KANTIAN REFLECTIONS ON ANIMAL MINDS

Ruth Millikan

A Priori
www. apriori. canterbury. ac. nz

December 2001

Kantian Reflections on Animal Minds

Speculating on the evolution of inner representations. Obviously the answers must eventually come out of empirical studies, but you have to know what to look for. Only a small part of the story today.

Humans probably have the whole hierarchy of earlier forms inside us. Most of our purposes and most of the facts we take into account may be represented far below the level of belief and desire and intention. (Lecture on cross-purposes). Sweet tastes may be intentional signals of nutritional value, for example, but they are not beliefs about nutritional value, nor is a sweet tooth a desire for nutrition.

1. Intentional Icons

Contra Lorenz: Intentional icons are not built into the organism but acquired by the individual in response to the environment. They exhibit a dimension or dimensions of possible variance running parallel to possible variances in the environment (fact representations). Or they exhibit variances that correspond to variations in what they are supposed to produce (goal representations). Or they do so anyway when all goes swimmingly!

2. Pushmi-pullyu Representations

The most primitive ICONS and SIGNALS are at once fact representations and goal representations, telling both what is the case and what to do about it. Thus each of the various chemical "messengers" that run in the blood stream "tells" about some particular state of the organism's physiology and, in stimulating a physiological response appropriate to that state, also "tells" what to do about it. The dances of honey bees are representations (not inner ones) at once of where the nectar is and of where the watching worker bees are to go. Representations having this sort of lumped double structure are "pushmi-pullyu" representations" after Hugh Lofting's mythical creature of that name (Millikan 1996).

At the opposite end of the spectrum from pure P-P signals and representations are human desires and beliefs.... Here fact-representing and goal-representing functions have completely separated. Our beliefs often concern facts that we have no notion how to use in action, and our desires include many we have no idea how to satisfy given actual conditions . Having separated fact representations from goal representations, it is necessary somehow to reassemble them for use. Practical inference is needed in order to do this.

What kinds of steps might there be in the evolution from systems employing only simple P-P signals and representations to a system employing beliefs, desires and inference?

Notice what the bee dance tells is the relation of the bee to the nectar. Another development to consider is the move from self oriented representations to perspectiveless ones.

3. Affordances

A step above mere behavior "releasers"—hunger triggers food seeking behavior, the open beak of the gosling triggers dropping food in— are perceptions of Gibsonian affordances. They are Pushmi-pullyus. Perceptions used by perception-action systems. Perception of aspects of the lay out of the environment that affording, say, passing through, climbing up on, grasping, eating, hitting with... "Possibilities for action". What the perceptual P-P representation shows is some relation that the animal bears to its environment that will afford something to the animal, granted its response is guided in the right way by that relation. Gibson claimed these relations are represented by invariances in the proximal stimulus. If so, being guided by the proximal stimulus is using it as an icon of, representation of, the animal's relation to the distal stimulus. It may well be then that from the point of view of physical science, what the animal perceives of the world are strangely complex and arbitrary properties and relations. Nor does the animal perceive these as fact about the world but, intrinsically, merely as being for use in a specific way (Heidegger would have said "zuhanden").

4. Perception and Cognition as Search Techniques

Affordances often come in chains. Let me explain. To be in a position such that a primary goal, such as having a fly in the stomach, is achievable by utilizing just one perceived affordance is a blissful condition. Call such a condition a "B- condition".

It is typical of organisms that do not move about that they merely wait for B-conditions to pass by them and then seize the moment (venus fly traps, barnacles). More sophisticated animals make an effort to maneuver themselves into B-conditions, for example, the frog has sense enough to sit in a place that attracts flies. The simplest way to attempt to maneuver oneself into some B-condition or another is, of course, just to wander about aimlessly hoping to bump into one. Better, one can use some sort of systematic search technique. One way to view the story of the evolution of perception and cognition is as a story about the acquisition of more and more sophisticated search techniques for maneuvering oneself into B-conditions. These are techniques for raising the probability of eventually getting into places or positions from which one can act immediately and productively, or getting out of conditions in which chances of disaster are high.

