III Ecological Psychology

The rules that govern behavior are not like laws enforced by an authority or decisions made by a commander: behavior is regular without being regulated. The question is how this can be.

—James J. Gibson, *The Ecological Approach to Visual Perception* (1979)

6 Information and Direct Perception

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 high-profile attacks from establishment cognitive scientists (e.g., Fodor and Pylyshyn 1981). But thanks to Turvey, Shaw, and Mace's work as theorists and, perhaps more important, 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 what I take to be improved versions of each of the four main components of it. In this chapter, I will describe theories of information[1] 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.[2] 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 percent 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, places an unbridgeable gap 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—noninferential, noncomputational—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 (Gibson 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 that specifies opportunities for behavior. In other words, the environment must contain information that specifies affordances. These views place significant constraints on the theory of information that Gibson can offer. First, because it is used in noninferential 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.[3] 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," Gibson 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 not, therefore, 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, nonfoggy 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 their environment by sampling the light.[4]

This story allows us to understand what it is for light (or other energy) to carry information, but it 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 information 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.[5]

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 nonobstructed surfaces. Fourth, and most important for Gibson's project, is that the light can contain information that specifies 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 significantly 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.[6] 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 (Gibson 1979) was intended as a 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 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 McIntyre 1974; Mace 1977; Turvey 1977; Turvey and Shaw 1979; Shaw, Turvey, and Mace 1982; Turvey, Shaw, Reed, and Mace 1981[7]). 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 view on this work.[8] The goal of 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" (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 dispositions of things in the environment to support particular behaviors, and effectivities are dispositions 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 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 affordances, 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, for example, those relating patterns in the the light of 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. Crucially, 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 1990a). 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 nontransitive 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 the environment is the way it is specifies that information is the way it is and that information is the way it is specifies that perception is the way it is, and that perception is the way it is specifies that the information 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 lawful guarantee of the presence of the information and also of the environment. So direct perception is perception that, by ecological law, is guaranteed accurate.

6.3 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 on 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 gray 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 each gray 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, ecological or otherwise, 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 lawlike 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 thus are 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 nonstandard objects. In fact, these things happen all the time even in the environments where the conventions in question are supposed to be most strongly enforced, for example, 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 comparatively still air. So there is information that there are cans on the shelf and that the president has said that 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 on the presence of such information, we must, according to the Turvey-Shaw-Mace view, perceive that there is Boddingtons 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 2007; Withagen and Chemero 2009). But according to 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 Michaels (2007), 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.

6.4 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, 245) define *perception itself* as direct and law-governed. 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 noninferential, where being noninferential 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 recalling Brian Cantwell Smith's notions of *effective* and *noneffective 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 *noneffective 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 noneffective tracking, though, also does not require mental representation. There are three reasons for this. First, noneffective 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 noneffectively. 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 nonrepresentational, perception even when tracking is noneffective.[9]

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 noneffective tracking are nonrepresentational, 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 noneffective, 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 lightbulb 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 lightbulb, 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 lightbulb via a nonspecifying optical variable (Withagen 2004; see also Withagen and Chemero 2009). A variable is nonspecifying 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 nonspecifying variables. Yet according to the Turvey-Shaw-Mace view, nonspecifying 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 nonspecifying 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.

6.5 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 6.1. 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,

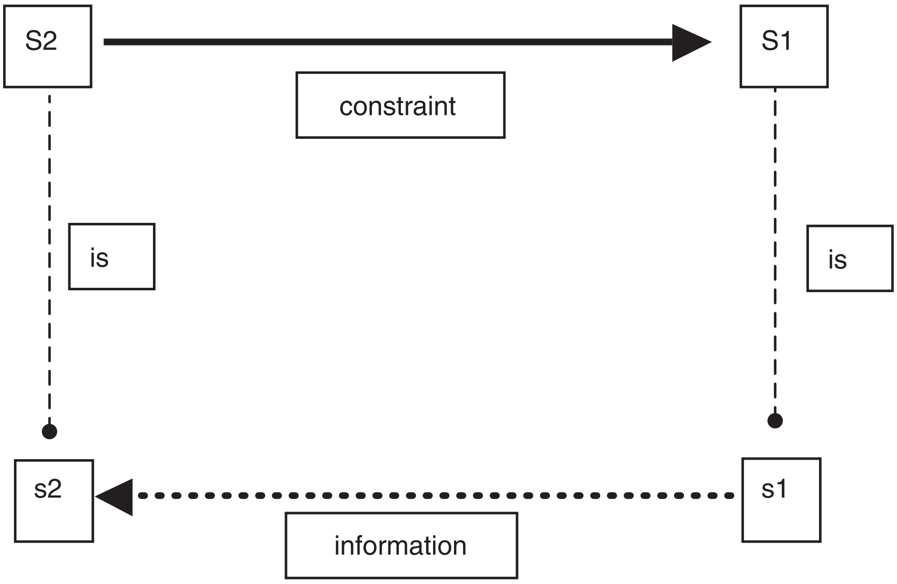

Figure 6.1
The information relationship. Lowercase *s*1 and *s*2 are tokens of capitalized types *S*1 and *S*2, respectively.

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 6.2.

For our purposes, 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.[10] 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 between situation types can hold in virtue of law-governed, causal connections, but they can also hold in virtue of customs, 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.

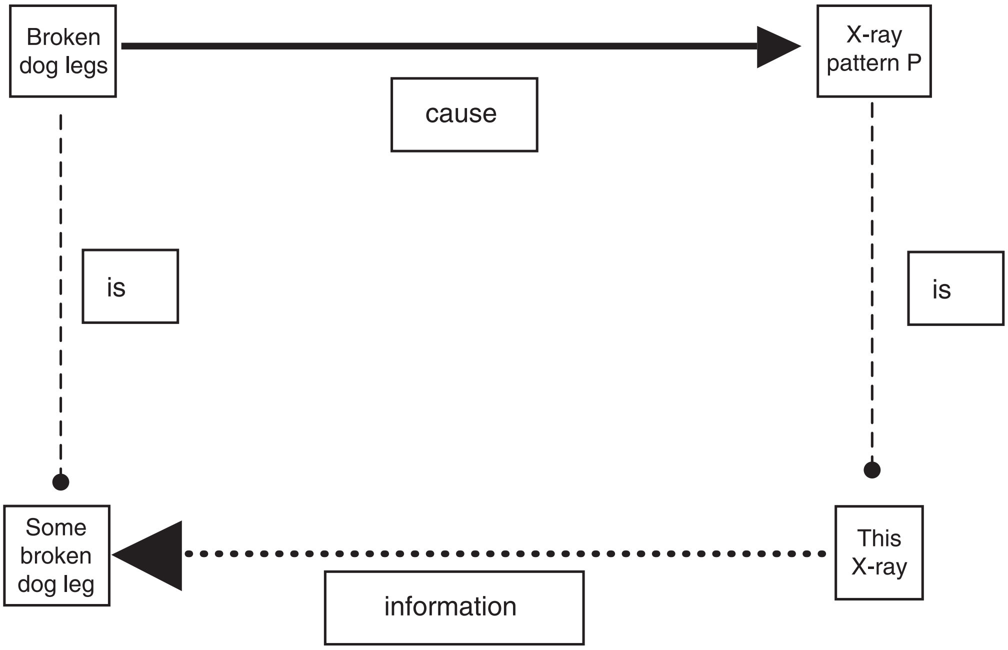

Figure 6.2
Information carried by an x-ray. *This X-ray* is a token of type *X-ray pattern P*; *Some broken dog leg* is a token of type *Broken dog legs*.

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 off other surfaces in the room). At any point in the room to 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 *s*1, 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 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, and hence the reachability of the beer can and drinkability of the beer for a person with eyes at that height.

Note that this example makes clear that on my view, but not Turvey-Shaw-Mace, constraints that connect situations are not limited to lawlike 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's presence can still carry information about the presence of beer. 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 (Dretske 1981) and has long been thought to be problematic. Situation theorists have typically argued that constraints need not be lawlike connections between situation types. Barwise and Seligman (1994, 1997), for example, have argued that the regularities that allow the flow of 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 lawlike, nonspecifying variables can carry information. Millikan also makes a valuable point concerning just how reliable nonspecifying variables need be. On her view, the correlation between two events need be just reliable enough 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 nonrepresentational, 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 important, I think, a child can directly perceive her mother, even though there are no laws of nature concerning individuals.

