3
The Embodied Mind

What does it mean to say that concepts and reason are embodied? This chapter takes a first step toward answering that question. It takes up the role that the perceptual and motor systems play in shaping particular kinds of concepts: color concepts, basic-level concepts, spatial-relations concepts, and aspectual (event-structuring) concepts.

Any reasoning you do using a concept requires that the neural structures of the brain carry out that reasoning. Accordingly, the architecture of your brain's neural networks determines what concepts you have and hence the kind of reasoning you can do. Neural modeling is the field that studies which configurations of neurons carry out the neural computations that we experience as particular forms of rational thought. It also studies how such neural configurations are learned.

Neural modeling can show in detail one aspect of what it means for the mind to be embodied: how particular configurations of neurons, operating according to principles of neural computation, compute what we experience as rational inferences. At this point the vague question "Can reason make use of the sensorimotor system?" becomes the technically answerable question "Can rational inferences be computed by the same neural architecture used in perception or bodily movement?" We now know that, in some cases, the answer to this question is yes. Those cases will be discussed in this chapter.

How the Body and Brain Shape Reason

We have inherited from the Western philosophical tradition a theory of faculty psychology, in which we have a "faculty" of reason that is separate from and independent of what we do with our bodies. In particular, reason is seen as independent of perception and bodily movement. In the Western tradition, this autonomous capacity of reason is regarded as what makes us essentially human, distinguishing us from all other animals. If reason were not autonomous, that is, not independent of perception, motion, emotion, and other bodily capacities, then the philosophical demarcation between us and all other animals would be less clearly drawn. This view was formulated prior to the emergence of evolutionary theory, which shows that human capacities grow out of animal capacities.

The evidence from cognitive science shows that classical faculty psychology is wrong. There is no such fully autonomous faculty of reason separate from and independent of bodily capacities such as perception and movement. The evidence supports, instead, an evolutionary view, in which reason uses and grows out of such bodily capacities. The result is a radically different view of what reason is and therefore of what a human being is. This chapter surveys some of the evidence for the view that reason is fundamentally embodied.

These findings of cognitive science are profoundly disquieting in two respects. First, they tell us that human reason is a form of animal reason, a reason inextricably tied to our bodies and the peculiarities of our brains. Second, these results tell us that our bodies, brains, and interactions with our environment provide the mostly unconscious basis for our everyday metaphysics, that is, our sense of what is real.

Cognitive science provides a new and important take on an age-old philosophical problem, the problem of what is real and how we can know it, if we can know it. Our sense of what is real begins with and depends crucially upon our bodies, especially our sensorimotor apparatus, which enables us to perceive, move, and manipulate, and the detailed structures of our brains, which have been shaped by both evolution and experience.

Neural Beings Must Categorize

Every living being categorizes. Even the amoeba categorizes the things it encounters into food or nonfood, what it moves toward or moves away from. The amoeba cannot choose whether to categorize; it just does. The same is true at every level of the animal world. Animals categorize food, predators, possible mates, members of their own species, and so on. How animals categorize depends upon their sensing apparatus and their ability to move themselves and to manipulate objects.

Categorization is therefore a consequence of how we are embodied. We have evolved to categorize; if we hadn't, we would not have survived. Categorization is, for the most part, not a product of conscious reasoning. We categorize as we do because we have the brains and bodies we have and because we interact in the world the way we do.

The first and most important thing to realize about categorization is that it is an inescapable consequence of our biological makeup. We are neural beings. Our brains each have 100 billion neurons and 100 trillion synaptic connections. It is common in the brain for information to be passed from one dense ensemble of neurons to another via a relatively sparse set of connections. Whenever this happens, the pattern of activation distributed over the first set of neurons is too great to be represented in a one-to-one manner in the sparse set of connections. Therefore, the sparse set of connections necessarily groups together certain input patterns in mapping them across to the output ensemble. Whenever a neural ensemble provides the same output with different inputs, there is neural categorization.