For the simplest animals the cardinal principle involved here is very elementary. Be constructed such that you can perceive affordances that will afford your probable placement in new positions from which you are likely to perceive new affordances that will afford your probable placement in newer positions from which ...and so forth...finally probably placing you in B-conditions. The trick is that this series of probabilities should have a product greater than the probability of B-conditions just happening along without any action on your part, the higher the probability the better. Thus the search domain is narrowed and then narrowed again. Affordances that afford affordances that afford affordances. The newborn baby's response to a touch on the cheek is to turn toward it, thus raising he probability of feeling a nipple on the mouth which will afford nourishment. Very simple animals show various kinds of taxis likely to take them into conditions where food affordances are prevalent or certain danger-avoidance affordances less likely to need to be utilized. And so forth. Never looking ahead, completely controlled by the current environment (although, typically, by aspects of the distal environment. Even releasers are control by the distal environment.)

Extremely complicated long and branching chains of affordances leading to the probability of finding one or another other affordances, leading to the probability of finding one or another...and so forth, may be grasped by some animals, resulting in surprisingly intricate, even flexible, behaviors. And it may be that correctly quantified increases in potentiations of response dispositions resulting from other relevant stimuli, other simple ques from the environment encountered along the way help account for the tendency of the animal to chose, from among equally available and relevant affordances, those objectively associated, in the animal's particular circumstances, with higher probabilities of eventual success. Gallistel (1980) speaks of such an animal as being governed by a "lattice hierarchy." Such an animal might be capable of navigating in the space-time-causal order from a great variety of different starting positions relative to its goals so as to reach them with reasonable probability.

On the other hand, such an animal might also be subject to failures that strike us as rather ridiculous. Dennett is fond of the digger wasp that can be sent into a behavioral loop from which it never emerges by moving the prey it has paralyzed a few inches away from the door of its nest every time it goes inside to inspect, preparatory to stashing the prey(Dennett 1984). I once watched a pair of hamsters repeatedly stumbling over one another as each returned a large cracker to its own corner over and over from the other one's corner just opposite. According to Gallistel, even in quite flexible animals, available behaviors are by no means always applied to relevant situations, even when the increment is very small. The ability to combine abilities. He tells, for example, of a hybrid species of lovebirds that repeatedly tried to carry strips of bark to be used for nest building by tucking them into their tail feathers, only invariably to lose them on the flight back to the nest. These birds were perfectly capable of carrying the strips safely in their beaks, but did so only 6% of the time (Gallistel 1980, 306-8). Even though the behaviors of an animal may be governed by many different signals and representations hierarchically arranged in ingeniously functional ways, I think it would be natural to say of an animal with nothing beyond this internal sort of architecture that it does not think.

Nor does introduction of the capacity to learn by trial and error, at least as conceived by the associationist tradition of psychology, add more than details to this general picture. In the end such learning merely affords quicker ways than genic selection to forge what are, however, basically the same kinds of behavior chains as before. Associative learning concerns the <u>genesis</u> of the lattice hierarchy only, and does not affect its basic structure. Programming that is done during ontogeny rather than phylogeny may permit more flexible adaptations in the behavior of individuals, but the resulting control structures operate in accordance with the same principles. That is, purely Skinnerian animals also do not think.

No thought. Also, notice, releasers, affordances, but no <u>facts</u>.5.

5. The First Pure Fact Representations

We humans use facts <u>with inference</u>. But there may be ways that facts could enter before inference and without disturbing the lattice hierarchy in any way. 2 ways.