6.6 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 the information available for perception specifies what is perceived, and (2) what is perceived specifies the information available for perception and 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 Boddington'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 fit. The environment-to-perception fit is at least partly causal, whereas 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

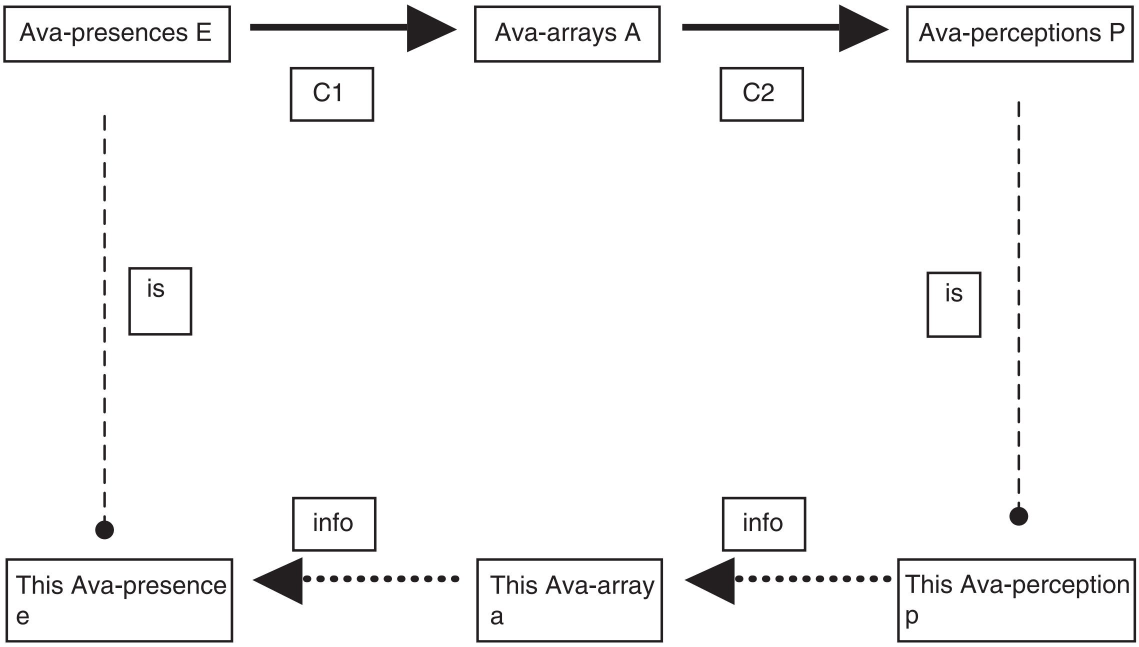

Figure 6.3

Information flow when my daughter is present. *This Ava-presence e*, *This Ava-array a*, and *This Ava-perception p* are tokens of types *Ava-presences E*, *Ava-arrays A*, and *Ava-perceptions P*, respectively. The top part of the diagram is analogous to Shaw's E > I > P; the bottom is analogous to his P > I > E.

can see this in the top part of figure 6.3. 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 6.3. 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 McIntyre 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.

6.7 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 make 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 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, and so on, 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*, for example, how fast something is moving, without computing it.

6.7.1 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, large sea birds that live along colder, northern coastlines, 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 thirty meters) and reach speeds of up to sixty 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 sixty miles per hour, 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. As you get 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 toward 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.[11]

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; Lee and Reddish 1981), 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 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 6.4 in which we have a decreasing distance between an object in the world, such as a fish, and an animal's eye.[12] 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 $v(t)$,

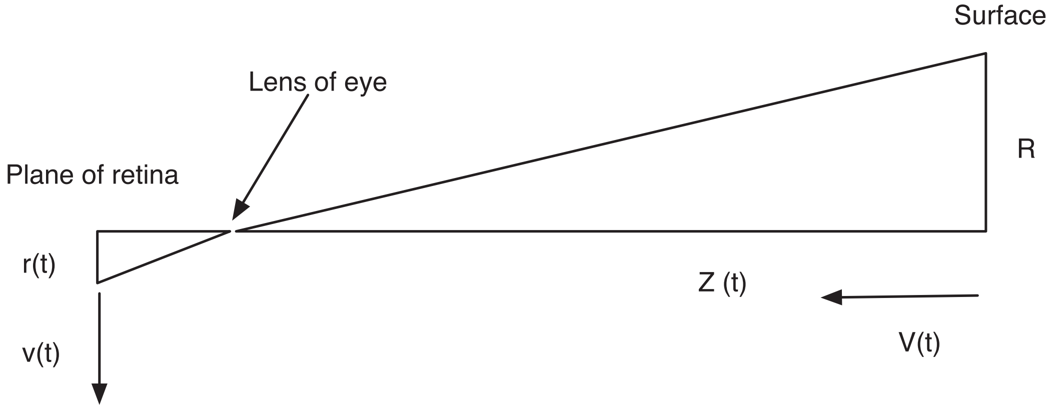

Figure 6.4
The optical expansion at the retina of the image of projected by object *R* as it moves toward the eye.

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

Because the triangles on each side of the lens in figure 6.4 are similar (and using a little suppressed calculus), we know that $r(t)/v(t)$ 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,

$$\tau = z(t)/V. \tag{6.2}$$

If V is constant,

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

which simplifies to

$$\tau = t. \tag{6.4}$$

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 you're trying to cross the street, how far away in meters an approaching car is matters much less than how soon it will hit 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 important, Lee and Reddish (1981) show that diving gannets are sensitive to τ and use it to determine when to fold their wings. They filmed diving gannets and showed that the time of wing retraction is better predicted by the hypothesis that gannets 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, and by humans hitting balls, somersaulting, long jumping, putting in golf, and steering. (See Lee 2006 for an overview.)

6.7.2 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 destination. 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 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; Warren 2004) 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 two different kinds of optic flow are 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 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.[13]

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.[14]

6.7.3 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 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.[15] Finally, you adapt the solution in the source representation to fit with the target representation. If the luggage carousel is downstairs at 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 nonhuman primates to suggest that there is a "profound disparity" (Thompson and Oden 2000) between humans and chimpanzees on one hand and monkeys on the other. In a series of studies (Oden, Thompson, and Premack 1990; Thompson and Oden 2000; Thompson, Oden, and Boysen 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 on 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 (1997, 2003, 2008) takes. Experience recognizing and manipulating public, perceptually accessible symbols leads animals to have new capacities that 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, Wasserman, and Young 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 sixteen pictorial icons that are either all identical (sixteen 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 sixteen identical icons or a different array of sixteen different icons. By successfully matching an array of sixteen ice cream cones to an array of sixteen 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 apparent 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 sixteen icons based on relations. But as one gradually decreases the number of icons in the array from sixteen to fifteen to fourteen and so on down to two, the ability of pigeons and baboons to correctly match arrays drops off, falling to near chance with arrays of four 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[16] are responding to the *entropy* in the arrays. As used here, entropy is an information-theoretic measure of disorder, calculated with this equation:

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

where A is a 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 sixteen different icons, the proportion of any icon will be $1/16 = .0625$, so

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

When an array has two different icons, the proportion of any icon will be $1/2 = .5$, so

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

Because the $\log_2 (1) = 0$, the entropy of an array of identical items, no matter what size, will be zero. 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 sixteen icons, the animals must discriminate between entropy values of zero (all icons the same) and four (all icons different), but with arrays of two icons, the animals must discriminate between entropy values zero and one. 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 that is counterintuitive 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 between entropy levels. A mathematician or computer scientist would say that entropy cannot be detected by a two-layer network. This is the case because, like XOR, entropy is not linearly separable. Indeed, with two icons, entropy is logically equivalent to XOR, and XOR famously requires hidden units. Thus, it might seem that attempting to use a two-layer network to demonstrate the direct perception of entropy is a waste of time. The key to seeing that it might not be 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

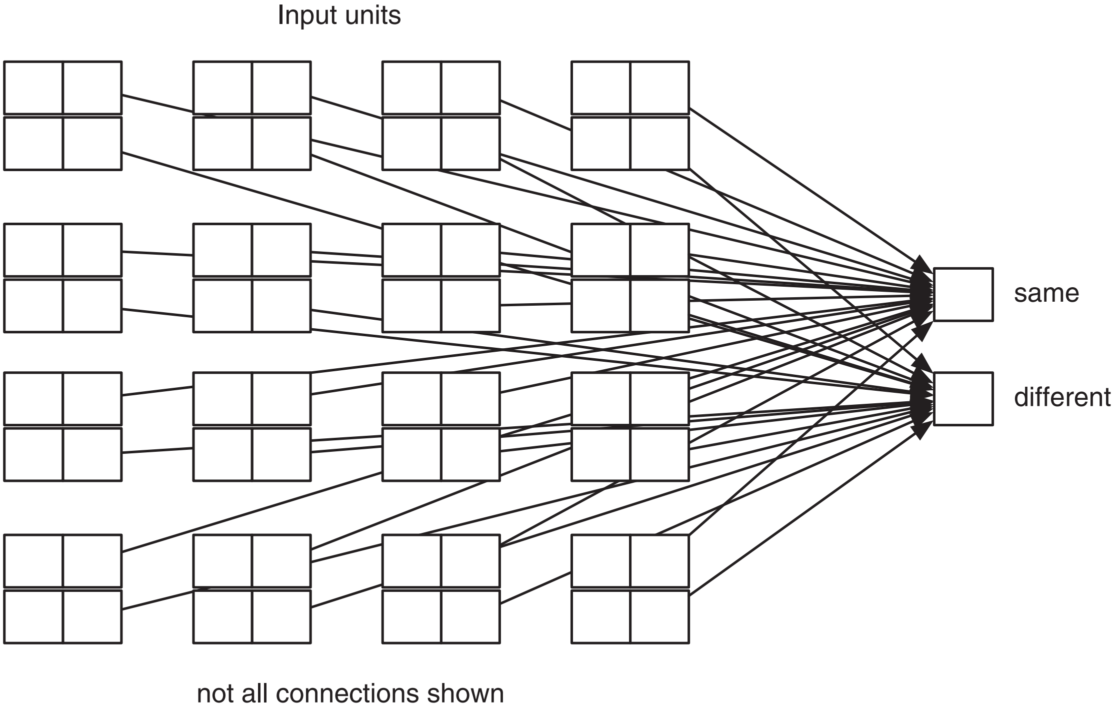

Figure 6.5

solve the problem." Clearly, for the purposes here, the animal behavior criterion is more appropriate.