To take a concrete example, each human eye has 100 million light-sensing cells, but only about 1 million fibers leading to the brain. Each incoming image must therefore be reduced in complexity by a factor of 100. That is, information in each fiber constitutes a "categorization" of the information from about 100 cells. Neural categorization of this sort exists throughout the brain, up through the highest levels of categories that we can be aware of. When we see trees, we see them as trees, not just as individual objects distinct from one another. The same with rocks, houses, windows, doors, and so on.

A small percentage of our categories have been formed by conscious acts of categorization, but most are formed automatically and unconsciously as a result of functioning in the world. Though we learn new categories regularly, we cannot make massive changes in our category systems through conscious acts of recategorization (though, through experience in the world, our categories are subject to unconscious reshaping and partial change). We do not, and cannot, have full conscious control over how we categorize. Even when we think we are deliberately forming new categories, our unconscious categories enter into our choice of possible conscious categories.

Most important, it is not just that our bodies and brains determine *that* we will categorize; they also determine what kinds of categories we will have and what their structure will be. Think of the properties of the human body that contribute to the peculiarities of our conceptual system. We have eyes and ears, arms and legs that work in certain very definite ways and not in others. We have a visual system, with topographic maps and orientation-sensitive cells, that provides structure for our ability to conceptualize spatial relations. Our abilities to move in the ways we do and to track the motion of other things give motion a major role in our conceptual system. The fact that we have muscles and use them to apply force in certain ways leads to the structure of our system of causal concepts. What is important is not just that we have bodies and that thought is somehow embodied. What is important is that the peculiar nature of our bodies shapes our very possibilities for conceptualization and categorization.

The Inseparability of Categories, Concepts, and Experience

Living systems must categorize. Since we are neural beings, our categories are formed through our embodiment. What that means is that the categories we form are *part of our experience!* They are the structures that differentiate aspects of our experience into discernible kinds. Categorization is thus not a purely intellectual matter, occurring after the fact of experience. Rather, the formation and use of categories is the stuff of experience. It is part of what our bodies and brains are constantly engaged in. We cannot, as some meditative traditions suggest, "get beyond" our categories and have a purely uncategorized and unconceptualized experience. Neural beings cannot do that.

What we call *concepts* are neural structures that allow us to mentally characterize our categories and reason about them. Human categories are typically conceptualized in more than one way, in terms of what are called *prototypes*. Each prototype is a neural structure that permits us to do some sort of inferential or imaginative task relative to a category. Typical-case prototypes are used in drawing inferences about category members in the absence of any special contextual information. Ideal-case prototypes allow us to evaluate category members relative to some conceptual standard. (To see the difference, compare the prototypes for the ideal husband and the typical husband.) Social stereotypes are used to make snap judgments, usually about people. Salient exemplars (well-known examples) are used for making probability judgments. (For a survey of kinds of conceptual prototypes, see A4, Lakoff 1987.) In short, prototype-based reasoning constitutes a large proportion of the actual reasoning that we do. Reasoning with prototypes is, indeed, so common that it is inconceivable that we could function for long without it.

Since most categories are matters of degree (e.g., tall people), we also have graded concepts characterizing degrees along some scale with norms of various kinds for extreme cases, normal cases, not quite normal cases, and so on. Such graded norms are described by what are called *linguistic hedges* (A4, Lakoff 1972), for example, *very, pretty, kind of, barely,* and so on. For the sake of imposing sharp distinctions, we develop what might be called *essence prototypes,* which conceptualize categories as if they were sharply defined and minimally distinguished from one another.

When we conceptualize categories in this way, we often envision them using a spatial metaphor, as if they were containers, with an interior, an exterior, and a boundary. When we conceptualize categories as containers, we also impose complex hierarchical systems on them, with some category-containers inside other category-containers. Conceptualizing categories as containers hides a great deal of category structure. It hides conceptual prototypes, the graded structures of categories, and the fuzziness of category boundaries.

