Chapter 6. Information and Direct Perception

6.0. Introduction.

The purpose of this chapter and the next is to describe Gibsonian ecological psychology and to show that it can serve as an appropriate theoretical backdrop for Radical Embodied Cognitive Science. It hardly makes sense to do so other than in the context of the theoretical work of Michael Turvey, Robert Shaw and William Mace. Since the 1970s, Turvey, Shaw, and Mace have worked on the formulation of a philosophically-sound and empirically-tractable version of James Gibson's ecological psychology. It is surely no exaggeration to say that without their theoretical work ecological psychology would have died on the vine because of the high-profile attacks from establishment cognitive scientists (Fodor and Pylyshyn 1981, Ullman 1981). But thanks to Turvey, Shaw and Mace's work as theorists and, perhaps more importantly, as teachers, ecological psychology is currently flourishing. A generation of students, having been trained by Turvey, Shaw and Mace at Trinity College and/or the University of Connecticut, are now distinguished experimental psychologists who train their own students in Turvey-Shaw-Mace ecological psychology. Despite the undeniable and lasting importance of Turvey, Shaw and Mace's theoretical contributions for psychology and the other cognitive sciences, their work has not received much attention from philosophers. It will get some of that attention in the next two chapters. I will point to shortcomings in the Turvey-Shaw-Mace approach to ecological psychology, and will offer I take to be improved versions of each of the four main components of it. In this chapter, I will describe theories of information and of direct perception that differ from the Turvey-Shaw-Mace account; in the next chapter I will tackle affordances and abilities.

Given the debt that those of us interested in ecological psychology owe to Turvey, Shaw and Mace, this, no doubt, seems ungrateful.[1] Perhaps it is. But I would argue that because of the success of the Turvey-Shaw-Mace approach to ecological psychology, the field has become a true contender in psychology, cognitive science and artificial intelligence. Given the stability of ecological psychology and its standing as a research program, it can withstand some questioning of the assumptions on which its current practice is founded. This is especially the case if the questioning is aimed at firming up foundations rather than tearing down the house.

6.1 Gibson on Direct Perception and Information

Gibson's posthumous *magnum opus*, *The Ecological Approach to Visual Perception* (1979) is perhaps alone among books about perception in devoting nearly 50% of its pages to discussion of the nature of the environment that animals perceive. This half of the book is a description of Gibson's theory of the information available for vision, which goes hand-in-hand with his theory of visual perception. There are two main points to Gibson's theory of perception. First, Gibson disagreed with the tradition that took the purpose of visual perception to be the internal reconstruction of the three-dimensional environment from two-dimensional inputs. Instead, the function of perception is the guidance of adaptive action. Second, Gibson (1966, 1979) rejected classical views of perception in which perception results from the addition or processing of information in the mind to physically-caused sensation; that is, he rejected perception as mental gymnastics. This information processing way of understanding perception, Gibson thought, puts an unbridgeable gap in place between the mind (where the information is added, and the perception happens) and the world (where the merely physical light causally interacts with the retina). Instead, Gibson argued, perception is a direct—non-inferential, non-computational—process, in which information is gathered or picked-up in active exploration the environment.

Combined, these two theses give rise to Gibson's most well-known contribution, his theory of affordances (1979; see chapter 7 for a detailed story about affordances). If perception is direct, no information is added in the mind; if perception also guides behavior, the environment must contain sufficient information for the animal to guide its behavior. That is, the environment must contain information specifying opportunities for behavior. In other words, the environment must contain information specifying affordances. These views place significant constraints on the theory of information that Gibson can offer. First, because it is used in non-inferential perception, information must be both ubiquitous in the environment and largely unambiguous; second, because perception also guides behavior, the information in the environment must specify opportunities for behavior, which is to say it must specify affordances. Although the theory of information outlined in Gibson 1979 does meet these criteria quite nicely, it is spelled out in too plainspoken a manner to be convincing to most philosophically inclined readers.[2] I will try to do better here.

The first thing to know about what Gibson meant when he used the word 'information' is that he was not talking about information as described by Shannon and Weaver. ("The information for perception, unhappily, cannot be defined and measured as Claude Shannon's information can be." 1979, 243.) The best first pass at an understanding of what Gibson *did* mean by 'information' is his distinction between stimulation and stimulus information. To see the difference, consider standing in a uniformly bright, densely fog-filled room. In such a room, your retinal cells are stimulated. The light in the room enters your eye and excites the rods and cones. But there is no information carried by the light that stimulates your retina. This is the case because the uniform white light that converges on the eye from the various parts of the room and is focused by the eye's lens does not specify the structure of the room. So stimulation, the excitement of sensory cells, is not in itself information and is, therefore, not sufficient for perception. The differences between the normal environment and the fog-filled room are instructive. In the fog-filled room, the light that converges on any point that could be occupied by an observer's head and eyes has been scattered by the fog. Thus, when it reaches the observer it has not come directly from any surface in the room, and hence cannot inform the subject about the surfaces in the room. In the more typical, non-foggy situation, the light that reaches any point in the room has been reflected off the room's surfaces. The chemical makeup, texture, and overall shape of the surfaces off which the light reflects determine the characteristics of the light. Since surfaces are interfaces of substances with the air in the room, the nature of the surfaces is, in turn, determined by the substances that make them up. This set of facts is what allows the light that converges at any point to carry information about the substances in the environment. It also allows animals whose heads occupy the point to learn about its environment by sampling the light.[3]

This story allows us to understand what it is for light (or other energy) to carry information, but says nothing about what sort of thing information is. When Gibson and his followers claim that information is ubiquitous, are they saying that in addition to the substances, objects and energies in the room, there is extra stuff, the information? Yes and no. Yes: information is a real, unproblematic aspect of the environment. But as is evident from the quotation from Gibson above, information it is not a kind of measurable, quantifiable stuff that exists alongside the objects or substances in the environment. Instead, information is a relational feature of the environment. In particular, the light converging on some point of observation is in a particular relationship to the surfaces in the room, that of having bounced off those surfaces and passed through a relatively transparent medium before arriving at the point. The information in the light *just is* this relation between the light and the environment.[4]

A few quick points about this. First, note that information relation between the light and the surfaces does not hold in the case of a fog-filled room. So the light in this case bears no information about layout of the environment. Second, it is worth noting that this way of understanding information allows it to be ubiquitous in the environment. Light reflected from surfaces in the environment converges at every point in the environment. Third, the information in the environment is more or less complete: the light converging at every point has reflected off *all* of the non-obstructed surfaces. Fourth, and most importantly for Gibson's project, is that the light can contain information specifying affordances. To see this, a little needs to be said about affordances. (Much more will be said in Chapter 7.)

Affordances are opportunities for behavior. Because different animals have different abilities, affordances are relative to the behavioral abilities of the animals that perceive them. In some cases, these abilities are importantly related to an animal's height. To take just two examples, Warren (1984) has established a relationship between leg length and stair climbing affordances and Jiang and Mark (1994) have established a relationship between eye height and the perception of gap crossing affordances.[5] Given the relationship between height and some affordances, information about height is also (partial) information about affordances. Remember that at every point in the environment reflected light converges from the surfaces in the environment. Among these surfaces is the ground, so one relatively obvious source of information concerning height is the light reflected from the ground beneath the point of observation. Sedgewick (1973) points out a less obvious source of information: the horizon cuts across objects at a height that is equal to the height of the point of observation. That is, whenever light is reflected to some point in the environment from the horizon and also from some object between that point and the horizon, the light will contain information about the height of the point of observation relative to the height of the object. Of course, information about the height of a point of observation is also information about the height of an animal. So, at least for the types of affordances that have some relationship to an animal's height (reaching, stair climbing, gap crossing), there is information in the light about the affordances. More generally, this means that information in light is not just about the things the light bounces off. It is also information about the perceiver and the relation between the perceiver and the environment. Gibson put this point by saying that proprioception and exteroception imply one another.

We will look at affordances in detail in the next chapter. For now, the following are the key points of this brief description of Gibson's theory of the information available in the environment for perception.

1. Information for perception is not Shannon-Weaver information.
2. Ontologically speaking, information is a relation between energy in the environment (light, vibrations, etc.) and the substances and surfaces in the environment.
3. Along with the substances and surfaces of the environment, the energy in the environment also contains information about animals that perceive it and about what is afforded to these animals.
4. Because of (3), information can be used by animals to guide behavior directly. That is, information about affordances can guide behavior without mental gymnastics

6.2 The Turvey-Shaw-Mace Approach

Gibson's ecological theory of vision (1979) was intended as a direct response to the increasing dominance of computational theories of mind. Unsurprisingly, Gibson's ideas were not widely accepted by cognitive scientists upon their appearance. Indeed, as noted above, they were subjected to withering criticism from an establishment in psychology that was committed to understanding perception and cognition as mental gymnastics. The ecological approach was not helped by Gibson's writing style, which, though and highly readable, was often imprecise.