The question, then, is can two-layer neural networks reliably make discriminations based on entropy at rates 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 sixty-four inputs (in sixteen sets of four) and two output units. See figure 6.5. Each set of four input units was used to make a binary representation of an icon. Thus, if we wished to present the network an array of sixteen identical icons, the inputs might be sixteen instances of "0010"; if we wished to present sixteen different icons to the network, each set of four would be different. Following the method of Young and Wasserman (1997), we trained the network to distinguish $\text{entropy} = 0$ (all icons identical) from $\text{entropy} > 0$ (at least one icon different from others) and to distinguish maximum entropy (all icons different from one another[17]) from other levels of entropy (at least two identical icons). We trained the network, first, with sixteen 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

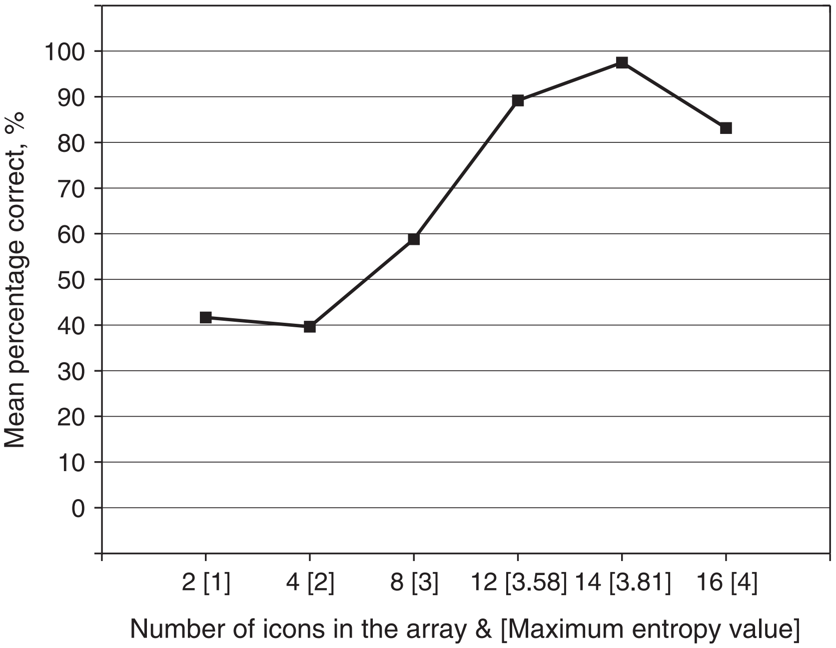

Figure 6.6
Mean percent correct entropy discriminations by six two-layer artificial neural networks as a function of number of icons in the array. Thanks to Dobri Dotov.

the two-layer network could discriminate entropy levels quite reliably with arrays of sixteen icons and that its performance deteriorated gradually as we reduced the number of icons, going to chance and then fluctuating wildly with arrays of five 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. I would argue that they 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 higher-level variable entropy is available and is sufficient to guide behavior without mental gymnastics.[18]

6.8 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 nonrepresentational 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.

7 Affordances, etc.

7.1 Direct Perception and Ontology

For radical embodied cognitive science to eschew mental representations, it must take perception to be direct, to be the pickup of information from the environment. Furthermore, animals must be able to use that information to guide action without complex processing, without mental gymnastics. This requires that perception be of affordances, or opportunities for behavior. Animals, that is, must be able to perceive what they can do directly. In the previous chapter, I explained how perception might be direct and gave a theory of the information available for perception. So far, though, I have said nothing about perceptual content, nothing, that is, about what animals actually perceive. This is where affordances come in. Following Gibson, I will maintain that animals perceive affordances directly. This leads to some ontological funny business. To see this, consider that the primary difference between direct and inferential theories of perception concerns the location of perceptual content. In inferential theories of perception, these meanings arise inside animals, based on their interactions with the physical environment. Light, for example, bumps into receptors causing a sensation. The animal (or rather its brain) performs inferences on the sensation, yielding a meaningful perception. In direct theories of perception, on the other hand, meaning is in the environment, and perception does not depend on meaning-conferring inferences. Instead the animal simply gathers information from a meaning-laden environment. The environment is meaning laden in that it contains affordances, and affordances are meaningful to animals. But if the environment contains meanings, then it cannot be merely physical. This places a heavy theoretical burden on radical embodied cognitive science, a burden so severe that it may outweigh all the advantages to conceiving perception as direct. Radical embodied cognitive science requires a new ontology, one that is at odds with today's physicalist, reductionist consensus that says the world just is the physical world, full stop. Without a coherent understanding of what the world is like, such that it can contain meanings and is not merely physical, direct perception is simply indefensible. Thus, like earlier theories that take perception to be direct (e.g., James 1912/1976; Heidegger 1927), Gibson's ecological psychology (Gibson 1966, 1979) includes an ontology, his theory of affordances.

Gibson's first cut at describing affordances is deceptively simple. "The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*, either for good or ill" (Gibson 1979, 127). An affordance, this seems to imply, is a resource that the environment offers any animal that has the capabilities to perceive and use it. As such, affordances are meaningful to animals—they provide opportunity for particular kinds of behavior. Thus affordances are properties of the environment, but taken relative to an animal. So far, so good. Unfortunately, two pages later, Gibson's valiant, plainspoken attempt to make clear how much his theory of affordances differs from standard physicalist, reductionist ontology ends up just being confusing.

> [A]n affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective–objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical, yet neither. An affordance points both ways, to the environment and to the observer. (1979, 129)

This description makes affordances seem like impossible, ghostly entities, entities that no respectable scientist (or analytic philosopher) could have as part of his or her ontology. The purpose of this chapter is to provide a description of affordances that makes them more ontologically respectable, yet still does justice to Gibson's conception and, in so doing, to say how radical embodied cognitive scientists ought to understand affordances.

7.2 Affordances 1.0

Mine is, of course, not the first attempt to develop a coherent theory of affordances. It is worthwhile to say a few things about previous attempts, in order to see what is different about the theory outlined here. Previous (post-Gibson) attempts to set out an ontology of affordances have typically assumed that affordances are properties of the environment (Turvey et al. 1981; Michaels and Carello 1981; Heft 1989, 2001; Turvey 1992; Reed 1996; Michaels 2000).[1] These authors agree that affordances are animal-relative properties of the environment. In particular, affordances are properties of the environment that have some significance to some animal's behavior. To the extent that there is disagreement among these authors it is over two things: what kind of animal-relative properties of the environment affordances are, and what it is about animals that affordances are relative to.

There are two different views concerning the type of animal-relative properties of the environment that are affordances. Edward Reed (1996) argues that affordances are resources in the environment, properties of objects that might be exploitable by some animal, and he links this understanding of affordances to evolution by natural selection. Indeed, Reed takes this linkage between affordances and natural selection to be the most important thing about Gibsonian ecological psychology.

> The fundamental hypothesis of ecological psychology . . . is that affordances and only the relative availability (or nonavailability) of affordances create selection pressure on animals; hence behavior is regulated with respect to the affordances of the environment for a given animal. (Reed 1996, 18)

The resources in the environment are the source of selection pressure on animals, causing them to evolve perceptual systems that can perceive those resources. Those resources that some species of animal evolve the ability to perceive are affordances for members of that species. This selectionist view of affordances, in which they are environmental resources that exist prior to the animals that come to perceive and use them, is also semiendorsed by Stoffregen (2000). (In later work, Stoffregen [2003] does not endorse this view, however.)

In contrast to this selectionist view of affordances, which ties them closely to evolution by natural selection, is the Turvey-Shaw-Mace view, described in great detail by Turvey (1992), in which ecological psychology is tied more closely to physics than to evolutionary biology (see also Turvey et al. 1981)[2]. As discussed chapter 6, according to Turvey, affordances are dispositional properties of the environment. Dispositional properties are tendencies to manifest some other property in certain circumstances. "Being fragile" is a common dispositional property. Something is fragile just in case it would break in certain circumstances, particularly circumstances in which it is struck sharply. Dispositional properties are only conceivable when paired with actualizing circumstances, circumstances in which the disposition becomes manifest—the glass is only fragile if there are possible circumstances in which it might shatter. To say that affordances are dispositional properties of the environment, then, is to say that the environment is such that in some circumstances, certain other properties will become manifest. So, for example, the affordance "being edible" is a property of things in the environment only if there are animals that are capable of eating and digesting those things.

Notice that unlike Reed's view of affordances as resources, Turvey's account of affordances as dispositions is nonselectionist. Dispositions depend on possible actualizing circumstances; for example, nothing is soluble if there are no solvents. If affordances are dispositions, they depend on the possible presence of animals that can actualize them. Affordances, in Turvey's preferred language, must be *complemented* by properties of animals. So, an object can be edible only if there are animals that can eat and digest it. Given this, contrary to Reed's fundamental hypothesis, affordances per se cannot exert selection pressure on animals. Properties of the environment are not affordances in the absence of complementary properties of animals.

Turvey's insistence that affordances must be complemented by properties of animals brings us to the second difference among accounts of affordances: if affordances are animal-relative, we should wonder what it is about animals that affordances are relative to. Turvey (1992) proposes that affordances are complemented by effectivities (Turvey et al. 1981; Shaw, Turvey, and Mace 1982; see also chapter 6 above). Effectivities, like affordances, are dispositions, and as such they must be complemented by properties that lead to their actualization. Effectivities are properties of animals that allow them to make use of affordances. Effectivities and affordances are, thus, inseparable according to Turvey (1992). They complement one another. Claire Michaels (2000) also endorses this view. Another candidate for the aspect of animals to which affordances are relative is body scale. This view of affordances, endorsed by Harry Heft (1989, 2001), is suggested by empirical studies of affordances, which follow Warren's (1984) classic study of stair-climbing affordances in quantifying affordances with π-numbers, which are ratios between measures of body scale and measures of an environmental property.[3] Thus Stoffregen's (2000) discussion of affordances focuses on their relation to body scale. Heft (1989) provides a second reason for taking body scale to be the property of animals to which the affordances of the environment are related. Understanding affordances as body related, Heft suggests, can do justice to the phenomenological insights of Merleau-Ponty (1962) and the profound influence those insights had on Gibson (on which see Heft 2001).