In short, we form extraordinarily rich conceptual structures for our categories and reason about them in many ways that are crucial for our everyday functioning. All of these conceptual structures are, of course, neural structures in our brains. This makes them embodied in the trivial sense that any mental construct is realized neurally. But there is a deeper and more important sense in which our concepts are embodied. What makes concepts concepts is their inferential capacity, their ability to be bound together in ways that yield inferences. *An embodied concept is a neural structure that is actually part of, or makes use of, the sensorimotor system of our brains. Much of conceptual inference is, therefore, sensorimotor inference.*

If concepts are, as we believe, embodied in this strong sense, the philosophical consequences are enormous. The locus of reason (conceptual inference) would be the same as the locus of perception and motor control, which are bodily functions. If this seems like a radical claim, it is radical only from the perspective of faculty psychology, a philosophy that posits a radical separation between rational abilities and the sensorimotor system. It is not at all radical from the point of view of the brain, which is the joint locus of reason, perception, and movement. The question from the viewpoint of the brain is whether conceptual inference makes use of the same brain structures as perceptual motor inference. In other words, does reason piggyback on perception and motor control? From the perspective of the brain, the locus of all three functions, it would be quite natural if it did.

Realism, Inference, and Embodiment

The question of what we take to be real and the question of how we reason are inextricably linked. Our categories of things in the world determine what we take to be real: trees, rocks, animals, people, buildings, and so on. Our concepts determine how we reason about those categories. In order to function *realistically* in the world, our categories and our forms of reason must "work" very well together; our concepts must characterize the structure of our categories sufficiently well enough for us to function.

Mainstream Western philosophy adds to this picture certain claims that we will argue are false. Not trivially false, but so false as to drastically distort our understanding of what human beings are, what the mind and reason are, what causation and morality are, and what our place is in the universe. Here are those claims:

1. Reality comes divided up into categories that exist independent of the specific properties of human minds, brains, or bodies.
2. The world has a rational structure: The relationships among categories in the world are characterized by a *transcendent* or *universal* reason, which is independent of any peculiarities of human minds, brains, and bodies.
3. The concepts used by mind-, brain-, and body-free reason correctly characterize the mind-, brain-, and body-free categories of reality.
4. Human reason is the capacity of the human mind to use transcendent reason, or at least a portion of it. Human reason may be performed by the human brain, but the structure of human reason is defined by transcendent reason, independent of human bodies or brains. Thus, the structure of human reason is disembodied.
5. Human concepts are the concepts of transcendent reason. They are therefore defined independent of human brains or bodies, and so they too are disembodied.
6. Human concepts therefore characterize the objective categories of mind-, brain, and body-free reality. That is, the world has a unique, fixed category structure, and we all know it and use it when we are reasoning correctly.
7. What makes us essentially human is our capacity for disembodied reason.

8. Since transcendent reason is culture-free, what makes us essentially human is not our capacity for culture or for interpersonal relations.
9. Since reason is disembodied, what makes us essentially human is not our relation to the material world. Our essential humanness has nothing to do with our connection to nature or to art or to music or to anything of the senses.

Much of the history of mainstream Western philosophy consists of exploring variations on these themes and drawing out the consequences of these claims. A given philosopher may not hold all of these tenets in the strong form that we have stated them; however, together these claims form a picture of concepts, reason, and the world that any student of philosophy will be familiar with. If they are false, then large parts of the Western philosophical tradition and many of our most common beliefs have to be rethought.

These tenets were not adopted on the basis of empirical evidence. They arose instead out of a priori philosophy. Contemporary cognitive science calls this entire philosophical worldview into serious question on empirical grounds. Here is the reason why cognitive science has a crucial bearing on these issues.