Enter Michael Turvey, Robert Shaw and William Mace. Along with a few colleagues, Turvey, Shaw and Mace wrote a series of papers outlining a detailed philosophical account of the ontology and epistemology of Gibson's ecological approach (Shaw and MacIntyre 1974; Mace 1977; Turvey 1977; Turvey and Shaw 1979; Shaw, Turvey and Mace 1980; Turvey, Shaw, Reed and Mace 1981[6]). The most complete and rigorous of these papers is Turvey et al's 1981 reply to criticism from Fodor and Pylyshyn, so I will focus my discussion of the Turvey-Shaw-Mace on this work.[7] The goal of this Turvey et al 1981, stated in the first sentence, is to provide a more precise explication of Gibson's work, specifically his claim that "there are ecological laws relating organisms to the affordances of the environment" (p. 237). There are four key notions here, which come in pairs: the first pair is affordance and effectivity; the second is ecological law and information. I will look at them in order, suppressing as much formalism as possible. On the Turvey-Shaw-Mace view, an object X affords an activity Y for an organism Z just in case there are dispositional properties of object X that are complemented by dispositional properties of organism Z, and the manifestation of those dispositional properties is the occurrence of activity Y. Conversely, an organism Z can effect the activity Y with respect to object X just in case there are dispositional properties of Z that are complemented by dispositional properties of object X, and the manifestation of those dispositional properties is the occurrence of activity Y. The idea here is that affordances, or opportunities for behavior, are tendencies of things in the environment to support particular behaviors and effectivities are abilities of animals to undertake those behaviors in the right circumstances. Thus, a copy of *Infinite Jest* has the affordance 'climbability' for mice in virtue of certain properties of the book (height, width, stability, etc) and of the mouse (muscle strength, flexibility, leg length, etc.); the mouse has the effectivity 'being-able-to-climb' in virtue of properties of the same properties of the mouse and the book. The dispositional affordance and effectivity complement one another in that the climbing-of-book-by-mouse occurs only when the climbability and the being-able-to-climb interact. This, according to the Turvey-Shaw-Mace view is what affordances and effectivities are.

To understand how organisms perceive and take advantage of affordance, and, in particular, how they do so directly, Turvey et al define information and natural law. As with affordances and effectivities, the definitions of information and ecological law interact. Ecological laws, according to the Turvey-Shaw-Mace view, are quite different than they are according to what they term the 'establishment/extensional analysis'. Most of the differences don't matter to us here, so I will focus on just one key point of ecological laws: their being bound to contexts. According to Turvey et al, ecological laws are defined only within settings and do not apply universally. Thus, the ecological laws relating to things in the niche of mice do not necessarily hold in outer space, or even in the niches of mackerel or fruit flies. So, instead of taking laws to be universal relationships between properties as the 'establishment/extensional analysis' does, Turvey et al say that properties-in-environments *specify*, or uniquely correspond to, other properties-in-environments. The most important ecological laws on the Turvey-Shaw-Mace view are those relating ambient energy to properties in the environment, e.g., those relating patterns in the optic array to affordances. Thus, in virtue of ecological laws, particular patterns of the ambient optic array specify the presence of affordances in particular environments. It is this specification that allows the arrays to *carry information* about the affordances: because there is a lawful connection between patterns in ambient energy and the properties specified by those patterns, organisms can learn, or be informed about, the properties by sensing the patterns. Of course, among the properties about which information is carried in the array are affordances.

Here's what we have so far: Ecological laws make it such that ambient arrays specify properties (including affordances) and this specification is what makes the arrays carriers of information. The presence of this kind of information underwrites direct perception. If the information required to guide behavior is available in the environment, then organisms can guide their behavior just by picking that information up. Ecological laws guarantee that if a particular pattern is present in the optic array in a mouse's niche, affordances for climbing by mice are also present. Hence perception of those properties can be direct. This view of direct perception is clearly represented by Shaw's principle of symmetry (Shaw and McIntyre 1974, Turvey 1990). We can represent the symmetry principle as follows. Let E ="The environment is the way it is", I = "The information is the way it is", and P = "Perception is the way it is". Also, let '>' stand for the logical relation of adjunction, a non-transitive conjunction that we can read as "specifies". Then, the symmetry principle is

$$[(E > I) \& (I > P)] \& [(P > I) \& (I > E)].$$

In English, this says that "That the environment is the way it is specificies that information is the way it is and that information is the way it is specifices that perception is the way it is, and that perception is the way it is specifies that the environment is the way it is and that information is the way it is specifies that the environment is the way it is." We can simplify this to say that the environment specifies the information, which specifies perception, and perception specifies the information, which specifies the environment. This principle is symmetrical in that the environment, information and perception determine one another. This, on the Turvey-Shaw-Mace view, is what it is for perception to be direct. By law, the environment determines the information, which determines the perception. This makes the perception a guarantee of the presence of the information and also of the environment. So direct perception is perception that, by ecological law, is guaranteed accurate.

6.2. Issues with the Turvey-Shaw-Mace Approach

The Turvey-Shaw-Mace approach is a sensible and faithful account of an epistemology and ontology to accompany Gibsonian ecological psychology. I think, though, that there are problems with the account. Over the last several years, I have developed an alternative ontological and epistemological background for ecological psychology, one that attempts to be equally faithful to Gibson's vision. I will restrict my comments here to differences concerning direct perception and information. I will have some critical comments about the Turvey-Shaw-Mace view of affordances in Chapter 7. The main problem with the Turvey-Shaw-Mace account of information is that, by insisting that information depends upon natural law, they have made it such that there is too little information available for direct perception. In particular, on the Turvey-Shaw-Mace view, there is no information about individuals, in social settings, or in natural language. I will discuss these in order.

On Individuals. Because Turvey, Shaw and Mace take direct perception to be infallible, they insist that it be underwritten by information, which is, in turn, underwritten by natural law.

They are careful to maintain that the laws in question are *ecological* laws, laws that hold only in particular niches. Thus, laws need not be universal in order to allow information to be carried in the environment. But, of course, ecological laws must still be general in that they apply to a variety of individuals. For example, there would be an ecological law that connects a particular optical structure, a visible texture, to the bark of a particular kind of tree: in the environment of grey squirrels, say, optical structure O is present only when light has reflected off a silver maple. Note that making the ecological law niche-specific makes it so that the presence of optical pattern O in other environments, where lighting conditions or tree species differ, doesn't affect O's information carrying in the squirrel's environment. So far so good, but in the each grey squirrel's environment there are a few trees that have special affordances in that, unlike most trees in the environment, they contain nests. There are no ecological laws relating these trees, as individuals, to properties of the optic array, so there is no information about these trees, as individuals, available to the squirrels. This, of course, does not apply only to trees. If information depends on laws, there is also no information about individual people available for perception. So although a human infant might have information available about humans, she has none about her mother. So, on the Turvey-Shaw-Mace view, either babies do not perceive their mothers (because the information for direct perception is unavailable) or they do not perceive them directly. I take it that either alternative is unacceptable to Radical Embodied Cognitive Scientists.

On Social and Linguistic Information. Another facet of the Turvey-Shaw-Mace requirement of law-like regularities for information to be present is that no information can be carried in virtue of conventions. Conventions hold, when they do, by public agreement or acquiescence and are thus easily violated. Because of an error at the factory or a practical joke a milk carton may not contain milk and a beer can may not contain beer. This is true in any context in which milk cartons and beer cans appear. Similarly, through ignorance or dishonesty spoken and written sentences can be false and words can be used to refer to non-standard objects. In fact, these things happen all the time even in the environments where the conventions in question are supposed to most strongly enforced, e.g., at the grocery store or Presidential press conferences. None of this is to imply that there is no information to be picked up at grocery stores or when the President speaks. Ecological laws determine the way that collections of aluminum cans in a cardboard box will structure fluorescent light and the way exhalations through vocal cords that pass by moving mouth, lips, tongue and teeth will structure the relatively still air. So there is information that there are cans on the shelf and that the President has said the he and the Prime Minister use the same toothpaste. But, because these things are merely conventionally determined and conventions may be violated, there is no information concerning the presence of beer or the President's toothpaste of choice. And since direct perception depends upon the presence of such information, we must, according to the Turvey-Shaw-Mace view, perceive that there is Bodingtons in the cans and that the President and Prime Minister use the same toothpaste either indirectly, or not at all.

Radical Embodied Cognitive Scientists require theories of information and direct perception that allow children to directly perceive their mothers and for beer cans to inform us about the presence of beer. This requires different accounts of what it is for perception to be direct and of the nature of information. Before presenting my alternative views of information and direct perception, I should point out that there is an active controversy in the ecological psychology community over what I'm calling the Turvey-Shaw-Mace view of information. In recent years, mounting empirical evidence gathered by ecological psychologists indicates that humans regularly use nonspecifying variables to perceive, in successful perception and in perceptual learning (Jacobs, Michaels and Runeson 2000; Jacobs, Runeson and Michaels 2001; Fajen 2005; Withagen and Michaels 2005; Jacobs and Michaels 2008; Withagen and Chemero 2008). But according ot the Turvey-Shaw-Mace view, a variable that does not specify (i.e., is not lawfully connected to) a particular environmental feature cannot carry information about that feature. There is mounting evidence, that is, that the Turvey-Shaw-Mace view of information is inadequate. So, even if you are unconvinced by the philosophical arguments I have offered against the Turvey-Shaw-Mace view, there are other compelling reasons to worry about it. Among those who have felt compelled to worry are Jacobs and Micheals (2008), who offer a theory of learning that attempts to rescue most of the Turvey-Shaw-Mace view. I am less confident that it is savable.