To summarize this brief discussion of some of the previous theoretical work on affordances, we can say the following. First, Turvey, Heft, Michaels, and Reed agree that affordances are animal-relative properties of the environment. Second, there is some disagreement over whether these properties exist independent of animals. This disagreement comes to an argument over whether affordances are resources that guide natural selection, or dispositional properties of the environment that must be complemented by some property of animals. Third, there is disagreement over whether the relevant properties of animals are abilities (or effectivities) or body scale.

In a paper published a few years ago (Chemero 2003a), I outlined a theory of affordances designed to avoid these two controversies. It did so by disagreeing with the premise on which they are based, the claim that affordances are animal-relative properties of the environment.[4] I argued that affordances are not properties of the environment; indeed, they are not even properties. Affordances, I argued, are relations between particular aspects of animals and particular aspects of situations. As I will explain, I still believe that this is a significant improvement over prior work, but it is nonetheless not sufficient as a theory of affordances.

7.3 A Few Critical Comments Regarding Affordances 1.0

In this section, I will argue very briefly against the idea that affordances are properties of the environment. This argument will have two parts. First, I will argue that affordances are not properties, or at least not always properties. Second, I will argue that affordances are not in the environment.

7.3.1 Affordances and Properties

In "What We Perceive When We Perceive Affordances" (Chemero 2001c), I argued that it is vital to distinguish between features and properties when discussing affordances. The purpose of that discussion was to counter Michaels's claim that perceiving ball-punching affordances (as in Michaels, Zeinstra, and Oudejans 2001) is perceiving something about oneself, not something about the environment. This, I argued, is true only if one fails to realize that there is a more primitive way of perceiving the environment, involving what Strawson called feature placing (Strawson 1959; Smith 1996). Feature placing is easiest to understand in contrast to the perception of objects with properties. Compare, for example, realizing that your car is dented to realizing that it's raining. In the former case, the perception of a property of the car, you must (a) perceive a particular entity; (b) know its identity, that it is your car; (c) know what it is to be dented; and (d) perceive that this particular entity (your car) has this particular property (being dented). In the latter case, the placing of a feature, there is no need to know anything about any particular entity. All that is necessary is the ability to recognize a feature of situations (raininess). To see this, consider that the "it" in "it is raining" is never the same thing; it refers to a situation (what's going on right here, right now) that will never appear again. We can ask what is dented, but we cannot ask what is raining.

Drawing attention to this distinction between placing features and perceiving properties of objects is relevant to the perception of affordances because Michaels (2000) argued that when we perceive ball-punching affordances, we perceive that "it's time to flex the elbow." This, she argues, is perceiving something more about yourself than about the environment. The recognition of feature placing calls this into question. Perceiving that it is time to flex the elbow is like perceiving that it is raining. It is a matter of perceiving that the situation as a whole has a certain feature, that the situation as a whole supports (perhaps demands) a certain kind of action. All of this is to say that perceiving affordances is placing features. And because features are not properties, views of affordances that take them to be properties can't be right.

7.3.2 Affordances and the Environment

If one accepts that affordances are not properties of objects, it is a small step to see that affordances cannot be properties, or even features, of the environment alone. I have just argued that affordances are features of whole situations. Animals are, of course, crucial parts of these whole situations, so perceiving something about the whole situation cannot be perceiving something about the environment, divorced from the animal. Thus, as Stoffregen (2003) suggests, affordances must belong to animal–environment systems, not just the environment. Though I agree with Stoffregen on this point, I'd like to argue for something more specific: that affordances are relations. To see this, consider Harry Heft's (2001) discussion of the relation between Gibson and the American naturalist William James.

In *Ecological Psychology in Context* (2001), Heft argues quite convincingly that Gibson's ecological psychology is a descendent of the radical empiricism of William James. To the radical empiricist, perception is direct because it is an act that includes the thing perceived. This leads to what James called "the problem of two minds." Suppose you and I both perceive the same pint of Guinness. The pint, according to radical empiricism, is part of both my perception and yours. But this leads to a problem of mereology: if the pint is part of both our perceptions, then our minds overlap. This, James thought, is in direct conflict with the (to him) obvious fact that our minds are private. The problem of two minds, then, is as follows. If perception is direct and two individuals can perceive the same object, then how can their minds be truly separate? James struggled with the problem of two minds throughout his later years, never reaching a satisfying resolution. This same problem affects any theory of direct perception, including Gibson's ecological psychology. Affordances are part of the act of perception, so if you and I both perceive the affordance "potability" of the pint of Guinness, our perceptions overlap. Our experiences, and hence our minds, are not private.

The solution to this problem is apparent in another of the main tenets of Jamesian radical empiricism. According to radical empiricism, everything that is experienced is equally real. Among the things we experience are relations between things; so relations are real, with the same status as the things that stand in relations. To solve the problem of two minds, suppose that perceivables are relations between perceivers and aspects of situations. If that is true, you and I can both perceive the potability of the Guinness, without our perceptions overlapping. You will perceive the relation between you and the pint, while I will perceive the relation between me and the pint, and our perceptions can remain private. The key to this solution, though, is that what we perceive, the affordance potability, is not in the environment alone. It is, instead, the relation between the perceiver and the environment. This point, that affordances are *relations*, is the key to the theory of affordances I will describe in the following sections.

Here, I call the view of affordances I am about to describe, the one which I set out in the 2003 paper mentioned above, "Affordances 1.1." In section 7.5, I argue that it resolves the issues I have been pointing to for Affordances 1.0. However, I have come to believe that Affordances 1.1 is not, in itself, sufficient as a theory of affordances. In section 7.6, I use Affordances 1.1 as the basis for a sketch of Affordances 2.0, a theory of affordances that meshes well with dynamical systems explanations and, hence, is more appropriate for radical embodied cognitive scientists.

7.4 Affordances 1.1

I have said several times that affordances are relations between animals and features of situations. I will now spell out in detail what that means. To begin, here is the basic logical structure of affordances, which will be expanded later.

Affords-ϕ (environment, organism), where ϕ is a behavior.

Translated literally into English, this means "The relation 'affords-ϕ' holds between 'environment' and 'organism'"; translated more loosely and colloquially, this means "The environment affords behavior ϕ for the organism." To get an idea of what this means, and what it means to say that affordances are relations, compare it to a more familiar relation.

Taller-than (Shaquille, Henry).

This says that Shaquille is taller than Henry. Notice first that the only objects in this relation are Shaquille and Henry. The taller-than relation is not inherent in either of them, but depends on both of them for its existence. Affords-ϕ is like taller-than in this respect: it is neither of the person, nor of the environment, but rather of their combination. Second, the affordance is not an extra thing in any of the usual senses of "thing." Yet it exists nonetheless, and, like the fact that Shaquille is taller than Henry, it is directly perceivable. (Remember the discussion of entropy in chapter 6.) Taking affordances to be relations, despite the fact that they are not things in the usual sense, is quite plausible in light of Heft's (2001) account of Gibson as a Jamesian radical empiricist (Chemero 2003c). As noted above, according to the radical empiricist relations are perceivable, and anything perceivable is real.

The formal definition of affordances as relations between organisms and environments is incomplete. In the next few section, I will fill it out.

7.4.1 The Environmental Relata

As discussed above, perceiving affordances is placing features, seeing that the situation allows a certain activity.[5] Thus the environmental relata in affordances must be features, not properties. The only further comment here is that this is in direct disagreement with Turvey, who pronounces that "There are only propertied things" (1992, 176). Situations are not things; features are not properties.

7.4.2 The Organismal Relata

Ever since Warren's (1984) groundbreaking experiments on stair climbing, it has been (tacitly) assumed by experimentalists that the aspect of animals that determines what the environment affords, the organismal relata in the affordance relation, are aspects of body scale. Warren, in attempting to quantify affordances for stair climbing, quantified them as unitless π numbers, the ratio between leg length and riser height. The affordance climbability is then identified as this ratio. Subsequent experiments identified affordances similarly, as ratios between body scale and some bit of the environment measurable in the same units. (See, e.g., work on gap crossing by Mark 1987; Burton 1992, 1994; Jiang and Mark 1994; Cornus, Montagne, and Laurent 1999; Mark et al. 1999.) Many experimentalists, I suspect, have not given much thought to this fact, simply assuming that what they are measuring just are affordances. Given the discussion above, it might seem natural to say that the affordance is expressed as the following relation:

affords-climbing (my leg length, riser height),

which is perceivable whenever the ratio of my leg length to the riser height is within a certain range. Doing so is a mistake: it must be remembered that body scale is just an easily quantifiable stand-in for ability. Most theoretical work on affordances does not make this mistake, pointing out that the animal-side counterparts of affordances are effectivities (e.g., Turvey 1992; Reed 1996; Stoffregen 2000; Michaels 2000; see above for discussion).[6]

Although body scale is easily measured, it is only occasionally a good placeholder for ability. In most cases, there is not a tight relation between body scale and ability. Indeed, recent research seems to be calling even the paradigm cases of body-scaled studies of affordances into question. Consider, for example, research by Cesari, Formenti, and Olivato (2003) on stair-climbing affordances. The experiments they report indicate that subjects perceive stair-climbing and descending affordances not as the ratio between leg length and riser height (as Warren 1984 holds) but rather as a relation between stepping ability and riser height. In the Cesari et al. study, subjects were asked to determine the highest step they could climb; this variable was called "perceived riser height." Subjects were then asked to (1) approach the steps from a distance of four meters as if they were going to climb them, (2) stop, and then (3) climb the stairs. The important variable here was distance from the subject's foot to the stair bottom when the subject stopped. It was found that different types of subject (children, young adults, older adults) had the same optimal ratio of distance from step to riser height, which is to say that they had the same ratio for the highest step they could climb. This ratio is a function of stair-climbing ability, not leg length. To see this, consider further results from the same set of studies. First, there was an important difference between older adults, on the one hand, and younger adults and children, on the other: older adults maintained the optimal ratio of distance from step to step height for steps as much as 10 percent shorter than the maximum steps they could climb, whereas in younger subjects the ratio changes significantly for steps 10 percent shorter than the maximum climbable height. Finally, Cesari et al.

found that older adults are significantly less flexible than younger adults and children.