At the heart of this worldview are tenets 4, 5, and 6—that human reason and human concepts are mind-, brain-, and body-free and characterize objective, external reality. If these tenets are false, the whole worldview collapses. Suppose human concepts and human reason are body- and brain-dependent. Suppose they are shaped as much by the body and brain as by reality. Then the body and brain are essential to our humanity. Moreover, our notion of what reality is changes. There is no reason whatever to believe that there is a disembodied reason or that the world comes neatly carved up into categories or that the categories of our mind are the categories of the world. If tenets 4, 5, and 6 are empirically incorrect, then we have a lot of rethinking to do about who we are and what our place is in the universe.

Embodied Concepts

In this chapter and the next, we will review some of the results of cognitive science research that bear on these issues. We will suggest, first, that human concepts are not just reflections of an external reality, but that they are crucially shaped by our bodies and brains, especially by our sensorimotor system. We will do so by looking at three kinds of concepts: color concepts, basic-level concepts, and spatial-relations concepts. After that, we will use studies of neural modeling to argue that certain human concepts and forms of conceptual reasoning make use of the sensorimotor system.

The philosophical stakes here are high. As we shall see in later chapters, these arguments have far-reaching implications for who we are and what our role in the world is.

Color Concepts

What could be simpler or more obvious than colors? The sky is blue. Fresh grass is green. Blood is red. The sun and moon are yellow. We see colors as inhering in things. Blue is in the sky, green in the grass, red in the blood, yellow in the sun. We see color, and yet it is false, as false as another thing we see, the moving sun rising past the edge of the stationary earth. Just as astronomy tells us that the earth moves around the sun, not the sun around a stationary earth, so cognitive science tells us that colors do not exist in the external world. Given the world, our bodies and brains have evolved to create color.

Our experience of color is created by a combination of four factors: wavelengths of reflected light, lighting conditions, and two aspects of our bodies: (1) the three kinds of color cones in our retinas, which absorb light of long, medium, and short wavelengths, and (2) the complex neural circuitry connected to those cones.

Here are some crucial things to bear in mind. One physical property of the surface of an object matters for color: its reflectance, that is, the relative percentages of high-, medium-, and low-frequency light that it reflects. That is a constant. But the actual wavelengths of light reflected by an object are not a constant. Take a banana. The wavelengths of light coming from the banana depend on the nature of the light illuminating it: tungsten or fluorescent, daylight on a sunny or a cloudy day, the light of dawn or dusk. Under different conditions the wavelengths of light coming from the banana will differ considerably, yet the color of the banana will be relatively constant; it will look pretty much the same. Color, then, is not just the perception of wavelength; color constancy depends on the brain's ability to compensate for variations in the light source. Moreover, there is not a one-to-one correspondence between reflectance and color; two different reflectances can both be perceived as the same red.

Another crucial thing to bear in mind is that light is not colored. Visible light is electromagnetic radiation, like radio waves, vibrating within a certain frequency range. It is not the kind of thing that could be colored. Only when this electromagnetic radiation impinges on our retinas are we able to see. We see a particular color when the surrounding lighting conditions are right, when radiation in a certain range impinges on our retina, and when our color cones absorb the radiation, producing an electrical signal that is appropriately processed by the neural circuitry of our brains. The qualitative experience that this produces in us is what we call "color."

One might suppose that color is an internal representation of the external reality of the reflectance properties of the surface of an object. If this were true, then the properties of colors and color categories would be representations of reflectances and categories of reflectances. But it is not true. Color concepts have internal structure, with certain colors being "focal." The category *red,* for instance, contains central red as well as noncentral, peripheral hues such as purplish red, pinkish red, and orangish red. The center-periphery structure of categories is a result of the neural response curves for color in our brains. Focal hues correspond to frequencies of maximal neural response. The internal structure of color categories is not out there in the surface reflectances. The same is true of the relationships among colors. The opposition between red and green or blue and yellow is a fact about our neural circuitry, not about the reflectance properties of surfaces. Color is not just the internal representation of external reflectance. And it is not a thing or a substance out there in the world.