6.3. An Alternative Approach to Direct Perception

On the Turvey-Shaw-Mace approach, direct perception is defined as perception that is grounded in ecological law, so is always accurate. Indeed, Turvey et al 1981 define *perception itself* as direct and law-governed (245). As argued above, this rules out information about, and so direct perception of, individuals and things partly determined by convention. To make it possible for these things to be perceived directly, we need a different understanding of direct perception. In this section, I describe perception as direct when and only when it is non-inferential, where being non-inferential does not guarantee accuracy. Direct perception is perception that does not involve mental representations. This understanding of direct perception, I would argue, is what Gibson had in mind. For example, he writes "When I assert that perception of the environment is direct, I mean that it is not mediated by retinal pictures, neural pictures, or mental pictures." (Gibson, 1979, 147)

We can get started in seeing what this kind of direct perception is by looking at Brian Cantwell Smith's notions of *effective* and *non-effective tracking*, already described in Chapter 3. An outfielder effectively tracks a fly ball when the light reflecting off the ball makes contact with her eyes, and she moves her eyes and head so as to maintain that contact. In terms of the physics of the situation, the ball, the outfielder and the intervening medium are just one connected thing. In effective tracking, that is, the outfielder, the ball and the light reflected from the ball to the outfielder form a single coupled system. No explanatory purchase is gained by invoking representations here: in effective tracking, any internal parts of the agent that one might call mental representations are causally coupled with their targets. This effective tracking is direct perception. We can also have direct perception during *non-effective tracking*. Often an animal must continue to track an object despite disruption of causal connection. The outfielder, that is, must be able to continue to track the fly ball even when the light reflected from it is (temporarily) unavailable, as when her head turns directly past the low, late-afternoon sun. This non-effective tracking, though, also does not require mental representation. There are three reasons for this. First, non-effective tracking could be accomplished just by causal connection and momentum. The head's momentum keeps it going that way, and the light coming directly from the sun no longer overwhelms that reflecting off the ball. Second, as Gibson points out, perception is an activity, and as such happens over time. So directly perceiving something may involve periods of time when it is being tracked effectively and periods when it is tracked non-effectively. Third, and this is getting ahead of myself because I haven't said what information is yet, there is still information in the light about something that is temporarily occluded. Thus we can have direct, that is non-representational, perception even when tracking is non-effective.[8]

There are two relevant consequences of taking tracking as the model of direct perception. First, we can see that perception is, by definition, direct. Perception is always a matter of tracking something that is present in the environment. Because animals are coupled to the perceived when they track it, there is never need to call upon representations during tracking. Effective and non-effective tracking are non-representational, hence direct. Explaining how we write novels or plan vacations might require invoking something like a representation in the sense of strong decoupling described in Chapter 3. But perception never does.

The second consequence of taking tracking as the model of direct perception is that perception can be direct and mistaken. First, and perhaps obviously, when tracking is non-effective, it is possible for the animal to lose track of its object. The fox might stop behind the rock, yet the bird's head and eyes might keep moving along the path that the fox was following. This kind of minor error is typically easily corrected, of course. Another possibility is when an animal is coupled with an inappropriate object. For example, the same optical pattern can be caused by a full moon and a light bulb on a cloudy night. And there will be the same sort of continuous column of disturbance connecting a moth to each. So the moth will be effectively tracking whichever of the two it happens to be connected with. When the moth is effectively tracking the light bulb, it is making a mistake. But this does not mean that it is tracking the bulb via a mental representation of the moon. For if it did, then it would also be tracking the moon via a mental representation of the moon when it was doing things correctly and perception would never be direct. Instead, the moth is directly perceiving the moon or misperceiving the light bulb via a non-specifying optical variable (Withagen 2004; see also Withagen and Chemero 2008). A variable is non-specifying when its presence is not one-one correlated with some object in the environment. Like the moth when it is coupled with the moon, many animals rely on non-specifying variables. Yet according to the Turvey-Shaw-Mace view, non-specifying variables do not carry information about the environment, and so cannot be used for perception, direct or otherwise. So to make sense of the moth's effective coupling with the moon as a case of direct perception, we need a different theory of information, according to which non-specifying variables can carry information. The same is true if we want to understand my perception of beer-presence in beer cans and meanings in words.

6.4. An Alternative Approach to Information

There is a theory of information that has considerable currency in cognitive science that is consistent with Gibsonian information: Barwise and Perry's (1981, 1983) *situation semantics*, discussed briefly in Chapter 2, and the extensions of it by Israel and Perry (1990), Devlin (1991), and Barwise and Seligman (1997). Situation semantics is a good candidate here because Barwise and Perry's realism about information was directly influenced by Gibson. Barwise and Perry (1981, 1983) developed situation semantics in order to, as they said, bring ontology back to semantics. That is, they were interested in a semantics based on how the world is, and not on minds, knowledge, or mental representations. Information according to this view is a part of the natural world, there to be exploited by animals, though it exists whether or not any animals actually do exploit it. According to situation semantics, information exists in *situations*, which are roughly local, incomplete possible worlds. Suppose we have situation token s1 which of type S1 and situation token s2 which is of type S2. Then situation token s1 carries information about situation token s2 just in case there is some *constraint* linking the type S2 to the type S1. Constraints are connections between situation types. See Figure 1. To use the classic situation semantics example (Barwise and Perry 1983, Israel and Perry 1990, Barwise and Seligman 1994), consider the set of all situations of type X, in which there is an x-ray with a pattern of type P. Because patterns of type P on x-rays are caused by veterinarians taking x-rays of dogs with broken legs, there will be a constraint connecting situations of type X with situations of type D, those in which there is a dog with a broken leg that visits a veterinarian. Given this, the fact that a situation x is of type X carries the information that there is a situation d (possibly identical to x) of type D in which some dog has a broken leg. See Figure 2.

For our purposes here, there are two things to note about this example. First, the constraint between the situation types is doing all the work. That is, the information that exists in the environment exists because of the constraint, and for some animal to use the information the animal must be aware of the constraint. This feature is true not just of the example of the unfortunate dog, but holds generally of information in situation semantics. The second point is that the constraint in the example holds because of a causal regularity that holds among dog bones, x-ray machines and x-rays. That is, the particular x-ray bears the information about the particular dog's leg because, given the laws of nature and the way x-ray machines are designed, broken dog legs *cause* x-rays with patterns of type P. This feature of the example does *not* hold more generally of information in situation semantics. That is, constraints that hold between situation types are not just law-governed, causal connections. Constraints can hold because of natural laws, conventions, and other regularities. So, a situation with smoke of a particular type can bear information about the existence of fire by natural law, but it can also bear information about the decisions of tribal elders by conventions governing the semantics of smoke signals.

Even given this very sketchy description of the nature of information in situation semantics, we can see that this view of information can capture the kind of information that Gibson was interested in. We can see this via an example. Imagine that there is a beer can on a table in a room that is brightly lit from an overhead source. Light from the source will reflect off the beer can (some directly from the overhead source, some that has already been reflected of other surfaces in the room). At any point in the room at which there is an uninterrupted path from the beer can, there will be light that has reflected off the beer can. Because of the natural laws governing the reflection of light off surfaces of particular textures, colors and chemical makeup, the light at any such point will be structured in a very particular way by its having reflected off the beer can. In situation s1, the light at point p has structure a of type A. Given the laws just mentioned, there is a constraint connecting the situations with light-structure type A to the beer-can-present situations of type B. So, the light structure at point p contains information about situation about token beer-can-presence b (of type B). Notice too that, because of conventional constraints governing the relationship between cans and their contents, beer-can-presence b being of type B carries information about beer-presence c of type C. Furthermore, the light at some point in the room from which the beer can is visible will contain information about the beer can's affordances. Take some point p, which is at my eye height. The light structure available at this point will contain not just information about the beer can and the beer, but also about the distance the point is from the ground, the relationship between that distance and the distance the beer can is from the ground, hence the reachability of the beer can and drinkability of the beer for a person with eyes at that height.

Note the this example makes clear that on my view, but not Turvey-Shaw-Mace, constraints that connect situations are not limited to law-like connections but can also be cultural or conventional in nature, the fact that some situation token contains information about some other token does not necessarily entail that the second situation token is factual. For example, the light at my point of observation contains information about the beer can and the beer can contains information about beer being present. Even though it's possible that, because of some error at the brewery that caused the can to be filled with water, there is no beer in the can, the beer can presence can still carry information about beer presence. But according to Turvey-Shaw-Mace, the connection between the states of affairs must be governed by natural law. So according to the Turvey-Shaw-Mace view, beer can presences don't carry information about beer presences, because the beer can is not connected by natural law with the presence of beer. This is also a feature of Dretske's theory of information (1981) and has long been thought to be problematic.[9] Situation theorists have typically argued that constraints need not be law-like connections between situation types. Barwise and Seligman (1994, 1997) for example have argued that the regularities that allow the flow if information must be reliable, but must also allow for exceptions. Millikan (2000) makes a similar point. She distinguishes between informationL (information carried in virtue of natural law) and informationC (information carried in virtue of correlation). Because constraints need only be reliable, and not law-like, non-specifying variables can carry information. Millikan also makes a valuable point concerning just how reliable non-specifying variables need be. On her view, the correlation between two events need be just reliable that some animal can use it to guide its behavior. Thus information-carrying connections between variables can be fully-specifying, marginally significant, or anything in between, depending on the type of behavior that the variable provides information for.