Put together, Cesari et al.'s results indicate quite strongly that the relevant animal-side variable for stair-climbing affordances is climbing ability. First, there is an optimal ratio of height to distance for stair climbing, and all subjects used this information to determine the tallest step they could climb. Less flexible older adults maintained this ratio even for steps lower than their highest climbable steps; young adults and children did not. Given the flexibility results, older adults have different stair-climbing abilities than young adults and children. They also use the ratio differently, choosing to maintain the optimal ratio even for situations in which they can climb stairs relatively easily. So, the ratio, which is the aspect of the environment perceived in determining climbability, is perceived in terms of ability.

Two more experiments, done in my lab, show similar results for gap-crossing affordances. In the first, we (unpublished data from experiments discussed in Chemero, Klein, and Cordeiro 2003) asked subjects to stand on a platform and judge whether or not they could step across a series of differently sized gaps onto another platform. We then measured the subjects' leg length and the length of an actual step they took on the floor (not on the platform). We calculated two π numbers: one is the ratio of leg length to the maximum gap size subjects judged they could cross; the second is the ratio of step length to the maximum gap size subjects judged they could cross. We found first that subjects perceived gap-crossing affordances very accurately: the ratio of step size to maximum gap judged crossable was equal to one. Second, we found that step size was much more highly correlated with maximum gap judged crossable than leg length was. In fact, partial regression revealed that the correlation between maximum gap judged crossable and leg length, with the effect of step length subtracted, was zero. So our subjects perceived gap-crossing affordances very accurately, and they did so in terms of their stepping abilities, not leg length. In another experiment (Fox and Chemero, unpublished data), we compared the perception of gap-crossing affordances by college students and senior citizens. As in the prior study, we measured leg length and step size, and determined the maximum gap the subjects judged they could step across. We found that the mean leg lengths for college students and senior citizens were the same, but that college students stepped farther and judged that they could cross larger gaps. For both groups, step size was highly correlated with maximum gap size judged crossable. But only college students had significant correlations between leg length and maximum gap judged crossable. Like the Cesari et al. experiments, these studies show that if one can separate body scale and ability, it becomes clear that affordances are functions of, and perceived in terms of, ability.

Affordances, then, are relations between abilities and features of the environment. Affordances, that is, have this structure:

Affords-ϕ (feature, ability).

7.4.3 Affordances, Abilities, and Dispositions

Thus far, I have been using the words "ability" and "effectivity" more or less interchangeably. There are two things about effectivities as they are typically discussed that makes them different from abilities, however. First, effectivities are defined as the organismal complement to affordances *qua* dispositional properties of the environment (Turvey et al. 1981; Shaw, Turvey, and Mace 1982; Warren 1984; Turvey 1992). I have been arguing that affordances are not properties of the environment; thus there is no need for the complementing property in the organism. Second, effectivities are defined as dispositions. Abilities are not dispositions.

The problem with seeing abilities as dispositions is that when coupled with the right enabling conditions, dispositions are guaranteed to become manifest. The soluble solid sugar will always dissolve in water in suitable conditions. This is not true of abilities. Having the ability to walk does not mean that one will not fall down even in the ideal conditions for walking.[7] This is to point out that there is something inherently normative about abilities. Individuals with abilities are supposed to behave in particular ways, and they may fail to do so. Dispositions, on the other hand, never fail; they simply are or are not in the appropriate circumstances to become manifest. A better way to understand abilities is as functions. Functions depend on an individual animal's developmental history or the evolutionary history of the species, both of which occur in the context of the environment. Given this, it is actually more appropriate to understand abilities, like affordances, as being inherent not in animals, but in animal–environment systems. That is, like affordances, abilities are relations.[8] Abilities come to play the role they do in the behavioral economy of the animal because, at some point in the past, they helped the animal (or its ancestor) to survive, reproduce, or flourish in its environment. Yet even in identical circumstances to those in which they were helpful in the past, abilities can fail to become manifest; there can, that is, be a malfunction. By taking abilities to be functions, we can account for the fact that even on a firm surface, with no wind, while perfectly healthy and sober, I may fail in my attempt to climb a step that affords climbing for me. This is inconceivable in the case of dispositions, which necessarily become manifest whenever their actualizing circumstances are present.[9]

This analysis of abilities, as functions and not as dispositions, has a further noteworthy consequence. Since functions depend on evolutionary history and affordances are partly constituted by functions, affordances are tied to evolution. This makes ecological psychology a branch of biology, and a truly ecological science (Withagen and Chemero 2009). But notice that it does so without being selectionist in the way Reed's (1996) understanding of affordances is.[10] That is, it does not assume that affordances are resources that exert selection pressure. I take it that being evolutionary and ecological but not selectionist is a positive feature of the theory of affordances outlined here. First, there are (admittedly highly controversial) reasons from theoretical evolutionary biology to be skeptical of selectionist views of evolution.[11] And, indeed, radical embodied cognitive scientists should align themselves with nonselectionist, developmental systems approaches in biology (Griffiths and Gray 2001; see section 7.6 below for more on the connection between radical embodied cognitive science and these approaches in biology). Second, a selectionist view of the relationship between affordances and animals fails to do justice to the mutuality of animal and environment (Gibson 1979). If it is affordances that exert selection pressure, it cannot be, as Gibson suggested, that animals imply niches (sets of affordances—see below) and vice versa. Rather, on the selectionist view, it is affordances that are in the driver's seat, and animals must conform to them over evolutionary history. On the view being offered here, there is true animal–environment mutuality. Affordances, which are the glue that holds the animal and environment together, exist only in virtue of selection pressure exerted on animals by the normal physical environment. They arise along with the abilities of animals to perceive and take advantage of them.

7.4.4 Perceiving Affordances

Any account of the ontology of affordances requires a story about perceiving affordances. Perception for the radical embodied cognitive scientist is direct and can be conceptualized as a relation between the perceiver and what is perceived. On the account of affordances outlined here, this relation looks like this.

Perceives [animal, affords-ϕ (feature, ability)].

This is the act of perception that is studied by the psychologist, from a third-person perspective. Typically, though, an animal is consciously aware only of the affordance relation, and not the constituent relata. That is, from the point of view of a behaving animal, the structure of the perception of affordances will be this:

Perceives [animal, affordance-for-ϕ].

This is surely the usual phenomenology of humans. I am normally not aware of anything about my climbing abilities or riser heights when I perceive that I can climb a step. Humans, however, can—with training and when so inclined—perceive things about their abilities and features of the physical environment. I suspect that most nonhuman animals are incapable of this.

7.5 Ecological Ontology

Affordances 1.1 are neither properties of the animal alone, nor properties of the environment alone. Instead they are relations between abilities and some feature of a situation. They are not kickable and often not easily localizable physically, but they are nonetheless perfectly real and perfectly perceivable. There are still unanswered questions about affordances. In the next sections, I will answer three of them.

7.5.1 Affordances and Niches

Gibson (1979) points out that a niche is the set of affordances for a particular animal. Different animals, with different abilities, may have physically colocated but nonetheless nonoverlapping niches. For example, a human and a bacterium may share a physical location (as when a bacterium is inside a human), but their niches will not overlap. As noted above, Gibson also suggests that this is the way to make sense of the mutuality of animals and environments. An animal's abilities imply an ecological niche. Conversely, an ecological niche implies an animal. Given the relational definition of affordances, we can make sense of these facts about niches.

Start by taking organisms to be sets of abilities. These abilities will be interconnected, of course. An animal cannot have the ability to run if it cannot maintain its posture, nor will it be able to climb a tree if it cannot affix itself to things (with suction, by grabbing, etc.). As Reed (1996) points out in his revealing analysis of action systems, all other abilities will depend on basic orienting abilities, abilities to maintain posture, and the like. There will also be a nested structure of abilities, in which larger abilities will be composed of smaller-scale abilities. Each of an animal's abilities will have a set of situations in which it can be exercised. But no larger-scale ability will be exercisable in situations in which its component smaller-scale abilities can't be exercised; similarly no ability will be exercisable in situations in which a more basic ability on which it depends cannot be exercised. Thus, if walking is leg swinging, falling, and catching yourself, walking will be impossible in situations in which one cannot swing a leg, or fall, or catch oneself. Walking will also not be possible in situations with no gravity, or too much gravity, or in which the atmosphere is in flames, because the basic orienting system on which walking depends is inoperable in these situations.

All this said, we can define an animal's niche as the set of situations in which one or more of its abilities can be exercised. To put this formally, start with the set of all possible situations, S. For each ability a_i there is a subset of S, s_i in which that ability can be exercised. Suppose an organism has abilities $a_1 \ldots a_n$. That organism's niche will be the union of $s_1 \ldots s_n$, for each ability $a_1 \ldots a_n$ that the organism has. This collection of situations forms the organism's cognitive, behavioral, and phenomenological niche.

7.5.2 Affordances and Events

In his 2000 paper "Affordances and Events," Stoffregen argues that events, conceived as changes in the physical layout, are not perceivable according to ecological psychology. This is the case, he argued, because what we perceive are affordances, and events and affordances are of different ontological kinds. In response, I (Chemero 2000b) offered a different understanding of events, an understanding according to which event perception is not problematic for ecological psychology. Perceivable events, I argued, are changes in the layout of affordances. A later paper by Chemero, Klein, and Cordeiro (2003) provides experimental evidence that events so described can be perceived. I will discuss these data in chapter 9. Here, I will limit discussion to how the definition of affordances outlined here impacts the theory of events just described.