To summarize, our color concepts, their internal structures, and the relationships between them are inextricably tied to our embodiment. They are a consequence of four interacting factors: lighting conditions, wavelengths of electromagnetic radiation, color cones, and neural processing. Colors as we see them, say, the red of blood or the blue of the sky, are not out there in the blood or the sky. Indeed, the sky is not even an object. It has no surface for the color to be in. And without a physical surface, the sky does not even have a surface reflectance to be detected as color. The sky is blue because the atmosphere transmits only a certain range of wavelengths of incoming light from the sun, and of the wavelengths it does transmit, it scatters some more than others. The effect is like a colored lightbulb that only lets certain wavelengths of light through the glass. Thus, the sky is blue for a very different reason than a painting of the sky is blue. What we perceive as blue does not characterize a single "thing" in the world, neither "blueness" nor wavelength reflectance.

Color concepts are "interactional"; they arise from the interactions of our bodies, our brains, the reflective properties of objects, and electromagnetic radiation. Colors are not objective; there is in the grass or the sky no greenness or blueness independent of retinas, color cones, neural circuitry, and brains.

Nor are colors purely subjective; they are neither a figment of our imaginations nor spontaneous creations of our brains.

The philosophical consequences are immediate. Since colors are not things or substances in the world, metaphysical realism fails. The meaning of the word *red* cannot be just the relation between the word and something in the world (say, a collection of wavelengths of light or a surface reflectance). An adequate theory of the conceptual structure of *red*, including an account of why it has the structure it has (with focal red, purplish red, orangish red, and so on) cannot be constructed solely from the spectral properties of surfaces. It must make reference to color cones and neural circuitry. Since the cones and neural circuitry are embodied, the internal conceptual properties of *red* are correspondingly embodied.

Subjectivism in its various forms—radical relativism and social constructionism—also fails to explain color, since color is created jointly by our biology and the world, not by our culture. This is not to say that color does not differ in its significance from culture to culture. It clearly does. Rather, color is a function of the world and our biology interacting.

Philosophically, color and color concepts make sense only in something like an embodied realism, a form of interactionism that is neither purely objective nor purely subjective. Color is also important for the "realism" of embodied realism. Evolution has worked with physical limitations: only certain wavelengths of light get through the atmosphere, only certain chemicals react to short, medium, and long wavelengths, and so on. We have evolved within these limitations to have the color systems we have, and they allow us to function well in the world. Plant life has been important to our evolution, and so the ability to place in one category the things that are green has apparent value for survival and flourishing. The same goes for blood and the color red, water and the sky and the color blue, and the sun and the moon and the color yellow. We have the color concepts we do because the physical limitations constraining evolution gave evolutionary advantages to beings with a color system that enabled them to function well in crucial respects.

Color, of course, does more than just help us recognize things in the world. It is an evolved aspect of the brain that plays many roles in our lives, cultural, aesthetic, and emotional. Thinking of color as merely the internal representation of the external reality of surface reflectance is not merely inaccurate; it misses most of the function of color in our lives.

At least since John Locke, philosophers have known that color is an interactional property of objects, what Locke called a "secondary quality" that does not exist in the object itself. Locke contrasted secondary qualities with "primary qualities," which were assumed to exist objectively in things independent of any perceiver. Primary qualities were seen as having metaphysical import, as determining what is real, while secondary qualities were seen as perceiver-dependent and therefore not constitutive of objective reality.

But giving up on color as a metaphysically real "primary quality" has profound philosophical consequences. It means abandoning the correspondence theory of truth, the idea that truth lies in the relationship between words and the metaphysically and objectively real world external to any perceiver. Since there is no color in the world in itself, a sentence like "Blood is red," which we all take to be true, would not be true according to the correspondence theory.