This works well with the theory of what it is for perception to be direct, outlined in section 6.3 above. Remember that according to this view perception is direct when it is non-representational, the result of an informational coupling between perceiver and perceived. This says nothing about what kind of constraint allows the information to be available. Since the situation semantics theory of information allows information to be present with merely reliable constraints, constraints that hold only sometimes can underwrite direct perception. So we *can* directly perceive beer-presence, given beer-can presence, despite occasional mix-ups at the brewery. And we can directly perceive the meaning in the spoken sentences despite the fact that people lie or misspeak. Most importantly, I think, a child can directly perceive her mother, even though there are no laws of nature concerning individuals.

6.5. Compare and Contrast: On Specification and Symmetry

I have already said that on the views of information and direct perception outlined here, there is information about, and so the possibility of direct perception of, individuals and socially-, culturally-, and conventionally-determined entities and states of affairs. This is already a marked difference between the view I outline and the Turvey-Shaw-Mace view. Even more striking, and perhaps more troubling to some ecological psychologists, is the effect the views I have outlined have on Shaw's principle of symmetry. Remember that the principle of symmetry is that (1) the environment specifies the information available for perception and that the information available for perception specifies what is perceived and (2) what is perceived specifies the information available for perception and that the information available for perception specifies the environment. There are, in other words, 1:1 correspondences between the environment and the information available for perception and between the information available for perception and what is perceived. This principle is taken to be the most important part of the Turvey-Shaw-Mace view of information and direct perception. Indeed as was noted above, information and direct perception are defined in terms of it. On the view described here, however, symmetry does not hold. This is the case because on my situation-semantics-derived view, information does not depend on 1:1 correspondences. To repeat the example, on my view, there could be information about beer at my point of observation because light arriving there has been reflected off an unopened Bodington's can, despite the possibility that there is actually no beer because the can might be full of something else. In fact, according to the view I've outlined, there is an important asymmetry at work here. The asymmetry in question here is partly an asymmetry in what we might call direction of it. The environment to perception fit is, at least partly, causal, while the perception to environment fit is primarily normative. The can being the way it is causes the light to be the way it is at my point of observation, which causes me to perceive the beer on the table. But my perception, via the structure of the light, that there is beer in the refrigerator in no way causes there to be beer in the refrigerator. Instead, my perception fails, is incorrect, if there is no beer.

A second way the asymmetry of direction of fit shows up can be brought to light diagrammatically. In situation semantics, constraints connecting types of situations allow tokens of those types to carry information. So for example, because of various constraints concerning the way light reflects off surfaces, there are causal constraints connecting the type of situation in which my daughter is present to situations in which the optic array is structured in a particular way, and because of the way light interacts with me and my visual system, there will be constraints connecting these optical array structurings and my perception of my daughter. That is, constraint C1 connects Ava-present situation type E with Ava-array situation type A and constraint C2 connects Ava-array situation type A with Ava-perception situation type P. Constraints C1 and C2 are, of course, primarily causal. We can see this in the top part of Figure 3. This part of the figure, and this direction of fit from environment to perception, corresponds to the first part of the symmetry principle, E > I > P. In contrast, consider the lower part of Figure 3. This depicts the relationship among tokens: this particular Ava-perception token p of type P is informative about a particular Ava-array token a of type A which is, in turn, informative about a particular Ava-presence token e of type E. This reflects a truism of situation theory: information "flows" among tokens in virtue of constraints among types. This lower part of the diagram corresponds to the second part of the symmetry principle, P > I > E.

We can, then, see another way in which the different directions of fit are different: the environment to perception direction of fit is due to constraints among types and the perception to environment direction of fit is due to an informational relationship among tokens. On this view, Shaw and MacIntyre were right that there is a two-way informational relationship between perception and the environment, but they were wrong in thinking that both directions of the relationship are the same.

6.6 Information All Around

For Radical Embodied Cognitive Science to be convincing, more is needed than that ecological information can be coherently defined: it must be ubiquitously available for direct perception, and it must be information of a kind that can guide behavior without requiring mental gymnastics. In other words, it must be argued that the stimulus is not at all impoverished, that all the information required to guide behavior is available in the environment. To begin to a case for this, I will briefly discuss two different types of research on environmental information: optic flow and visual entropy. Before beginning, I should point out that each of these is a higher-order variable, which is to say that each is relational and takes time to perceive. Most of the variables that are of interest to ecological psychologists are higher-order. The guiding assumption is that perception is an activity involving orienting sensory organs, scanning, and the like and that activities take time. This means that perception is not just of simple quantities like mass, wavelength, position, etc., but also of comparatively complex relations, ratios, velocities, and accelerations. There is information available in the environment to perceive each of these properties directly. That is, given the temporal extendedness of the activity of perception one can simply *see*, e.g., how fast something is moving, without computing it.

6.6.1 Optic Flow and the Variable τ.

Many readers of this book will have seen the documentary film *Winged Migration*. One of the many, many wonderful things to be seen in this film is of direct relevance to us here. The film depicts diving gannets. Gannets are large sea birds that live in colder, northern coastlines, and are of interest to us because of the way they fish. Gannets are able to catch fish at much greater depths than other birds typically can, even pursuing them under water, because they dive down to the water from heights of around 100 feet (approximately 30 meters) and reach speeds of up to 60 miles (approximately 100 kilometers) per hour. Such a dive represents an extraordinary coordination problem. Diving gannets must keep their wings spread for as long as possible in order to maintain and adjust their heading toward a target fish in windy conditions. But hitting the water with spread wings would be catastrophic, at 60 mph wing bones would break. The question here is how gannets manage to retract their wings at the last possible moment, so as to hit the water at the right location and avoid injury. One possibility is that gannets perform a computation: using a stored representation of the expected size of prey fish, compute distance from the surface of the water; then compute time-to-contact with the surface from this distance, using internally represented laws of motion (mass, acceleration due to gravity, and friction are constants). This, it turns out, is not what gannets do. Gannets rely on *optic flow*, the patterns of motion available at the eyes of any moving observer.

The easiest way to understand optic flow is to remember what happens when one plays a first person video game. Moving your character around in its virtual environment causes a changing pattern on your monitor that, if the game is well-designed, gives you the sensation of actually moving around in the environment. This temporally extended onscreen pattern is a simulation of optic flow. Consider a familiar video game scenario: your virtual car is heading toward a fatal collision with, let's say, a brick wall.

1. As your car approaches the wall, the image of the wall on your monitor expands.
2. When you get close enough, individual bricks will become visible.
3. As you continue toward your virtual crash, the image of the wall will cover the entire monitor, and images of individual bricks will expand.

4. Getting closer to the wall, the images of the bricks will expand so that only a few of them are actually able to fit on the monitor, and they will appear textured.
5. Moving closer still, the images of the texture elements on the bricks will expand as well;
6. Then there is the loud crash noise and the cracked virtual windshield.

Back in the real world and less dramatically, the same phenomenon, called *looming*, happens constantly. As any animal moves about its environment, the images of objects or texture elements that the animal is moving towards will expand at the animal's eyes. This is often described by saying that optic flow is *centrifugal* in the direction of locomotion: texture elements radiate out from the center of your field of view as you move toward an object.[10]

Detecting centrifugal optic flow is very important, of course, but it is not sufficient to guide the gannet in drawing in its wings. David Lee (1980), however, demonstrated that properties of centrifugal optic flow can be sufficient to guide behavior by defining the higher order optical variable τ. τ is the ratio of the size of a projected image to the rate of change of the rate of change of the image's size. Using a little geometry and calculus, Lee showed that τ, a feature of the optic array available at the eye, is sufficient to guide the gannet's behavior without the use of internal computations. Imagine a situation as pictured in figure 6.3 in which we have a decreasing distance between an object in the world, such as a fish, and an animal's eye.[11] Suppose the distance between the eye and the object is changing at constant velocity V and that at time t the object is at distance z(t). At time t, the object will project an image of a size r(t) proportional to its size R, and as the distance between R and the animal decreases the projected image r(t)'s size will increase at velocity v(t). τ is the ratio of size of the image r(t) to rate of change of the size of the image v(t),

6.1 $\tau = r(t) / v(t)$.

Because the triangles on each side of the lens in figure 6.3 are similar (and using a little suppressed calculus), we know that the r(t)/v(t) is the same as the ratio of the objects distance z(t) and the rate at which it is moving toward the animal V. Thus, 6.2 $\tau = z(t) / V$.

If V is constant,

6.3 $\tau = z(t) / [z(t) / t]$,

which simplifies to

6.4 $\tau = t$.

So if V is constant, τ is equal to the time remaining until contact between the eye and the object.