First is the use of multipurpose representations that represent always the same kind of world affair but afford the animal different possibilities for action given different motivations. This could happen merely as a side effect of economic construction of the perceptual apparatuses. If you eat both mice and frogs, it is not economical to have completely different perceptual processing mechanisms, for example, separate pairs of eyes, for perceiving these. Similarly, if you eat mice and flee from snakes. If you have a complex structure such as an eye, clearly you should use that same eye for as many of your various purposes as it can be made relevant to. You should avoid specialized adjustments that will make it unsuitable for multipurpose use. Consider the obstacles confronting the design of any apparatus with a sophisticated capacity reliably to make representations showing affordances of <u>distal</u> objects. To be as useful as possible, such an apparatus must enable recognition of the affording distal object or property and its relevant relation to the animal over as wide a range of object-animal relations as possible (not just dead center under the animal's nose) under a variety of mediating conditions (under various lighting conditions, sound echo conditions, etc.), despite distractive intrusions affecting proximal stimulation ("static" such as wind noise or shadows or reflections). It will be easiest to make it do this if, in the first instance at least, it registers simple objective physical properties and relations that fall under uniform physical laws rather than disjunctive gerrymandered properties and relations, even though the latter may be more immediately useful for certain entirely specific tasks. (Akins on perception of heat and cold) Also, the registration of simple objective physical properties and relations is more likely to be useful in guiding a variety of different activities. Consider, for example, visual perception of the arrangement, sizes, shapes, textures, orientations and relative distances of the objects in ones vicinity. This kind of information can be put to innumerable uses in the guidance of action. It is best, then, if at some stage of processing at least, the eye produces representations that are not merely of gerrymandered single use distal affairs. But the more purposes it serves, the more disjunctive, hence indeterminate, is the "pullyu" aspect of the P-P representations it produces. The representation that has a dozen or a hundred uses depending on the particular state of potentiation of the nervous system, all of which uses require it to be aligned with the world in exactly the same way—the semantics is the same— becomes at the limit an any-purpose, hence purely fact-representing, representation. (It doesn't follow that every aspect is represented moment by moment.)

This leads immediately to another principle, the production of surplus information. The more versatile a perceptual apparatus becomes, the more likely it is to be relying on quite general principles in producing its representations, and this is likely to result in a lot more simply coded natural information being captured than is consumed by the specific uses for which it was designed. If you build a visual system so that it that can see mice, frogs, snakes and also conspecifics, then it undoubtedly brings in enough information to see many other medium sized objects as well. One only needs then to design into the animal some principles or mechanisms by which experiments can be made in the use of this extra information, for example, principles by which it searches for patterns of association involving this information, and you have an animal that employs completely general purpose representations and employs them by design. It is designed to perceive any of certain general kinds of facts for as yet unspecified uses. You have an animal that harbors pure fact representations.

6. The Construction of Objective Space

Notice, however, what these fact representations represent. They represent, always, only relations that things in its environment bear to the animal. They do not show how things are independently of the animal, in relation to one another. For direct guidance of behavior by the environment, the organism only needs an awareness of environmental relations to itself. Moreover, the representations that such an animal uses do not represent its relation to the environment explicitly. The sentence "there is a mouse a yard in front of Tabby" is articulated to show mouseness, the relation a yard in front of and also Tabby, each explicitly. But Tabby's perceptual view of the mouse as needed for stalking it does not show Tabby explicitly. No significant transforms of it show the relation of anything other than Tabby to the mouse, nor do any transforms of it omit Tabby. Tabby's perceptual views of the world always concern Tabby since they always concern her position in the world, yet they never, as it were, mention Tabby.

Now for the Kantian part of my story. There are better strategies than constructing a lattice hierarchy network to search for convenient and safe paths from wherever one happens to be into B-conditions. These strategies involve the construction of inner representations of various aspects of the world as they exist apart from the animal's special position in it, representations of the world, as it were, "in itself" rather than relative to the animal. The strategy requires, however, a rudimentary form of inference. I'll start with a simple illustration, namely, the advantages of employing cognitive spatial maps. Then I'll suggest some richer applications of the principles involved.

Suppose that you wanted to find your way home, but that all you had to go by was a collection of memories showing, from the point of view only or your own past perceptions, paths you had actually taken at one time or another from one place to another. Perhaps these memories form an associative network, —this model suggested by Deutch— the want-to-go-home signal potentiating nodes representing the various places you have been, with attraction lessening as the number of links from home in the chain increases and also as the lengths of the individual links increases. You take the path that sends the strongest signal to the node representing the place that you now are in. The trouble with this arrangement —as with any arrangement resting on a RECORD MERELY OF PREVIOUS ORDERINGS OF ONE'S OWN EXPERIENCE— is that this way of representing the paths tells nothing about the general geometry of the underlying space in which they lie. True, you would have enough information to get home from where you are, but you would be very lucky if this information happened to put you on a direct route. What you need to have mapped to tell how to get home fastest is how the various paths lie relative not to your own past history but how they lie relative to one another in Euclidean space. You need to know how the paths twist and turn, at what angles they intersect with one another, and so forth, within that space.