Since the correspondence theory of truth is the one thing many philosophers are not willing to give up, they go to extraordinary lengths to salvage it. Some attempt to see color as the internal representation of external reflectance of surfaces, and to say that "Blood is red" is true if and only if blood has such and such a surface reflectance. As we have seen, the same reasoning cannot work for "The sky is blue," since the sky cannot have a surface reflectance. Some philosophers have even been willing on these grounds to say that "The sky is blue" is false, granting that the sky has no surface reflectance but trying to keep the correspondence theory nonetheless. They claim that those of us who think that it is true that the sky is blue are simply being fooled by an optical illusion! Getting philosophers to give up on the correspondence theory of truth will not be easy. (For a thorough discussion of the details of the color debate in philosophy, see Thompson [A5, 1995]. For an account of the general philosophical implications of color research, see Varela, Thompson, and Rosch [C2, 1991], who argue, as we do, that color is interactional in nature and hence neither objective nor subjective. Defenses of objectivism and subjectivism can be found in Hilbert [A5, 1987, 1992] and Hardin [A5, 1988].)

As we are about to see, color is the tip of the iceberg. What Locke recognized as perceiver-dependence is a fully general phenomenon. Cognitive science and neuroscience suggest that the world as we *know* it contains *no* primary qualities in Locke's sense, because the qualities of things as we can experience and comprehend them depend crucially on our neural makeup, our bodily interactions with them, and our purposes and interests. For real human beings, the only realism is an embodied realism.

Basic-Level Categories

Why has metaphysical realism been so popular over the centuries? Why is it so common to feel that our concepts reflect the world as it is—that our categories of mind fit the categories of the world? One reason is that we have evolved to form at least one important class of categories that optimally fit our bodily experiences of entities and certain extremely important differences in the natural environment—what are called *basic-level categories.*

Our perceptual systems have no problem distinguishing cows from horses, goats from cats, or elephants from giraffes. In the natural world, the categories we distinguish among most readily are the folk versions of biological genera, namely, those that have evolved significantly distinct shapes so as to take advantage of different features of their environments. Go one level down in the biological hierarchy and it is a lot harder to distinguish one species of elephant from another (A4, Berlin et al. 1974). It's the same for physical objects. It's easy to tell cars from boats or trains, but a lot less easy to tell one kind of car from another.

Consider the categories *chair* and *car,* which are "in the middle" of the category hierarchies *furniture–chair–rocking chair* and *vehicle–car–sports car.* In the mid-1970s, Brent Berlin, Eleanor Rosch, Carolyn Mervis, and their coworkers discovered that such mid-level categories are cognitively "basic"—that is, they have a kind of cognitive priority, as contrasted with "superordinate" categories like *furniture* and *vehicle* and with "subordinate" categories like *rocking chair* and *sports car* (A4, Berlin et al. 1974; Mervis and Rosch 1981).

The Body-Based Properties of Basic-Level Categories

Basic-level categories are distinguished from superordinate categories by aspects of our bodies, brains, and minds: mental images, gestalt perception, motor programs, and knowledge structure. The basic level, as Berlin and Rosch found, is characterized by at least four conditions.

Condition 1: It is the highest level at which a single mental image can represent the entire category. For example, you can get a mental image of a chair. You can get mental images of other categories at the basic level such as tables and beds. But you cannot get a mental image of a general piece of furniture—a thing that is not a chair, table, or bed, but something more general. Similarly, you can get a mental image of a car. You can also get mental images of opposing categories at this level such as trains, boats, and planes. But you cannot get a mental image of a generalized vehicle—a thing that is not a car, train, boat, or plane, but a vehicle in general. The basic level is the highest level at which we have mental images that stand for the entire category.

Condition 2: It is the highest level at which category members have similarly perceived overall shapes. You can recognize a chair or a car by its overall shape. There is no overall shape that you can assign to a generalized piece of furniture or a vehicle so that you could recognize the category from that shape. The basic level is the highest level at which category members are recognized by gestalt perception (perception of overall shape).

Condition 3: It is the highest level at which a person uses similar motor actions for interacting with category members. You have motor programs for interacting with objects at the basic level—for interacting with chairs, tables, and beds. But you have no motor programs for interacting with generalized pieces of furniture.