There are several things here worth noting. First note that τ does not give information about the absolute distance of an object. Instead, it gives information about time-to-contact with the object, which is relevant to guiding movement. When trying to cross the street, how far away in meters an approaching car is matters much less than how soon it will you. Second, note that τ need not be computed by the gannet. It is available at the retina. τ, in other words, can be perceived directly. So, τ provides important information for the control of action in the environment, and it provides that information without requiring mental gymnastics. That is, sensitivity to the ratio of optical angle to the expansion of optical angle is sensitivity to the timing of approaching collision. Third, and most importantly, Lee and Reddish (1981) show that diving gannet's are sensitive to τ and use it to determine when to fold their wings. They filmed diving gannets and showed that the times of wing retraction is better predicted by the hypothesis that gannets are pick up information using τ than by the hypotheses that gannets compute time-to-contact or retract wings at some particular height or velocity. Finally, there is evidence that τ and τ-derived variables are used to undertake a variety of visually-guided actions. Indeed, Lee's lab alone has shown that τ is used by landing pigeons and hummingbirds, for hitting balls, somersaulting, long jumping, putting in golf and steering. (See Lee 2006 for an overview.)

6.6.3 Optic Flow and Information Processing

Optic flow has many other features than the sort of expansion in the direction of heading that is captured by τ, and these other features have seemed to many to call mental representations back into the picture. Consider walking toward a goal. Imagine that you are in a parking lot and want to walk toward your car. It would seem that you could use optic flow and the variable τ to do so by walking so that the center of visual expansion is your car. If the only variety of optic flow were this visual expansion, this would be a successful strategy. But in addition to walking toward the car, you will be moving your eyes. So in addition to the optical expansion, you will have rotational optic flow from moving your eyes and the overall optic flow will be the vector sum of two components: flow from your locomotion and flow from your eye movement. If centrifugal expansion of the object you're walking toward is just one component of your optic flow, it would seem that optic flow is insufficient to determine (and maintain) your direction of locomotion. In fact, it would seem that a mental computation would be necessary to subtract the effect of the eye movement on the information available for perception. This sort of worry is the motivation behind motor theories of perception (Grush 1997; Hurley 1998; Ebenholtz 2001; Mandik 2005), the idea in which is that in order to effectively subtract the optic flow generated by eye movements, one uses a mental representation of the eye movement. This representation, sometimes called an *efference copy* and sometimes called *extraretinal information*, can be used to generate a prediction of the optic flow that would be generated by the eye movement, which predicted optic flow can be subtracted from the actual optic flow, leaving behind the optic flow generated by heading. If this is correct, information available in the environment is not sufficient to guide you to your car (or any target); it must be supplemented by mental representations of your eye movements.

Do we need extraretinal information to subtract out optic flow from eye movements to control our locomotion? There is evidence that indicates that we do not. Warren and Hannon (1988, 1990; Warren 2003) performed a series of experiments to determine whether optic flow is sufficient to determine the direction of locomotion, or whether extraretinal information is required. Subjects watched a monitor displaying simulated optic flow, and were asked to determine the direction of locomotion. In these experiments there are two different kinds of optic flow simulated. In one case, the flow on the monitor simulates motion toward a target. In this case, subjects are also asked to track an object following a continuous path along the monitor. Thus, these subjects have optic flow generated by simulated locomotion and their own actual eye movement. In the other case, the flow on a monitor simulates both locomotion toward a target and optic flow generated by eye movements tracking an object on the monitor. So in the second case, the subjects have optic flow generated by simulated locomotion and simulated eye movement. In both cases, the optic flow is the same, but only in the first case (with a real eye movement) could there be any extraretinal information or efference copy. If extraretinal information is necessary for perceiving direction of locomotion (i.e., if optic flow is not sufficient), subjects with real eye movements should determine direction of heading much more accurately than subjects with simulated eye movements. In fact, however, both sets of subjects perceived direction of heading equally accurately, which indicates that the environmental information is sufficient and need not be supplemented by mental representations of eye movement. Indeed, many subjects with simulated eye movement reported experiencing illusory eye movements. This is a hint that our awareness of voluntary eye movements comes from the environment and not from internal representations of the movements. That is, perhaps we know what we're doing primarily by seeing ourselves do it.[12]

It seems, then, that we do not need mental gymnastics to use optic flow to tell the direction of our locomotion, but the preceding discussion does supply a sense in which perception involves information processing. The information available in the optical variable τ is only available to animals that are moving. Thus one might say, following Rowlands 2006, that sometimes animals process information by acting in the world. There are countless examples of this sort of information processing via activity, most of which are less exotic than τ. We turn our heads, changing the positions of our ears, to generate differences in the arrival times of sounds, and hence information about the direction of the sound. We lean when surveying a scene, and in so doing generate a motion parallax, and hence information about the distances of objects. And on and on. This is what Radical Embodied Cognitive Scientists mean when they claim that perception and action are tightly intertwined, and that perception is, in part, action. Action changes the information available to an animal's perceptual systems, and sometimes the action actually generates information. Thus there is a sense in which perception-action as studied by Radical Embodied Cognitive Scientists involves information processing, but it is a variety of information processing that does not involve mental gymnastics.

6.6.3 Detecting Entropy and Perceiving Sameness

Analogical reasoning has been of special interest in the cognitive sciences, at least in part because it is often taken to be the one uniquely human cognitive ability (e.g., by Lakoff and Johnson 1980, 1999). And, indeed, analogical reasoning is taken to require Olympic-level mental gymnastics. It is typically thought that for analogical reasoning to occur, there must be representations of a stored base situation and the current target situation (i.e., the situation to be reasoned about right now). The analogy itself is the represented relation between those two representations. So imagine that you have arrived at an unfamiliar airport, say Charles de Gaulle in Paris, and are interested in finding your luggage. First, you form a mental representation of the current airport, including representations of many of its features. You recall a representation of a familiar airport, say Philadelphia International Airport, one in which you know where the luggage carousel is. You then compare the representation of the familiar airport with that of the unfamiliar airport, putting all the relevant parts of the representations in correspondence.[13] Finally, you adapt the solution in the source representation to fit with the target representation. If the luggage carousel is downstairs of the terminal in Philadelphia, you look for it downstairs at de Gaulle. The difficult part in this, of course, is determining which represented source in memory has enough relevant similarities to the target. There are many sorts of similarities that are relevant. There can be similarities among attributes (both the car and the apple are red), similarities among relations (breakfast is before lunch and the primary is before the general election), and similarities among similarities among relations, and so on. Furthermore, in many cases, it is necessary to ignore lower-order similarities and differences among attributes to attend to higher-order similarities and differences among relations. Thus it would seem that analogical reasoning requires detailed mental representations and complicated procedures for retrieving and comparing them.

Although it does not bear out claims that humans *alone* are capable of analogical reasoning, research by Roger Thompson and colleagues on analogical reasoning in non-human primates to suggest that there is a "profound disparity" (Thompson and Oden 1996, 2000) between humans and chimpanzees on one hand and monkeys on the other. In a series of studies (Oden, Thompson, & Premack, 1988; Thompson & Oden, 2000; Thompson, Oden & Boysen, 1997; Washburn, Thompson & Oden, 1997), it was shown that humans and chimpanzees can match pairs of relations and that monkeys cannot. In the studies, adult humans and language-trained chimpanzees are shown to be able to match samples based upon the relations among the objects in the samples, while ignoring properties of the individual objects. That is, they would match a pair of quarters (relation = same) with a pair of nickels (relation = same), rather than with a quarter and a dime (relation = different). Furthermore, infant humans and chimpanzees are able to recognize sameness and difference. Capuchin monkeys could do neither. Thompson has used this data to argue that humans and chimps, but not monkeys, have the ability to form the higher-order representations required for analogical reasoning. This is the profound disparity, and it can be seen as giving some comfort to the proponent of Radical Embodied Cognitive Science. If only humans and language-trained chimpanzees are capable of matching relations between relations, perhaps only humans and language-trained chimpanzees form representations. A natural hypothesis to explain this is that there is something about learning a public language that imparts representational capacities that were otherwise not there, leaving most cognition of most animals a matter of interaction with their environments. This is the line that Andy Clark takes (1997, 2003). Experience recognizing and manipulating public, perceptually accessible symbols leads animals to have new capacities which clearly require representational explanation. These animals internalize the symbols, and learn to manipulate them internally in the same way that they did externally.

Things, alas, are more complicated. First, there is mounting evidence that the profound disparity does not hold up, so whatever processes are required in humans and language-trained chimps seem called for in other species. Second, it turns out that analogical reasoning does not require complicated representational processes: pigeons and baboons, at least, can perceive similarity among relations just by picking up information in a higher-order environmental variable. A series of experiments by Ed Wasserman and his colleagues has shown that both pigeons and baboons can perceive sameness and difference in arrays of icons (Young and Wasserman 1997, 2000; Fagot, Young and Wassermann 2001; Wasserman, Young and Cook 2004 is a review). Both the baboons and pigeons learned a relational matching task in which they were shown an array of 16 pictorial icons that are either all identical (16 pictures of an ice cream cone) or all different (one picture each of an ice cream cone, a bus, a football…), and asked to match them to either a different array of 16 identical icons or a different array of 16 different icons. By successfully matching an array of 16 ice cream cones to an array of 16 footballs, the pigeons and baboons show that they can ignore surface differences (ice cream cones vs. footballs) and match the arrays according to the relations among them. As Fagot, Young and Wasserman 2001 point out, successful matching is, in essence, analogical reasoning. The animals must use relevant similarities between two things to guide their behavior, while ignoring both irrelevant similarities and differences, and they must do so by attending to higher-order properties of the arrays (sameness or difference of the entire array of icons) rather than the surface features (the identities of the individual pictures in the array). This suggests that the profound disparity does not hold up, indicating that many animals are capable of analogical cognition.