Assume that affordances are relations between abilities and features of environmental situations and that events are changes in the layout of affordances in the animal–environment system. How, then, do events happen? Equivalently, how do affordances change? Most changes in relations between the abilities of animals and environmental situations will be changes in environmental situations. Most events, that is, will result from changes in the physical environment. If the glass of water spills, the affordance drinkability disappears because my drinking abilities are not appropriate for spilled water; once the apple falls from the tree it is edible, because my being able to grasp the apple is a necessary condition for my being able to eat it. In cases such as these, there are events, changes in affordances, without changes in abilities. There can also be changes in affordances without changes in the features of the environment. The very same stair no longer affords climbing to an individual whose stepping abilities have decayed as a result of aging. Since abilities typically change more slowly than the environment, these events will happen less frequently than events that result from changes in the environment.

7.5.3 Do Affordances Exist without Animals?

For all the noise ecological psychologists make about being realists, it is not obvious at the outset that ecological psychology is not a form of idealism, in which perceivables exist only when they are perceived. It is a small step from this to a global idealism, in which the world disappears whenever I close my eyes. (See chapter 9 for a more general discussion of realism and radical embodied cognitive science.) Reed's (1996) conception of affordances as resources that exert selection pressure avoids this issue by making it the case that affordances exist unproblematically, even without animals capable of perceiving them. Other understandings of affordances must face this problem. For Turvey et al. (1981), Warren (1984), Turvey (1992), and Michaels (2000), who claim that affordances must be complemented by the effectivities of animals, the status of affordances is unclear in the absence of animals. Similarly, if affordances are relations between abilities and situations, as in Affordances 1.1, affordances depend in some sense on animals. The questions that must be answered are: In what sense do affordances depend on animals? Do affordances exist without animals?

As a first pass at answering these questions, I will once again coopt some terminology from Dennett (1998). Dennett distinguishes between things that are lovely and things that are suspect. To see the distinction, consider that a female hippopotamus in a zoo might be lovely, even if no male hippopotamus has ever seen her. She is lovely just in case if a male hippopotamus were to see her, he would find her to be so. The key is that being lovely depends on a potential observer, not an actual act of observation. Compare this to being suspect. To be suspect, something actually has to be under suspicion. Being suspect requires an actual observer. Whether affordances exist without animals is a matter of whether affordances are lovely or suspect. Affordances, we can see, are lovely. A feature of some situation might exist just as it is even if there are no animals. There will be affordances in which that feature takes part as long as some animal exists with the appropriate ability. This is the case even if that animal is nowhere in the vicinity of the situation that affords something to it.

Affordances do not disappear when there is no local animal to perceive and take advantage of them. They are perfectly real entities that can be objectively studied and are in no way figments of the imagination of the animal that perceives them. So radical embodied cognitive science is not a form of idealism. But affordances do depend on the existence of some animal that could perceive them, if the right conditions were met. Because affordances, the primary perceivables according to ecological psychology, depend in this way on animals, the ontology of ecological psychology is not a simple form of realism. It is a form of realism about the world as it is perceived and experienced—affordances, which are inherently meaningful, are in the world, and not merely projected onto it by animals. I will say a lot more about how this can be true in chapter 9.

7.6 Affordances 2.0

When considered within the confines of the ongoing theoretical debate among ecological psychologists about the nature of affordances, the theory of affordances just outlined seems to me to be a significant improvement over prior attempts to give a formal theory of affordances. (Of course, it would seem that way to its author.) In retrospect, however, it seems inadequate to the practice of radical embodied cognitive scientists. It is inadequate because radical embodied cognitive scientists are dynamicists, and in their actual experimental practices they understand affordances dynamically. Yet the formal theories of affordances offered by Gibson and later ecological psychologists (Affordances 1.0) and my attempt to improve on them (Affordances 1.1) define affordances statically. Perhaps this occurred because the view of affordances that forms the basis for this discussion was originally formulated by Turvey et al. (1981) at a time before dynamical systems modeling had so thoroughly infected ecological psychology. Now, however, it seems clear that we (ecological psychologists, radical embodied cognitive scientists) need a theory of affordances that is dynamical root and branch. My Affordances 1.1 is not that. Radical embodied cognitive science requires Affordances 2.0, a dynamical theory of affordances.

To formulate Affordances 2.0, start with Affordances 1.1, according to which affordances are relations between abilities to perceive and act and features of the environment. Then consider the interaction over time between an animal's sensorimotor abilities, that is, its embodied capacities for perception and action, and its niche, that is, the set of affordances available to it. This is depicted in figure 7.1. Over developmental time, an animal's sensorimotor abilities select its niche—the animal will become

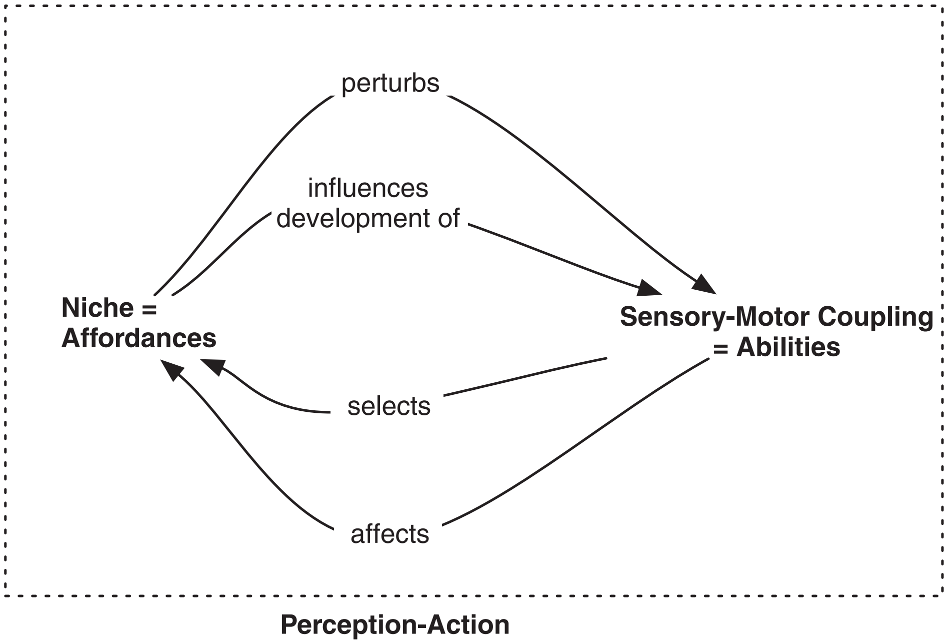

Figure 7.1

selectively sensitive to information relevant to the things it is able to do. Also over developmental time, the niche will strongly influence the development of the animal's ability to perceive and act. Over the shorter time scales of behavior, the animal's sensorimotor abilities manifest themselves in embodied action that causes changes in the layout of available affordances, and these affordances will change the way abilities are exercised in action. The key point here is that affordances and abilities are not just defined in terms of one another as in the dispositional and relational views discussed above, but causally interact in real time and are causally dependent on one another.

There are three noteworthy consequences of the shift to Affordances 2.0.

First, this is not so much a new way of understanding affordances as a critique of prior attempts to come up with a *definition* of the term "affordance." Ecological psychologists have always been aware of, indeed keenly interested in, the interaction of affordances and abilities in real time. As noted above, radical embodied cognitive scientists (including ecological psychologists) study perception and action dynamically. Affordances 2.0 is an attempt to develop a theoretical understanding of affordances that is more in line with the experimental and explanatory practices of ecological psychologists. (See Chemero 2008.)

Second, notice that this reconceptualization of affordances is a variety of niche construction that occurs over shorter time scales and in which the constructed niche is an animal's individual behavioral, cognitive, and phenomenological niche. In more standard biological niche construction, the activity of some population of organisms alters, sometimes dramatically, the population's own ecological niche as well as those of other organisms (Odling-Smee, Laland, and Feldman 2003). These animal-caused alterations to niches have profound and wide-reaching effects over evolutionary time. And, indeed, the population of organisms and the niche are so tightly coupled that Griffiths and Gray (2001) recommend that they form a unified developmental system that is to be modeled with just one variable *Œ*. The dynamics of this variable are specified in the following equation

$$d\textit{Œ}_{\text{pop}}/dt = f(\textit{Œ}_{\text{pop}}, E) \tag{7.1}$$

in which $\textit{Œ}_{\text{pop}}$ is the coupled organism–niche system for the population and E is the physical environment. The variety of niche construction sketched in Affordances 2.0 is an equally tightly coupled animal–environment system. It differs from the much-discussed biological case in two ways. First, the constructed niche is for an individual organism, not for a population. Second, it occurs over shorter time scales—an animal's activities alter the world as the animal experiences it, and these alterations to the phenomenological-cognitive-behavioral niche, in turn, affect the animal's behavior and the development of its abilities to perceive and act, which further alter the phenomenological-cognitive-behavioral niche, and on and on. Affordances 2.0, therefore, emphasizes the connections between radical embodied cognitive science and its natural allies in biology, that is, developmental systems and niche construction.