It is for this reason, presumably, that even some quite lowly creatures apparently make maps, make inner representations, of the locales where they live. For example, Gallistel and Gould & Gould argue that there is excellent evidence that bees do this. Bees apparently record the positions of various landmarks in their locale relative to other landmarks, rather than relative to themselves, in a medium having a topology and metric isomorphic to Euclidean space. Using a map you can be guided directly from one place to another regardless of whether you have traveled any part of the route before. Thus a bee, when transported by any route to any location in its territory, knows how to fly directly home, or to another location it "had in mind," as soon as it has taken its bearings. The bee knows how to take short cuts. (Ded reckoning?)

To make a map of an area requires not merely representing connections between places one has oneself happened to go, but taking account of the general geometry of space so as to leave empty areas of the right kind on the map for the places one has not happened to go. A "tabula rasa" is the traditional term for a mind that comes into the world with no preconceptions. Any actual blank tablet, however, has a definite geometry. On the customary kind of tablet, only two dimensional Euclidean figures can be drawn. The bees' tabula rasa is apparently a blank isomorph of Euclidean space in at`least two dimensions, waiting to be filled in with landmarks. It is the bee's version of Kant's pure intuition of space. (What if bees just look? Recognize landmarks?)

Imagine the bee's cognitive map as like a road map with gas stations, motels, good restaurants and roadside tables marked on it —the hive and good nectar-gathering sites and good places to colonize if necessary. The map might also show the current position of the bee itself, as the animated maps displayed in the front of some transcontinental airline coaches show the position of the airplane one is in. But even if the bees map was like that, notice that it could not stand alone as a guide to action. The bee will need to perceive its position relative to its environment directly as well, even if only to keep its image on the map in the right place. In order to use a map, you have to know where you are. The bee's map must be joined to its perception by identifying a place on the map with a place as directly perceived, so that the two together can guide action. And only if this overlapping of content, this shared middle term, this same place being represented in each representation, is recognized could these two inner representations be joined together to yield the relation of the bee to its destination. But the joining of two inner representations, pivoting on a middle term to yield new information or direction, is nothing more nor less than mediate inference, in this case, practical mediate inference. Bees must make inferences! (Connectionist nets?)

7. The Construction of an Objective World

The lesson can be generalized. It seems likely that the best way to find direct routes through the spatial-temporal-causal order from wherever you happen to be toward B-conditions, toward conditions in which you can act to satisfy needs immediately, is to begin to construct representations, "maps," of the relations of the various aspects of the objective order <u>to one another</u>, rather than merely to yourself. Remembering the scenery that has happened to pass by on your own historical, private, wiggly, space-time line is not sufficient. Associative conditioning is not enough. Direct guidance by the im-

mediately given environment is not enough. To grasp possibilities for different kinds of safe and efficient action, to anticipate new paths through the objective causal order to B-conditions, it is necessary to anticipate something of the abstract structure of that order <u>in itself</u>. Rather than relying on mere associative conditioning, Inner representations of the objective world need to be constructed. Then, just as the bee finds the direct route home by searching its inner map rather than its environment, direct routes toward B-conditions can be located by searching among one's representations of the world. Searching in one's head for paths to B-conditions is much safer as well as quicker than searching outdoors.