What lesson should be drawn from the failure of the profound disparity? One possibility is that animals other than humans and language-trained chimpanzees can reason analogically because the mental gymnastics required for analogical reasoning are not the result of learning a public language. Another possibility is that reasoning analogically does not require mental gymnastics. The details of the experiments on pigeons and baboons indicate that the latter of these is the case. As just described, pigeons and baboons are quite capable of learning to match arrays of 16 icons based on relations. But as one gradually decreases the number of icons in the array from 16 to 15 to 14 and so on down to 2, the abilities of pigeons and baboons to correctly match arrays drops off, falling to near chance with arrays of 4 and fewer icons. This should be a surprise to those who assume that this sort of analogical matching requires representation of each of the icons in an array, so that they can be compared with one another to arrive at the representation of the relational property "all the same" or "all different" of the array, which represented relational properties must be stored for comparison with the represented relational properties of the other two icons before a response can be made. If this were the case, it should be more difficult to represent and make comparisons with larger arrays than with smaller ones because larger arrays will present greater computational loads. Yet larger arrays are easier for pigeons and baboons.

To explain this phenomenon, Young and Wasserman (1997) suggest that pigeons[14] are responding to the *entropy* in the arrays. As used here, entropy is an information-theoretic measure of disorder, calculated with this equation

6.5 $$H(A) = -\sum_{a \in A} p_a \log_2 p_a,$$

where A is variable, a is a possible value of that variable, and p_a is the proportion instances of a among observed values of the variable. For the non-mathematically-inclined, the key point here is that the maximum possible entropy of a variable increases as the number of bits in the signal increases; while the minimum possible is always 0. For example, when an array has 16 different icons, the proportion of any icon will be 1/16 = .0625, so

6.6 $H(A) = -.0625 \text{ X } \log_2 (.0625) \text{ X } 16 = 4$.

When an array has 2 different icons, the proportion of any icon will be ½ = .5, so

6.7 $H(A) = -.5 \text{ X } \log_2 (.5) \text{ X } 2 = 1$.

Because the $\log_2 (1) = 0$, the entropy of an array of identical items, no matter what size, will be 0. This explains why it is easier for pigeons and baboons to match samples based on sameness and difference when arrays are larger. In arrays of 16 icons, the animals must discriminate between entropy values of 0 (all icons the same) and 4 (all icons different), but with arrays of 2 icons, the animals must discriminate between entropy values 0 and 1. Pigeons and baboons, then, have a hard time with smaller arrays because the differences in entropy on which they make their discriminations are smaller. This accounts for the gradual decrease in performance as the number of icons in the array is reduced, an effect which is counter-intuitive if one assumes that the task requires that animals must explicitly represent and compute over the icons in each array to determine whether they are all the same or all different, and then match the results of those computations in order to act appropriately.

The upshot of this is that the higher-order variable entropy carries sufficient information for animals to perceive sameness and difference, and to engage in a variety of analogical reasoning, all without mental gymnastics. One might wonder, however, how it is that the higher-order variable entropy can be perceived directly. It is a logarithmic function, after all. Don't animals need to compute it? One way to find out is to use neural network simulations. If entropy can be detected without computations over representations, a neural network without hidden layers ought to be able to make discriminations of entropy levels. A mathematician or computer scientist would say that entropy cannot be detected by a 2-layer network. This is the case because, like XOR, entropy is not linearly separable. Indeed, with two icons, entropy logically equlvalent to XOR, and XOR famously requires hidden units. Thus, it might seem that attempting to use a 2-layer network to demonstrate the direct perception of entropy is a waste of time. The key to seeing that it is not a waste of time is to realize that, according to computer scientists, pigeons and baboons cannot make discriminations based on entropy. 'Being able to solve a problem' in computer science means being guaranteed to come up with the right answer every time. In contrast, in animal behavior, 'being able to solve a problem' means reliably coming up with the right answer at rates significantly greater than chance. So whether a neural network or an animal can solve a problem depends on what you mean by 'being able to solve the problem'. Clearly, for the purposes here, the animal behavior criterion is more appropriate.

The question, then, is can 2-layer neural networks make discriminations based on entropy at rates reliably significantly greater than chance. We have shown that they can (Silansky and Chemero 2002; Dotov and Chemero, 2006). Using MATLAB, we built a neural network with 64 inputs (in 16 sets of 4) and two output units. Each set of 4 input units was used to make a binary representation of an icon. Thus, if we wished to present the network an array of 16 identical icons, the inputs might be 16 instances of '0010'; if we wished to present 16 different icons to the network, each set of 4 would be different. Following the method of Young and Wasserman 1997, we trained the network to distinguish entropy = 0 (all icons identical) from entropy > 0 (at least one icon different from others) and to distinguish maximum entropy (all icons different from one another[15]) from other levels of entropy (at least 2 identical icons). We trained the network, first, with 16 icon arrays until further training did not produce improvements in performance. We then repeated this process, gradually reducing the number of icons until there were just two. Our results were qualitatively similar to the data found with pigeons and baboons. In particular, we found that the 2-layer network could discriminate entropy levels quite reliably with arrays of 16 icons and that it's performance deteriorated gradually as one reduced the number of icons, going to chance and then fluctuating wildly with arrays of 5 and fewer icons. See figure 6.6.

The simulation results suggest very strongly that pigeons and baboons perceive sameness and difference by directly perceiving entropy. They do show definitively that it is possible to achieve performance that is qualitatively very similar to that exhibited by pigeons and baboons without manipulating representations. They show, that is, that information about sameness and difference in the form of the higher-level variable entropy is available and is sufficient to guide behavior without processing.

6.7 Wrap Up

The purpose of this chapter has been to begin to outline a Gibsonian theory of perception and cognition to serve as a background theory for Radical Embodied Cognitive Science. So far, I've given a theory of what it is for perception to be direct, and provided a little evidence suggesting that perception might actually be direct. Direct perception is the non-representational use of information in the guidance of behavior. Suggesting that perception is direct involved saying what information is, showing that there's plenty of it around for animals to use, and showing that animals actually do use it. So far so good. But from the point of view of the Radical Embodied Cognitive Scientist, the most important information is information about affordances, and I haven't yet said much about what affordances are. This happens in chapter 7.

Figure 2

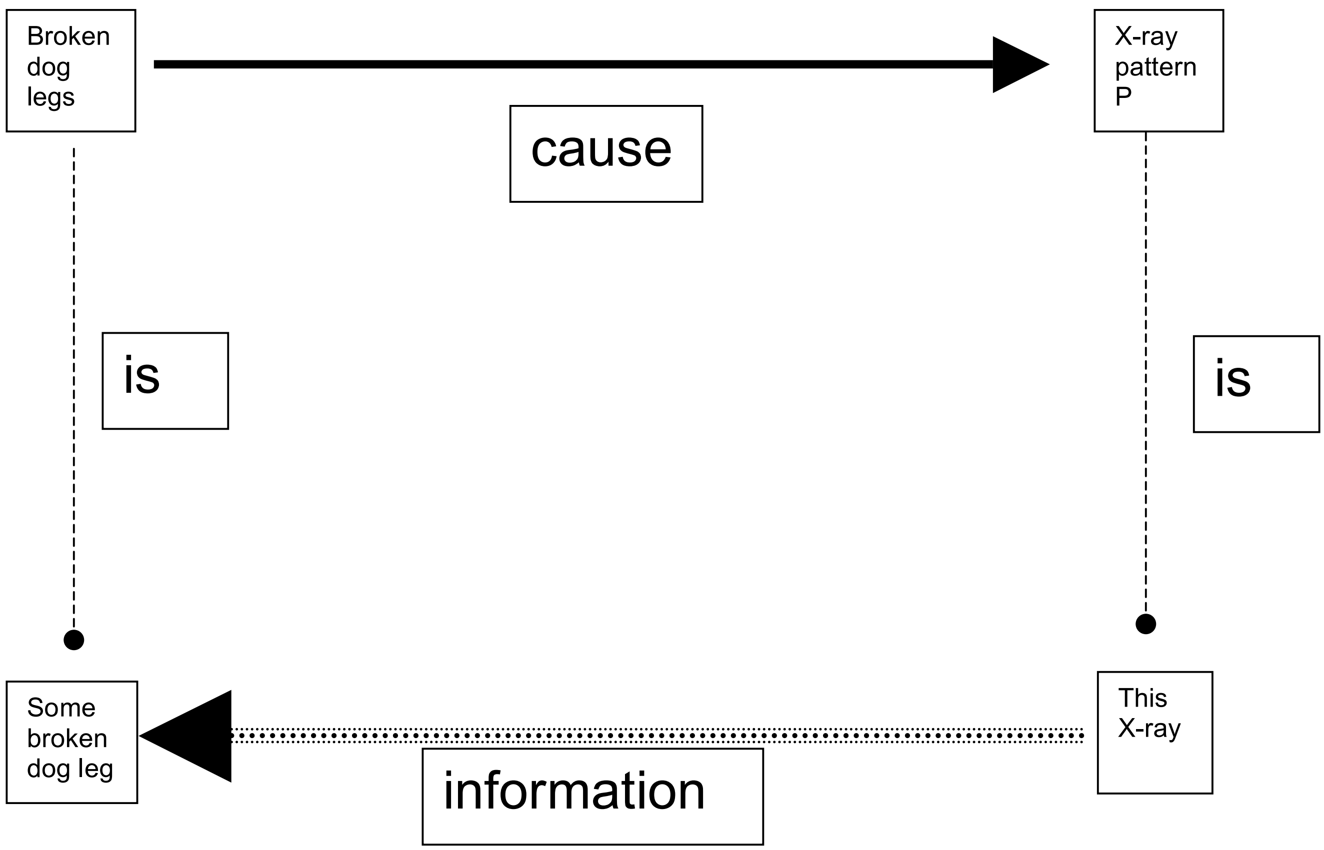

Broken dog legs
cause
X-ray pattern P
is
is
Some broken dog leg
information
This X-ray