Third, this reconceptualization of affordances is explicitly formulated to make the natural, but largely unmade, connections between ecological psychology and another form of radical embodied cognitive science: the burgeoning enactivist movement in the cognitive sciences (Varela, Thompson, and Rosch 1991; Thompson 2007). Figure 7.2 is an expanded version of figure 7.1, expanded to show the connection between organisms and sensorimotor coupling, as understood by enactivists, and Affordances 2.0. Enactivists view the organism as a self-organizing, autonomous, autopoietic system. In this system, the nervous system generates neuronal assemblies that make sensorimotor abilities possible, and these sensorimotor abilities modulate the dynamics of the nervous system. Combining Affordances 2.0 with enactivist studies of the organism makes radical embodied cognitive science a fully dynamical science of the entire brain–body–environment system: nonrepresentational neurodynamic studies of the nervous system and sensorimotor abilities (Cosmelli, Lachaux, and Thompson 2007;

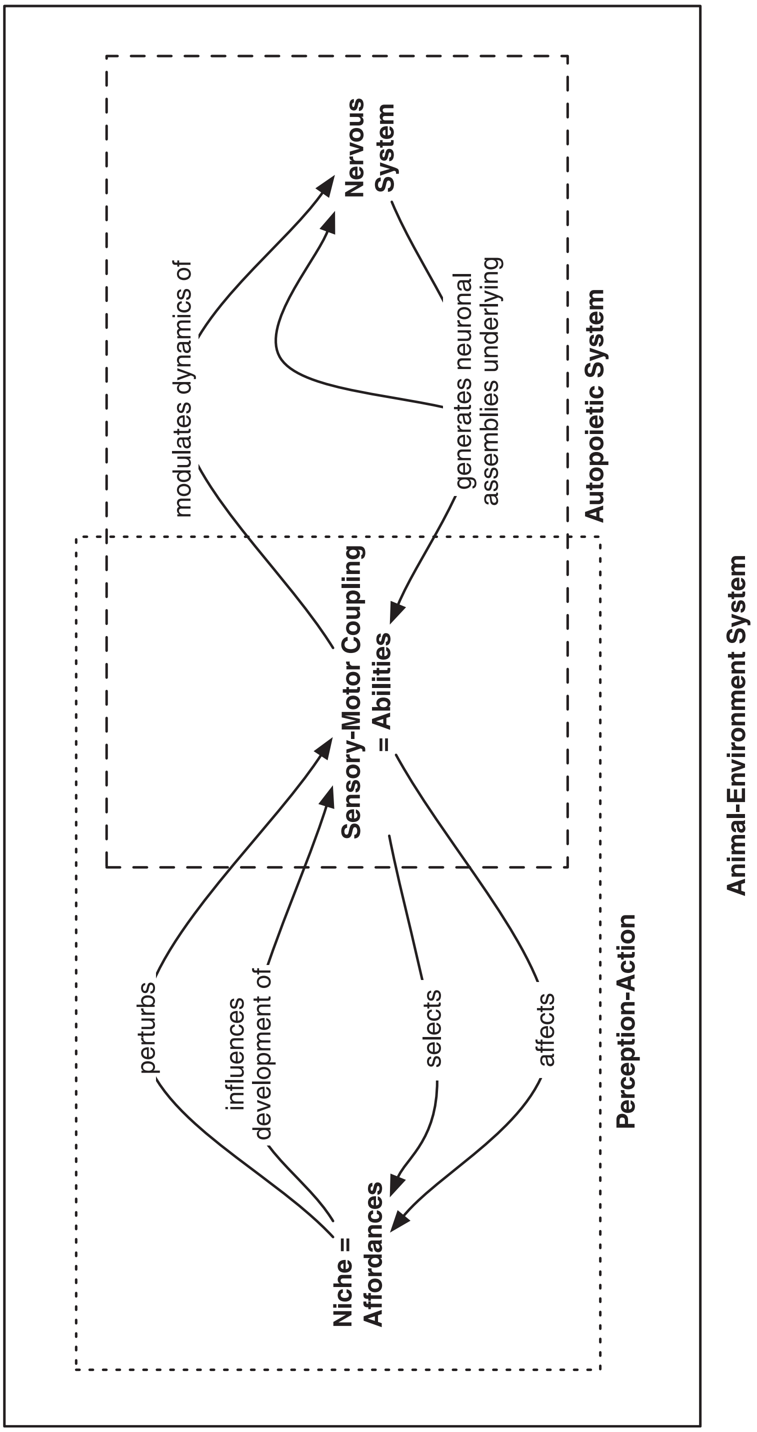
Niche =
Affordances
perturbs
influences
development of
selects
affects
Sensory-Motor Coupling
= Abilities
modulates dynamics of
generates neuronal
assemblies underlying
Nervous
System
Autopoietic System
Perception-Action
Animal-Environment System

Figure 7.2

Thompson and Varela 2001) match up with ecological psychological studies of affordances and sensorimotor abilities. Obviously, much more work is required to genuinely integrate ecological and enactive cognitive science under the banner of radical embodied cognitive science. These two approaches have more in common than their proponents realize. Combining them could make radical embodied cognitive science a much more significant force in the cognitive science community than either the ecological or enactive movements are separately.

7.7 Information about Affordances

So far I've given some details about what affordances are, how they relate to other important pieces of an ecological ontology, and how they ought to be understood dynamically. But I haven't directly addressed what Gibson called the central question for the theory of affordances: "The central question for the theory of affordances is not whether they exist and are real, but whether information is available in the ambient light for [directly] perceiving them" (1979, 140). I have already discussed several cases that imply that there is information available in the ambient light for the perception of affordances without mental gymnastics. For example, τ and its temporal derivatives provide information about affordances for locomotion, pursuit, and collision avoidance. (See also the discussions of gap-crossing and stair-climbing affordances in this chapter.) In this section, I will discuss one more piece of evidence that there is information in the environment that enables the direct perception of affordances: perception of moveability by dynamic touch. I will discuss dynamic touch (a.k.a. haptic perception) here for three reasons. First, it is an area of very active research in ecological psychology, but is not well known to most cognitive scientists. Second, cognitive scientists and philosophers tend to focus too narrowly on visual perception, ignoring the other senses. Third, perception by dynamic touch is a temporally extended process, which shows the value of a dynamic conception of affordances.

7.7.1 Perception by Dynamic Touch

Right now, you are holding a book.[12] You can see the book, but if you were to close your eyes, you would still have considerable information about the book. The book is exerting mechanical pressure on the portions of your fingers that are touching it. If you move your fingers along the surface of the book, you can feel its texture. Notice that only a small portion of the book is actually in contact with only parts of your fingers. Yet, even with your eyes closed, you can tell how big the book is: not just how heavy, but also how long, wide, and thick. You can also tell its orientation. (Is it parallel to the floor? Perpendicular? At some other angle?) This is possible because your muscles are working so that you can hold the book. You are applying pressure with your fingers to keep the book from slipping through them, and you are also working with your wrist and the rest of your arm to fight gravity and keep the book in place, in a position where you would easily be able to see it were your eyes open. This is dynamic touch.

Dynamic touch is *action*. Holding the book takes work from your muscles, and this work causes pressure and deformation to muscles and tendons, stimulating mechanoreceptors just as light stimulates retinal cells. This stimulation is the primary source of information for dynamic touch. (See Gibson 1962, 1966; Turvey 1996; Carello and Turvey 2000.) Notice too that to perceive the book by dynamic touch, you have to heft it; that is, you have to intentionally move it around, actively exploring the way it exerts forces on the muscles of your hands, wrists, and arms. As you move the book, the forces it exerts on your body change, which changes the way you experience the book and the affordances for continued active exploration of the book.

Solomon and Turvey (1988) studied the perception of length by dynamic touch.[13] Subjects were seated, placed their arm on a table so that they could move only their wrist and hand, and asked to grab the end of a rod occluded by a curtain. They judged the rod's length by moving a visible object (a block of wood on wheels) to the distance of the rod's length with their other hand. Subjects were quite good at this task, but what information were they using? Length per se cannot be perceived by dynamic touch, because length cannot affect mechanoreceptors. Solomon and Turvey showed that the length of rods is perceived by sensitivity to their rotational inertia, their resistance to turning about the wrist. Mathematically speaking, rotational inertia is quantified as the primary moment of inertia, which is approximately

$$I_1 = 1/3\ m \cdot L^2, \tag{7.2}$$

where m is the mass of the rod and L is its length.[14] Solomon and Turvey's results suggest that this moment of inertia is the higher-order variable that provides information about the length of a wielded rod, and it does so in a form that mechanoreceptors are sensitive to. But think again about the book you're holding. If you hold it in just one hand and use just your wrist, you can rotate it up and down, side to side, and you can twist it. That is, you can rotate it in three dimensions. Notice too that it will resist your attempts to move it differently in different directions. So just one moment of inertia will not provide enough information to determine an object's length. Instead, one must use the *inertia tensor*, a three by three matrix that can be used to determine three principal moments of inertia, called eigenvalues. These three moments of inertia are the source of information for perception by dynamic touch (Turvey 1996). In fact, the values of these three moments of inertia define an ellipsoid, a three-dimensional virtual shape that resembles the shape of a wielded object. For example, a sphere generates rotational inertia so that the three principal moments of inertia are equal. Burton, Turvey, and Solomon (1990) showed that subjects could use this information to perceive shape: their subjects could visually select objects that matched occluded objects they wielded (by holding a stick attached at the base, so they could not feel the object's shape).

Later work on dynamic touch showed that subjects could perceive a wide variety of properties of objects based on the inertial tensor. For reviews, see Turvey 1996; Pagano and Turvey 1998; and Carello and Turvey 2000. Remember, though, that the issue for ecological psychology is whether there is information available for affordances, and so far I have been discussing object properties such as mass, length, and shape. I will now briefly discuss research showing that there is information for perceiving affordances by dynamic touch.