An story has been circulating that some snakes detect mice for purposes of striking by <u>sight</u>, follow the trail of the dying mouse by <u>smell</u>, turn the mouse to swallow it <u>by feel</u>, and are incapable of recognizing a mouse for any of these three purposes through any modality but the assigned one. (Probably apocryphal!) Thus the mouse shows itself to the snake as three separate affordances that are not integrated into one object (first a strike me, then a follow me, then a swollow me). It would be for the snake as though these three aspects of the mouse were entirely separate pieces of the world that just happened to lie juxtaposed on the time line of the snake's experience. The mouse is not a single object with multiple properties but merely a string of associated affordances. Unlike such a snake, the animal that constructs representations of its objective world must gather together the fragments of the objective world it encounters and glue them together into objects. We might compare an archeologist who reconstructs ancient objects from a few broken fragments. But gluing pieces of the world together requires some sort of schematic plan of the general architecture it should have, its geometry and ontology, schemata for its representation.

This reconstruction would have to be done, I suppose, rather as Kant suggested. The animal must implicitly grasp, for example, that the world contains reidentifiable <u>substances</u> (Aristotelian) each with many <u>attributes</u>, entering into causal relations, and so forth. Just as the bees must have the pure intuition of Euclidean space, it must grasp the most abstract principles of the world's ontology, and then attempt to flesh out this schema with details relevant to its particular possibilities for action. What kinds of general schemas for world structure might be available to an animal? What aspects of the world might it find useful to reconstruct? Again, some of these basic aspects of the world's ontology were probably forms recognized by Kant. Besides space and time, of particular importance, I believe, are the categories of substance and accident, cause and effect, and the ontology that makes negative judgments possible. (I won't venture whether this kind of basic knowledge of world ontology is endogenous or not.)

Still, no representation of the objective world taken by itself can guide action. Representations of the objective world, even if they have one's destination clearly marked on them, are powerless to guide action. To be useful, they must be joined to representations showing part of that same world structure but from the present point of view of the animal. This joining of two representations to yield new information or instruction is a kind of practical mediate inference. (Intentionality does not require inference, but objectivity does.)

(That an animal makes inferences, however, does not yet imply that it is "rational" as that term is commonly understood. For example, its information may be tightly encapsulated and unavailable for general use. (Recall Gallistel's lovebirds.) We must not assume that if something is represented in one part of an animal's system it is available in others, or that access by one part of a system to information gathered in another is a transitive relation. (remember that a sweet taste in the mouth is not a belief about nutrition. Also the animal may be unable to represent negation, hence unable to recognize contradictions in order to avoid them.)

An animal that uses objective maps of its world has to make maps of its world. So it will probably devote some energies specifically to this purpose, exploring and prospecting. This will be "theoretical" activity, in Kant's sense.

8. Representing Invariance in Substances

One kind of useful objective construction is reconstruction of an object or space in three dimensions from those fragments of the energy it structures that are either accidentally encountered or searched out by the animal. This is the sort of construction that David Marr tried to explain in his theory of vision, and it is well known that Marr had to postulate that to accomplish this task the animal makes certain implicit or unrepresented assumptions about general properties of the layout of its environment. An animal that can reconstruct objects and spaces in this way will be able to utilize affordances that show themselves directly only from perspectives other than its own. For example, it may grasp that an object affords climbing up on or that a space affords passing through if approached from a different angle. It may also discover affordances that depend on properties of the whole, such as overall shape or volume. Such an animal will also be in a far better position to reidentify places, objects and ("natural") kinds of objects as it encounters them in different orientations.

No representation of the objective world can be used to guide action without joining it to representations from perception, and this joining is done by finding a middle term, by identifying something showing in one representation with the same thing showing in another. Thus, for example, the bee needs to be able to recognize the same place from a variety of perspectives if it is to use its map effectively. Also, reidentification is required in order to construct maps and other representations. The bee will know where to place a new landmark on its map by noting its relation to old landmarks already on the map, so it must be able to reidentify these old landmarks. In order to glue fragments of a broken object together you must be able to recognize when two fragments fit together, which requires identifying the same surface shape in the convex and in the concave. All the places where the glue goes when reassembling the world, are properties or entities that need to be identified as the same ones again encountered from different perspectives. Roger Shepard claims we have tripartite color vision because that is what is necessary not for color constancy (we don't have color constancy) but for object constancy, for reidentifying objects (I would add stuffs, and many natural kinds). Similarly, the ability to reidentify objects, kinds and stuffs, from a variety of perspectives by shape is central to nearly every other reconstruction task. Many other kinds of perceptual constancies serve the same purpose.

