapped. Joining two representations in this manner to yield a representation of which way to go --that is, to yield a PPR-- looks a lot like mediate practical inference. Indeed, there is even a middle term. The same location has to be represented twice, once in its relation to other things not currently perceived, and once in relation to the perceiver as where the perceiver is now.

Do bees then actually make inferences? Perhaps so. Or perhaps the phenomenon is more parallel to they way a connectionist net may be able to fill in the rest of a configuration on which it has been well trained via Hebbian learning when presented with only a portion of that configuration. In either case, we should not let ourselves be carried away. That an animal can join one kind of representation with another or complete a partial representation for some specific kind of purpose does not make the animal rational. You are able to join visual representations from one eye with those from the other, using the overlap as a middle term, and thus derive representations of depth, but that is not what makes you rational.

Similarly, that an animal can collect and later use one kind of purely factual information, information about the space it lives in, has no implications for whether it can represent any other detached facts. That it collects and remembers information about local spaces depends on the fact that this kind of information has, often enough, been <u>used</u> during evolutionary history, used, indeed, in entirely specific ways. Similarly, many species of birds can remember hundreds and even thousands of caching places in which they have left food for future use. It does not follow that they are capable of collecting and remembering any other kinds of facts. Nor does it follow that they can use knowledge of these facts for any purpose other than finding food again when they are hungry. It is likely that the representations of fact that these animals collect are entirely dedicated to very specific uses. They are to be used for completing PPRs of very special predetermined kinds.

Parallel to the way in which animals collect specified kinds of factual information for predetermined uses, they may also collect certain kinds of skills out of the context of serious use. Young mammals, in particular, do a lot of playing. But once again, the things that they play at are always closely related to future uses. Animal play develops not arbitrary skills but skills for which the species has historically had uses.

Now it is true that through rigorous and careful step by step training by humans, individuals of many higher species can laboriously be brought to recognize perceptual affordances of kinds quite remote from any they were specifically designed to learn. They have some capacity to recombine their abilities to learn to recognize objects and to remember successful perception-induced response sequences so as to produce behavior patterns of kinds fairly remote from any anticipated in the histories of their species. There are three things that I strongly suspect they are not able to do, however, or to do at all well, but that humans seem to do quite easily. One is to represent pure facts that concern situations or objects of a sort that have not yet proved to be of use either to the animal or to prior members of its species. The second is represent facts about world affairs that have entirely unknown relations to the animal. The third is to be motivated by representations that do not originate from the animal's perception of its current needs and/or current environment.

Concerning the third, notice that the motivating representations we have been discussing are all PPRs. Typically, the indicative faces of these PPRs represent facts about the animal's current needs, coupled with facts that concern its immediate environment, as joined perhaps to some stored knowledge of the relation of the immediately present part of the environment to the wider environment which helps to fill out the animal's`perception of its current relation to more distal affording situations, places or objects. Even our most respected and intensively studied relatives, the monkeys and apes, seem to derive their motivation entirely from perception of the current situation. Thus, for example, Merlin Donald summarizes the literature on signing in apes: "...the 'meaning' of an ASL sign to an ape is simply the episodic representation of the events in which it has been rewarded..." (1991, p. 154) and "The use of signing in apes is restricted to situations in which the eliciting stimulus and the reward are clearly specified and present, or at least very close" (p. 152). No dog, I suspect, or even chimp, wonders where its next meal is coming from unless it is already hungry, nor does it wonder how it will cope next winter. Of course, appropriate migrating behaviors are elicited, in certain species, by natural signs that current food sources are running out, or by natural signs correlated with the immanent approach of winter. The indicative facets of the PPRs that are responses to these natural signs indeed do, though quite inarticulately, concern the future. These PPRs will produce appropriate behaviors only in the event that these future events are indeed imminent. What this shows, however, is only that animals are sometimes capable of perceiving the future, things temporally distal, just as they are capable of perceiving things spatially distal. Similarly, you must perceive the future in order to position yourself to catch a ball now in midair. It does not follow that you, or the animal, has left the level of PPRs.

But some human mental representations seem be free both from the yoke of historical usefulness and from the necessity of representing relations to self. And some motivating representations seem to be free of the bonds of currently perceived affordances. Unlike other animals, we represent and remember thousands of facts of kinds for which neither we nor our ancestors have yet found practical uses. The non-fiction sections of libraries are repositories, largely, for immense collections of such facts. We are able to interpret natural signs and also linguistic signs of world affairs that are distant from us both in time and place. We think about both the past and the distant future. We interpret signs of distal affairs and remember these facts even when we have no idea what relations these affairs bear to us. I know, for example, that gerbils come from the desert, but I have no idea what desert, or what use my knowledge of this fact about gerbils might have. We are adept at learning to interpret new kinds of signs, not just human language signs, but at learning how to read meters and scopes and information filtered through a multitude of other instruments. Apart from us, perhaps only apes can learn to interpret even visual information reflected from a mirror, and then only for guiding current activity. We notice and remember not just what we can cause, or what causes something we want, but what causes what, quite out of context.

We also spend huge amounts of energy and time developing skills, both physical and intellectual, for which neither we nor our ancestors knew any practical uses. We practice bouncing balls, juggling, manipulating Rubik's cubes, riding skate boards, cracking our knuckles, wiggling our ears, blowing bubbles, whistling through our teeth, spinning around to make ourselves dizzy (children often love this), and so forth and so forth.