Figure 1

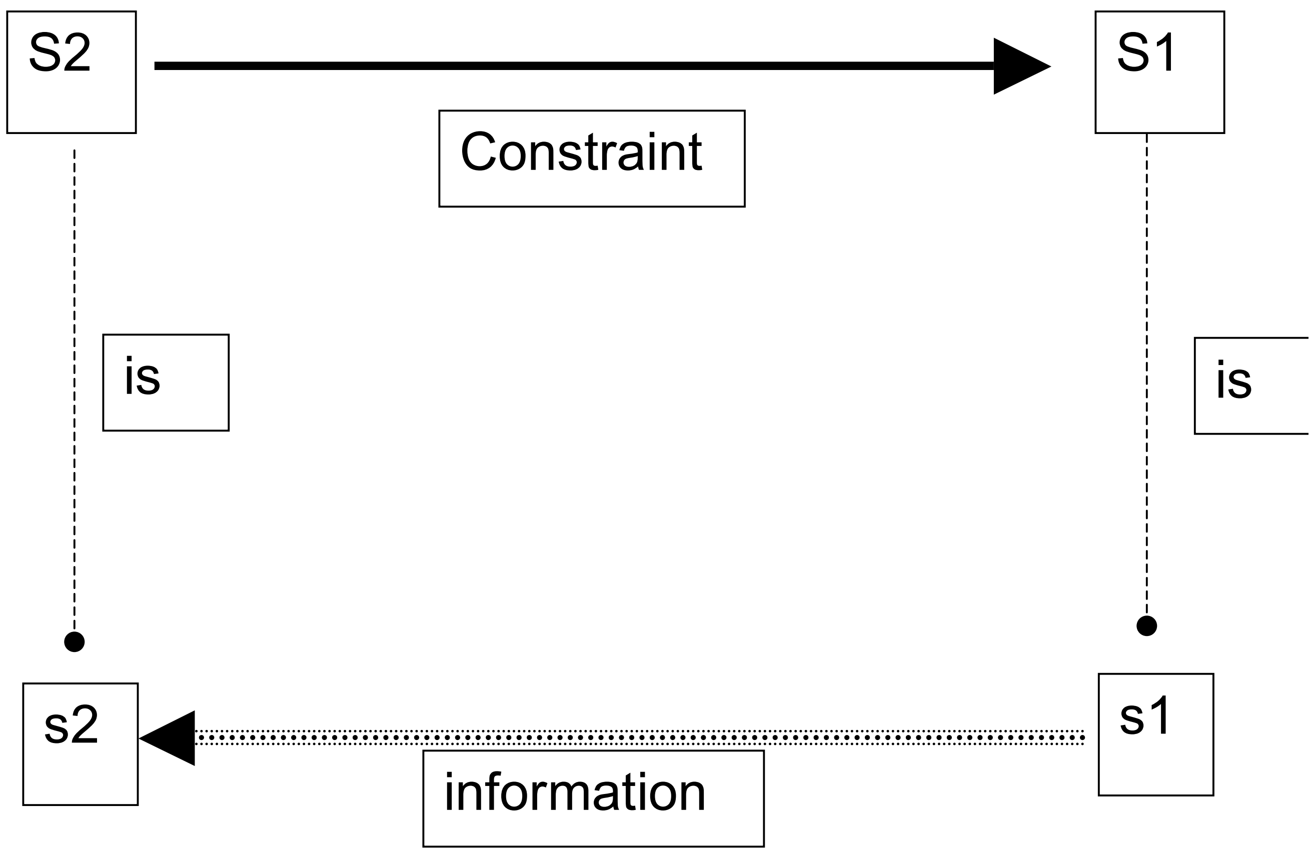

Figure 3. The top part of the diagram is analogous to Shaw's E > I > P; the bottom is analogous to his P > I > E.