But we should turn the coin over. It should never be taken for granted that any animal recognizes when different kinds of its representations map over the very same portion of the world, as the (apocryphal) snake did not grasp that it chases, strikes and swallows is the very same thing. Similarly — Jerry Fodor likes to remind us of this— Oedipus had no idea that his thought "Mother" and his thought "Jocasta" had the same referent. Knowing how to reidentify a thing through all of its possible manifestations is clearly impossible. No animal could be perfect in this regard. It can always happen that an animal harbors inner representations that overlap in content without grasping this. The capacity to reidentify objective substances and to perform mediate inference is, I believe, the core of conceptual representation as distinguished from, for example, representation in perception for action. On Clear and Confused Ideas.

It can be helpful to for an animal to have an objective representation of various items in its immediate vicinity and of their relations to one another.

But objects change and they come and go. For this reason, no permanent detailed mapping of them is possible, or not without adding the dimension of dated time, and what is the practical use of a map of the past? Probably we're the only animals that can use such a map. Other animals store away knowledge only of the most stable structures in their environments. Or if they store knowledge of world features that change, they simply throw it away when outdated. (Bees and which flowers are good this week—Gallistel.) More exactly, I believe, they store away knowledge of "substances," using something like the Aristotelian sense of that term. For example, ordinary individuals, various stuffs such as water, wind, rain and rock, and natural and historical kinds such as animal and plant species (Aristotle's secondary substances). "Substances" are distinguished by the fact that one can learn things about them on one encounter that will remain true with some reliability on other encounters. Thus you expect the sourness of one lemon from having tasted another lemon and you are ready for John's sourness on one day from having experienced it on days before. That's like learning to recognize, from whatever perspective one currently happens to have on an object, the existence of affordances available from other points of view. The trick is to grasp how to locate and quickly reidentify objective substances about which relatively stable knowledge can be had, despite great variety in a substance's possible manifestations to one's various senses. And, of course, one also needs to grasp quickly what kinds of stable knowledge can be gathered about each. If a ripe blueberry is edible, another will likely be edible too. But if one fox is rambling in the open, that does not preclude that others will be waiting quietly in the brush. There is evidence from the child development literature for a boost from endogenous factors toward recognizing categories of substances relevant to these tasks.

It should be clear that no animal is going to map more than a very small portion of its world or more than a small proportion of that portion's objective aspects. In general, an animal might be expected to reconstruct only certain aspects that are closely relevant to its needs. Nor do we have to think of the project as the progressive construction of a giant multi-dimensional model of the world in the animal's head: combinatorial explosion. Surely although the animal puts some fragments together it merely stores others, carefully preserving those aspects that mark samenesses for it so it can join them up later if needed. That is, it prepares materials for later use in inference. Compare having a map of the whole of a city but in book form so that one must find various overlapping pieces and join them together to find the relations among parts of the city that are distant from one another.

9. Representing Invariance in Processes

Turning from substance and accident to the category of causality, besides representing substance and accident sameness, animals have differing abilities to represent constancies in world processes. Of course, every moving animal actively uses constancies in causal processes. Being guided by a perceived affordance when in conditions normal for its realization is participating in a stable causal process. But knowing how to follow out fruitful affordances on propitious occasions is not jet knowing about or representing the constancies that make this possible. But we humans, at least, do represent constancies in processes. We are not just stimulated to pursue affordances when we happen to encounter them. We remember what turns into what, and what happens if you do what.

In the simplest cases, we, and probably other animals too, merely think ahead to what will happen if we utilize a beckoning affordance, and react to the anticipated outcome with an advance or a withdrawal. Perception shows the possibilities and the animal chooses among them. Even the chimps seem not to have advanced beyond this. Even their smartest antics seem to remain situation-bound in motivation. But we humans seem to do something more. Earlier I suggested ways in which representations of pure fact might emerge from pushmi-pullyus. In humans, pure goal representations have emerged as well. How are our goal representations used?