Similarly, we collect dreams of things we would like to do or have done, places we would like to go, things we would like to have or to be able to build, without having any notion, yet, how to fulfill these dreams. Certainly these dreams are not currently perceived affordances. Nor are they representations of currently perceived needs. In short, we appear to be compulsive collectors of all kinds of junk!

Looking at the evolution of these strange capacities and behaviors, it is clear, of course, that although many or even most of them may never find uses, the general disposition to collect junk does find uses. If you have enough storage space and a good enough retrieval system, some pieces of that junk may well come in handy sometime, though there was perhaps no way to tell in advance which pieces. But it is not just that we have bigger storage barns than do neighboring species, bigger brains, although that may be part of it. What we really are alone in having, I suspect, is what Dennett likes to call "Popperian" minds.

We have the capacity and disposition to play games in our minds, entirely divorced from current perception, tinkering with the collected junk to see what might be built out of it that would be useful or help fulfill otherwise empty dreams. We do trials and make errors in our heads. We learn in our heads. It is because we can do this that we can represent desires and goals of kinds that neither we nor our species have ever realized. These desires are imperative representations designed for a job, namely, to become fulfilled someday by means of lucky tinkering. It is because dreams and desires of this kind <u>are</u> sometimes fulfilled that our cognitive mechanisms have been designed to produce them. Indeed, this is what makes them be (intenTional) representations. Without this they would have no biological uses, hence could not be representations at all.

What exactly is the lesson then? If an animal lacks the capacity to form mental representations having certain kinds of content, obviously it can not learn to understand signs, either conventional or natural, that carry those contents. But perhaps most of what we humans convey with signs is of a kind that, for animals without Popperian powers, would be utterly useless for them to represent. And for an animal to represent what it can have no use for representing is actually a contradiction in terms.[4]

References:

Anscombe GEM (1957) Intention, Ithaca: Cornell University Press.

St. Augustine (1986) On Christian Doctrine bk II (c. 427). Trans D.W. Robertson Jr. New York: Macmillan.

Dennett DC (1996) Kinds of Minds, New York: Basic Books, HarperCollinsPublishers.

Dretske F (1981) Knowledge and the Flow of Information, Cambridge MA: MIT Press.

Dretske F (1991) Explaining Behavior, Cambridge MA: MIT Press.

Donald M (1991) Origins of the Modern Mind, Cambridge MA: Harvard University Press.

Dyer F C (1996) Spatial memory and navigation by honey bees on the scale of the foraging range. J Exp Bio 199: 147-54.

Gallistel CR (1990) The Organization of Learning, Cambridge MA: The MIT Press 1990

Gibson JJ (1966) The Senses Considered as Perceptual Systems, Boston, MA: Houghton Mifflin.

Gibson EJ (1977.) The theory of affordances. In: Perceiving, Acting, and Knowing (Shaw RE, Bransford J, Eds), Hillsdale, NJ: Erlbaum.

Grice P (1957) Meaning. The Philosophical Review 66, 377-388.

Gould J (1982) Ethology; The Mechanisms and Evolution of Behavior, New York: W.W. Norton.

Gould J (1986) The locale map of honey bees: do insects have cognitive maps? Science 232: 861-3.

Gould J and Gould C (1988) The Honey Bee. New York: Scientific American Library.

Kant I (1798) Anthropology From a Pragmatic Point of View bk. I, sec 29, trans. V.L. Dowdell (Rudnick H, ed) Carbondale: Southern Illinois University Press.

Lorenz K and Tinbergen N (1939) Taxis und Instinkthandlung in der Eirollbewegung der Graugans. Teirpsychologie 2: 1-29.

McFarland, D. (1981) The Oxford Companion to Animal Behavior, Oxford; New York: Oxford University Press.

Millikan RG (1996) Pushmi-pullyu representations. In: Philosophical Perspectives vol. IX (Tomberlin J ed), pp.185-200. Atascadero CA: Ridgeview Publishing. *Reprinted* in Mind and Morals (May L, Friedman M, eds), pp145-161. Cambridge MA: MIT Press,.

Millikan RG (forthcoming) Varieties of Meaning; The Jean Nicod Lectures

2002, Cambridge MA: MIT Press.

Norman J (2003?) Two visual systems and two theories of perception: an attempt to reconcile the constructivist and ecological approaches. Behavioral and Brain Sciences 24 (6): XXX-XXX.

Ockham, William of (1495) Super Quatuor Libros Sentenarium Subtilissimae Earumdenque Decisiones II qu. 25, trans. S. Tornay. In: Ockham: Studies and Selections. La Salle: open Court, 1938.

Searle J (1983) Intentionality, Cambridge UK: Cambridge University Press.

Tinbergen N (1951) The Study of Instinct, Oxford: Clarendon Press.

Wehner R and Menzel R (1990) Do insects have cognitive maps? Ann Rev Neurosci 13: 403-14.

Wehner R and Menzel R (1990) Insect navigation: use of maps or Ariadne's thread? Ethol Ecol Evol 2: 27-48.

---

1. For a discussion of the difference between inference and translation, see (Millikan 2004, chapter Nine).

2. On perceptual constancies, see any elementary textbook on general psychology or perception and cognition.

3. See, for example, (Gould 1986; Gould and Gould 1988). For dissent, see (Wehner and Menzel 1990; Dyer 1996).

4. For a more detailed discussion of many of the matters addressed in this essay, see (Millikan, forthcoming).