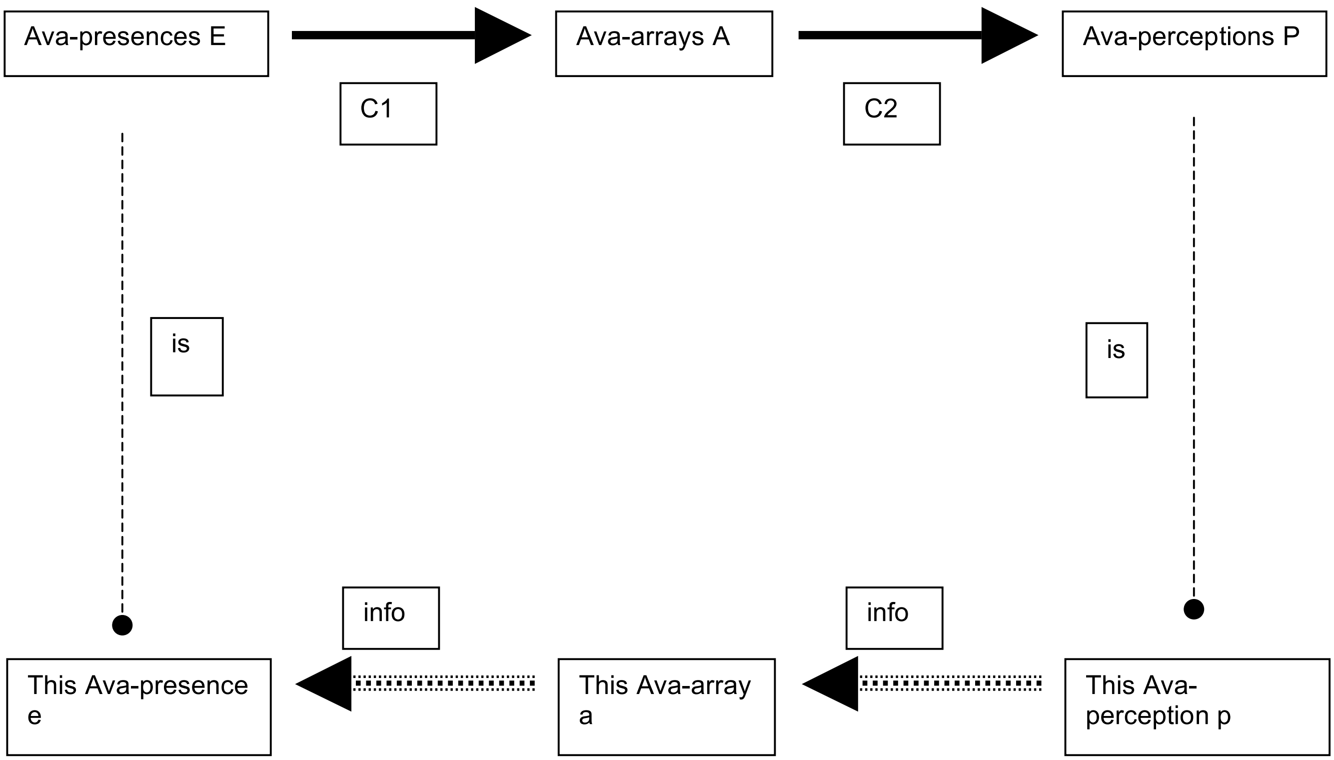

Figure 6.4

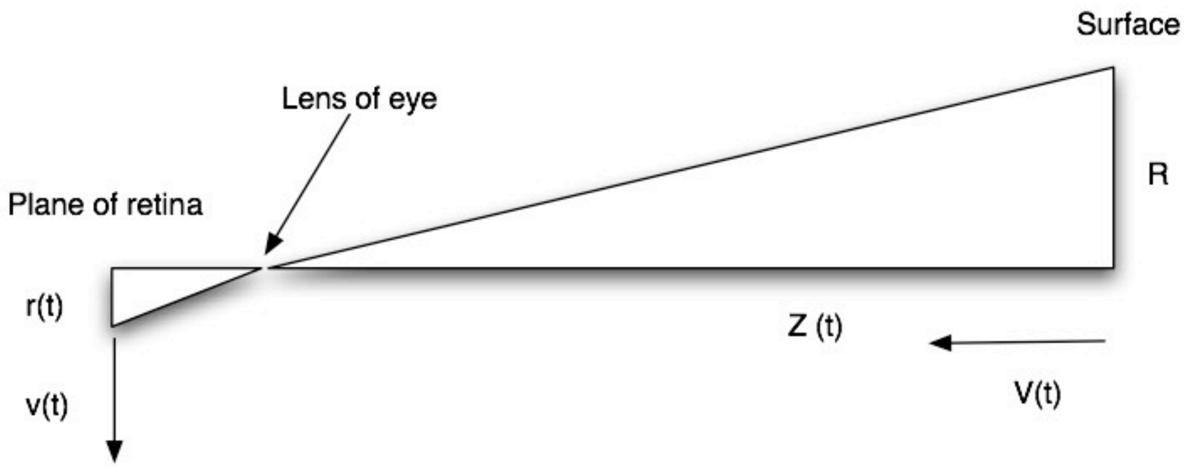

Figure 6.5

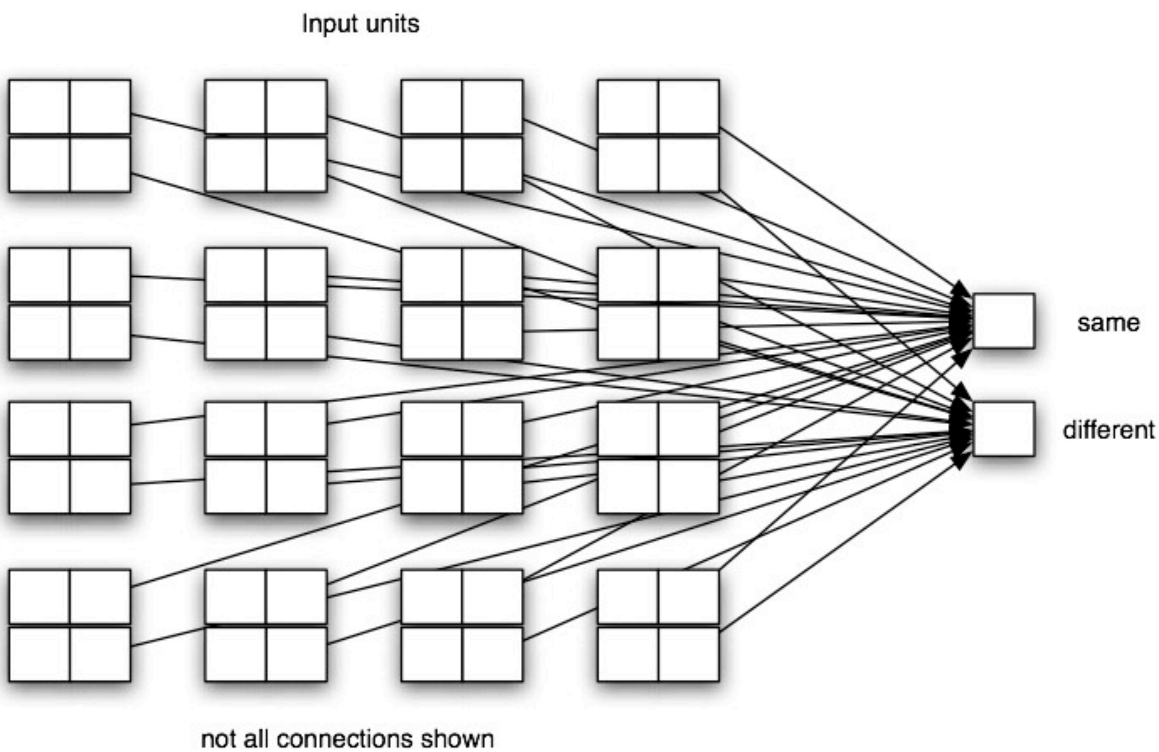

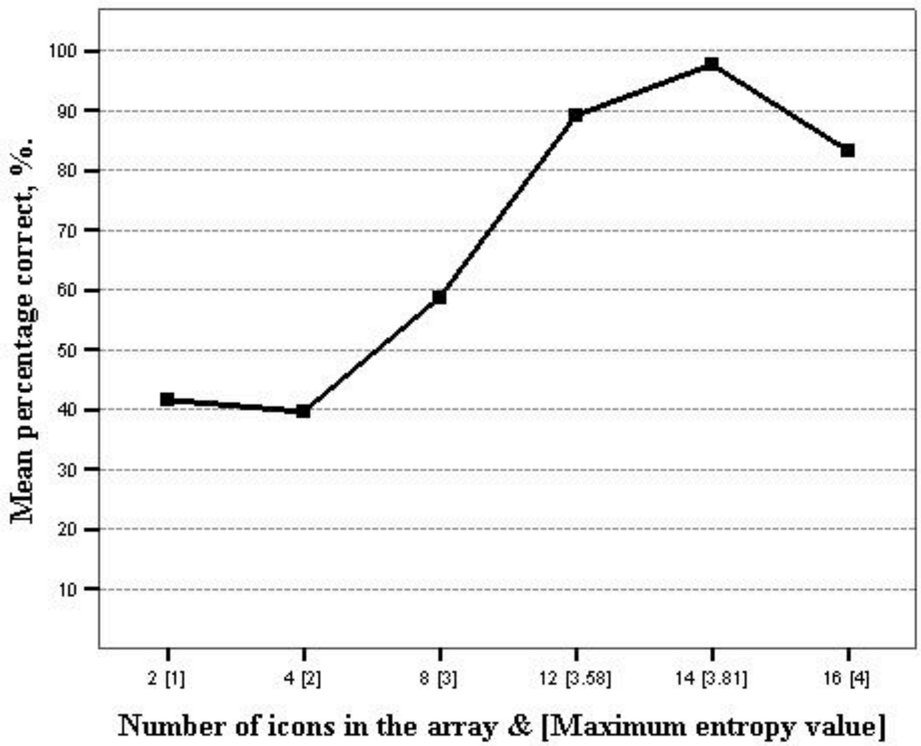

Figure 6.6 Mean percent correct entropy discriminations by six 2-layer artificial neural networks as a function of number of icons in the array. Thanks Dobri Dotov

---

[1] I should also point out that I owe them a personal debt. Though I was never formally a student of Shaw, Turvey, or Mace, each has been patient corrector of my misinterpretations and has even encouraged me in the development of my competing views. They still think I'm wrong.

[2] Reed 1996 also tries to give a philosophically-sound account of Gibson's theory of information and affordances. This chapter is an attempt to improve on Reed's work.

[3] Note that there is still stimulation of retinal cells in this case. Stimulation is necessary, but not sufficient for perception.

[4] Fodor and Pylyshyn (1981) agree with this point about the relational nature of information as Gibson understands it. They disagree with more or less everything else in Gibson 1979.

[5] The exact nature of the relationship between height and other aspects of body scale and affordances is a matter of dispute. See Chemero 2003 and Chapter 7 below.

[6] A quick note on Edward Reed: Although Reed was an author on the paper on cognition and spent his career working on a philosophically-sound version of Gibson's ecological psychology, I think it makes more sense to speak of the Turvey-Shaw-Mace view and not the 'Turvey-Shaw-Reed-Mace view'. This is because after working on the 1981 paper, Reed developed views that diverged both from that presented in the 1981 paper and from the one I'm presenting here. See Withagen and Chemero 200X

[7] Warren, a student of Shaw and Turvey, spells out the Turvey-Shaw-Mace view clearly and in detail in his "Direct Perception: The view from here" (2007).

[8] I should point out that there are some who would argue that there are mental representations involved, even in effective tracking. Chapters 3, 4 and 5, I hope, show that this is unhelpful.

[9] Note that everything said here about Turvey-Shaw-Mace is also true of Dretske's classic probability-based theory of information (1981).

[10] Though it is not directly relevant here, it is worth pointing out that optic flow is centripetal in the direction you are moving away from. Think about looking out the back of a moving car: the image size of objects decreases as you move away from them, which is to say that texture elements move toward the center of your field of view.

[11] This demonstration and the accompanying figure are based on Bruce et al (19XX) and differ slightly from Lee 1980.

[12] It must be noted that there is also evidence that subjects fail to judge heading accurately when the simulated eye movements are comparatively large. This would seem to indicate that extraretinal information is necessary, at least for large eye movements. But in the sort of simulations used in these experiments, optic flow with large, simulated eye movements is consistent with what one would see if one were following a curved path without eye movements. The "errors" subjects make in perceiving heading with simulated eye movements are consistent with correctly perceiving following a curved path, and subjects in experiments with simulated large eye movements often report that they are moving along a curved path. (See Warren 2003 for a review.) Thus it seems that there are two kinds of responses to optic flow with simulated eye movements: when the simulated eye movements are small, subjects perceive that they are following a straight path and that they've moved their eyes; when the simulated eye movements are large, subjects perceive that they are following a curved path and have not moved their eyes. Neither response indicates that extraretinal information or efference copy is necessary.

[13] The word 'relevant' in this sentence (well, not this very sentence, but the one to which this footnote is attached) should make you worried about the frame problem. If it does not, read Dennett 1987 immediately. If you think connectionist networks are the right way to solve the problem, read Haselager 1997 or Haselager and Van Rappard 1998.

[14] Fagot, Wasserman and Young 2001 use entropy to explain similar performances by baboons.

[15] The actual value of the maximum entropy varies depending on the number of icons in the array.