When we act on the world, transforming it in some way, what we do has a causal outcome, one that we may be able to anticipate in thought. Similarly, when an animal roams about, the direction it takes from a given place has a "spatial outcome", one that it can anticipate if it has a mental map in its head. Suppose then that the animal's goal is to arrive at a certain place. It's spatial destination is marked on its mental map. It perceives the place where it now is, identifies this place on the map, and joining percept with map, heads straight to its destination. Now suppose instead that it has as its goal to be in a certain situation in its world to which it must traverse not just spatially but causally. It wants, say, to be sheltered in a certain sort of house. It has a goal representation that represents this aimed-for objective situation. It is to aim for this situation in the causal, not just the spatial, order. How will it use its knowledge of what leads to what in the causal order to direct its aim so that it starts off in the right causal direction? The difficulty here is that unlike ordinary space, the logical space of possible causal outcomes in time is not a connected space with a definite geometry. It is a space in which possibilities diverge and then diverge again in infinite variety. There can be no analogue here of dead reckoning.

In order to use knowledge of constancies in causal outcomes in a sophisticated way to govern behavior, one must become, what Dennett calls, "Popperian." One must make trials and register successes and failures in one's head, imagining one by one various alternative chains leading from one's current situation to others until hitting on one or another causal route to its goal. Though I have argued that bees too make inferences, to be capable of this sort of inference is far removed from the bees.

10. A Common Code

What eventually emerges in Homo sapiens is the ability to recognize and to map causal processes initiated either by the thinker or by extrinsic events, and the ability to represent a layout of ongoing events many of which occur in places at a great remove from the thinker. The bee, we suppose, constructs at least a two dimensional space containing enduring places and some temporary features (what color and where the best flowers are this week). Parts of this map are revised or updated quite frequently, but it does not represent a temporal dimension, at least not one with absolute dates (Gallistel 1990). A human, on the other hand, constructs a four dimensional map of a dated world in progress, mapping both things that endure (substances, places) and also what happens, both in his or her own locale and in other places. Many of these facts are represented apart from any known relevance to the thinker's practical interests, and inferences are made from these facts to further facts of the same disinterested sort. (JUNK!) Ultimately, the point of all this cognitive activity is to join up at crucial points with perception so as to guide action. But a more immediate aim often is merely the efficient production of representations of more and more of the world.

Now the perceptual representations that guide immediate action need to be rich in specific kinds of information, showing the organism's exact relations to a number of aspects of its current environment directly as they unfold during action. These representations may need to have variable structure of a kind that conforms closely to the variable structure of the organism-environment relations that need to be taken mapped to immediate action, current motion. Because they need to be constructed quickly, they may be constructed by modular systems that are relatively cognitively impenetrable (Fodor 1989). The job of the disinterested fact representations of cognition is not this, but instead, easy participation in mediate inference processes. This job makes its own special demands. There is no way to specify in advance in what specific kinds of inferences such a representation may need to be used. Facts are collected for whatever, if anything, they may prove to be useful for. While the representations of perception may need to be cast in highly structured multidimensional media suitable to the immediate purposes to which they are dedicated, cognitive representations may need to be cast in a simpler uniform medium that makes them easy to compare and combine. The ideal fact representation would be one that could be combined with any other fact representation having an overlapping content, a potential middle term in common.

Whether or not information can interact in inference depends not on its content but on its vehicle. Putting it graphically, if the first premise of an inference is represented with a mental Venn diagram and the second with a mental sentence, it is hard to see what inference rules could apply to yield a conclusion. Similarly, one might suppose, if the information coming in through the various senses were not translated into something like a common medium for the purposes of theoretical and practical inference, it could not interact in a flexible way. Possibly this is the fundamental difference between inner representations that are more "perceptual" and those that are more "cognitive."

In any event, an important question when studying the mental life of any fact-collecting species must concern the degree and the kind of interaction in inference that can occur among the varieties of representations it collects. Whether or not intentional contents can interact in inference does not depend on satisfaction conditions, but on how the content is articulated and represented, stored and retrieved, and especially, to what degree content identity is clearly marked across representations with overlapping content.

No rationality without negation.