7.7.2 Dynamic Touch and the Size-Weight Illusion

One of the riddles I remember from when I was a child went as follows: Which is heavier, a pound of feathers or a pound of lead? An eight-year-old who thinks he is smart will quickly answer that the pound of lead is heavier. Most eight-year-olds will make this mistake only once. When asked to make this judgment by touch rather than by knowledge of the materials, however, they will continue to make the mistake.[15] This is called the *size-weight illusion* (Charpentier 1891; Murray et al. 1999). Given two objects of equal mass, people (both children and adults) judge the one with a smaller diameter to be heavier. That is, they judge a comparatively small pound of lead to be heavier than a comparatively large pound of feathers. This illusion has typically been taken to be the result of mental gymnastics: one judges an object's size and uses this judgment to (erroneously) correct one's judgment of weight.

To investigate the size-weight illusion, Amazeen and Turvey (1996) created what are called *tensor objects*. Tensor objects are composed of two rods connected to form a "plus" sign, with a third rod attached perpendicular to the point at which the two rods forming the "plus" sign intersect. Metal

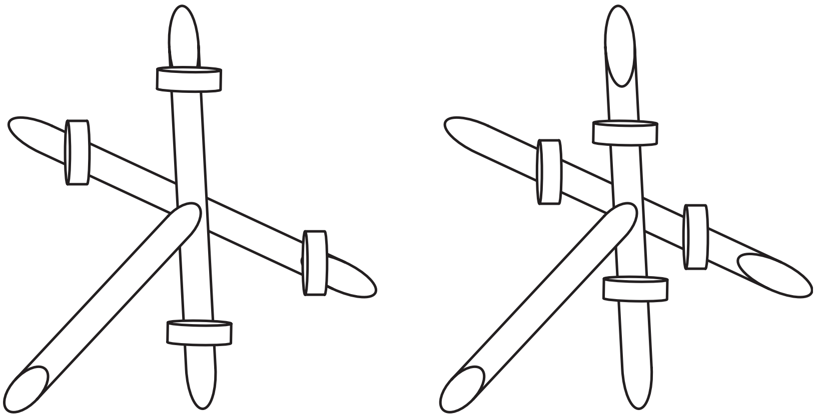

Figure 7.3
A tensor object (Amazeen and Turvey 1996).

rings are attached at different locations on the rods on the object so as to create tensor objects with different weight distributions. (See figure 7.3.) Different weight distribution means different moments of inertia (i.e., different pressures felt at the wrist joint). These tensor objects allowed Amazeen and Turvey to construct a series of objects of identical size and weight, but with different moments of inertia designed to mimic the stimuli typically used in experiments that produce the size-weight illusion. They found that subjects judged the heaviness of these tensor objects as predicted by their inertia tensors, despite the fact that they were the same size and weight. This occurred both when subjects wielded the objects occluded behind a curtain and when they could see the objects (covered tightly with paper to make their volume apparent but hide the distribution of the metal weights). Thus, Amazeen and Turvey showed that subjects do not perceive the weight of objects by judging their size (whether by touch alone or by vision and touch) and mentally combining that judgment with their felt force due to gravity. Instead, they use the information available at the wrist, as determined by the object's inertial tensor.

The point of the above is that humans do not misperceive weight by mentally calculating it, computationally combining size and force due to gravity. Instead, they correctly perceive the information in the inertial tensor, which does not carry information about weight alone. (Remember from formula 7.1 that the principal moment of inertia is a function of both length and mass.) What, then, is the information in the inertial tensor about? More recent work by Shockley, Carello, and Turvey (2004) indicates that these subjects perceive the affordance *moveability*. In particular, they showed that subjects who misjudge the weight of an object when falling prey to the size-weight illusion make nearly identical judgments about whether the object is moveable. That is, the question "On a scale of 1 to 100, where 50 is the control object, how heavy is this object?" gets the same answer as "On a scale of 1 to 100, where 50 is the control object, how easy is it to move this object?" So the size-weight illusion occurs because subjects are actually basing their judgments on perception of moveability, not weight, and their judgments about moveability are accurate. Indeed, it makes perfect sense that moveability is what subjects perceive by dynamic touch. Many, many experiments have confirmed that an object's inertial tensor is the source of the information for dynamic touch. As noted above, the eigenvalues of an object's inertial tensor carry information about the object's rotational inertia, which is to say they carry information about the object's tendency to resist rotation in a particular direction, which is to say they carry information about how difficult it is to move the object.

There is, then, information about the affordance moveability available for direct perception. I have also provided descriptions of empirical investigations into the directly perceivable information available for affordances for climbing, gap crossing, guidance of locomotion, and analogical reasoning. The answer to Gibson's central question is clearly "Yes": information about affordances is available in the environment.

7.7.3 Is Dynamic Touch a Special Case?

I have just described a case in which one can perceive affordances using dynamic touch.[16] I could just as easily have described another line of research in which one perceives the affordances of tools by dynamic touch (Carello et al. 1999; Cooper, Carello, and Turvey 1999; Wagman and Carello 2001, 2003; Wagman and Taylor 2004). In fact, the majority of the experimental research on perception of affordances done in recent years has been dynamic touch research. It is worth pausing to consider why that is the case. On the face of it, it might seem that perception by dynamic touch is a special case. Perception by dynamic touch is a matter of detecting information that centers on the wrist. This information is picked up by mechanoreceptors in the muscles and tendons of the hand, wrist, and arm. So in perception by dynamic touch, the information for perception is centered on the location of the action that is to be undertaken. Furthermore, because the eigenvalues of the inertial tensor carry information about resistance to rotation, the information available at the wrist is already "formatted" for use in action. Dynamic touch is the ideal kind of perception to call upon if one wants to explain links between perception and action without resorting to mental gymnastics. Because of the colocation of receptors and effectors and because no transformations are needed, the information available at the wrist is directly usable for controlling action.

The worry one might have at this point is that dynamic touch is the *only* kind of perception in which this is true. Compare dynamic touch with vision. It would seem that we see with our eyes and act on visual information with the rest of our bodies. This would put some distance between visual perception and action, leading to the worry that dynamic touch is a special case. There are two ways in which dynamic touch seems special. First, the information gathering and action occur at the same location of the body. Second, the information is correctly formatted for action. We have seen already that this second feature of dynamic touch does not differentiate it from vision: visual information *can* be in the right format for guiding action (e.g., τ and collision, discussed in chapter 6). Furthermore, if one understands vision correctly, it is simply false that vision and visually guided action happen at different places in the body. To see this, consider the distinction Gibson made between sensory modalities and perceptual systems (Gibson 1966). A sensory modality is defined anatomically, in terms of a collection of energy-specific receptor cells that make it up and the brain areas they are connected to. Perceptual systems, on the other hand, are defined functionally, in terms of information-gathering activity. Perceptual systems include energy-specific receptors and brain areas as proper parts, but also include parts of the organism that adjust, modify, or orient the receptors in active exploration. The human visual system, for example, includes the eyes and the canonical visual neural pathways along with the muscles and brain areas involved in eye movements and orientation of the head and neck, not to mention the head and neck themselves—all the parts of a human that take part in the activity of exploring the environment by looking around. In fact, Reed (1996) argues that perceptual systems are a variety of action system. Vision, the act of looking and seeing, is carried out by the entire visual system—an action system whose function is looking—and not merely by the visual sensory modality.

Understanding that perception is accomplished by perceptual systems, not sensory modalities, makes clear that dynamic touch is not a special case and not the only sense for which the information pickup and information-guided action are colocated. The same is true of the visual system, which uses information available in light to direct saccades and scan, focus, and track with the eyes, but also to control squinting, turn and rotate the head and neck to point both eyes at something of interest, crane to see over or around something, and so on. So information pickup by looking and the activity of looking *do* happen at the same place, but that place is the multipart visual system.

The haptic system, the system whose functions include dynamic touch, is not a special direct-perception-friendly perceptual system, but there is a way in which it differs from other perceptual systems. The difference, though, is in the sensory modalities commonly associated with the systems. The energy-specific receptors of the visual, auditory, and gustatory-olfactory systems are localized in organs on the head (in humans). In contrast, the receptors commonly associated with the sensory modality touch (the nerves in the skin, mechanoreceptors), are spread across the body. Furthermore, for these receptors to be activated, you have to actually be in contact with the object, usually acting upon it. This spatial diffuseness of the sensory modality and the fact that its receptors are only stimulated by contact with objects being sensed makes the tight connections between perception and action more obvious for dynamic touch than for other senses. But these special features of touch are special features of the sensory modality, not the perceptual system of which the sensory modality is a proper part, and it is perceptual systems, not sensory modalities, that pick up information about affordances.

7.8 Part III, Outro

This ends our introduction to Gibsonian ecological psychology, which I recommend as a nonrepresentational guide to discovery for radical embodied cognitive science. I have sketched a picture of animals as active agents, interacting with a world replete with information, and indeed generating information with their actions, including information about affordances. Perception and action, on this view, are tightly interconnected. Indeed, perception is a variety of action, and a good deal of action is done in the service of perception. The coupling of perception and action and the availability of information about affordances allow animals to guide their behavior without resorting to mental gymnastics. As noted in chapter 5, this theory of the nature of animals and their activity meshes perfectly with dynamical systems theory as a modeling tool. Again, as noted in chapter 5, this is unsurprising, given that dynamical systems theory was introduced as a modeling tool for psychology by ecological psychologists.

Here, then, is radical embodied cognitive science: Animals are active perceivers of and actors in an information-rich environment, and some of the information in the environment, the information to which animals are especially attuned, is about affordances. Unified animal–environment systems are to be modeled using the tools of dynamical systems theory. There is no need to posit representations of the environment inside the animal (or computations thereupon) because animals and environments are taken, both in theory and models, to be coupled.

In the final part of the book, I will examine how radical embodied cognitive science interacts with some traditional philosophical problems.