Condition 4: It is the level at which most of our knowledge is organized. You have a lot of knowledge at the basic level. Think for a moment of all that you know about cars versus what you know about vehicles. You know a handful of things about vehicles in general, but a great many things about cars. You know much less about lower-level categories, unless you are an expert.

As a result of these characteristics, the basic level has other priorities over the superordinate and subordinate levels: It is named and understood earlier by children, enters a language earlier in its history, has the shortest primary lexemes, and is identified faster by subjects. The basic level also tends to be used in neutral contexts, that is, contexts in which there is no explicit indication of which level is most appropriate. From the perspective of an overall theory of the human mind, these are important properties of concepts and cannot be ignored.

The Philosophical Significance of the Basic Level

The philosophical significance of these results follows directly. First, the division between basic-level and nonbasic-level categories is body-based, that is, based on gestalt perception, motor programs, and mental images. Because of this, classical metaphysical realism cannot be right, since the properties of categories are mediated by the body rather than determined directly by a mind-independent reality.

Second, the basic level is that level at which people interact optimally with their environments, given the kinds of bodies and brains they have and the kinds of environments they inhabit. How is this possible? The best answer we know, suggested by Tversky and Hemenway (A4, 1984), is that the properties that make for basic-level categories are responses to the part-whole structure of physical beings and objects. Gestalt perception is about overall part-whole structure, as is mental imagery. The use of motor schemas to interact with objects depends significantly on their overall part-whole structure. Moreover, the functions something can perform, and hence what we know about it, likewise depend to a significant degree on part-whole structure. That is why there is a basic-level category structure with respect to which we can function optimally.

Third, basic-level categorization tells us why metaphysical realism makes sense for so many people, where it seems to work, and where it goes wrong. Metaphysical realism seems to work primarily at the basic level. If you look only at examples of basic-level categories, at the level of category where we interact optimally with the world, then it appears as if our conceptual categories fit the categories of the world. If you look at categories at other levels, it does not (A4, Berlin et al. 1974). It is not surprising, therefore, that philosophical discussions about the relationship between our categories and things in the world tend to use basic-level examples. Philosophical examples like "The cat is on the mat" or "The boy hit the ball" typically use basic-level categories like *cat, mat, boy,* and *ball* or basic-level substances like *water* and *gold.* It is no accident that philosophers do not try to make their argument with things farther down on the biological taxonomy: brown-capped chickadees, brown-headed nuthatches, Bewick's wrens, bushtits, and so on.

The basic level, of course, is not just about objects. There are basic-level actions, actions for which we have conventional mental images and motor programs, like swimming, walking, and grasping. We also have basic-level social concepts, like families, clubs, and baseball teams, as well as basic-level social actions, like arguing. And there are basic emotions, like happiness, anger, and sadness.

Fourth, the properties of the basic level explain an important aspect of the stability of scientific knowledge. For basic-level physical objects and basic-level actions or relations, the link between human categories and divisions of things in the world is quite accurate. We can think of scientific instruments as extending these basic-level abilities to perceive, image, and intervene. Telescopes, microscopes, cameras, and delicate probing instruments of all sorts extend our capacity for basic-level perception, imaging, and intervention. Such instruments allow us to greatly extend the range of our categories of mind to fit important distinctions in the world.

For basic-level categories, the idea that our categories of mind fit the categories of the world is not that far off. When our basic-level capacities are extended by scientific instrumentation, our ability to select useful real-world divisions is improved. Basic-level categories are the source of our most stable knowledge, and the technological capacity to extend them allows us to extend our stable knowledge.

In summary, our categories arise from the fact that we are neural beings, from the nature of our bodily capacities, from our experience interacting in the world, and from our evolved capacity for basic-level categorization—a level at which we optimally interact with the world. Evolution has not required us to be as accurate above and below the basic level as at the basic level, and so we are not.

There is a reason why our basic-level categorization and evolution match up. In the natural world, basic-level categories of organisms are genera. That means that they are for the most part determined by their overall part-whole structure. The part-whole structure of a class of organisms is, significantly, what determines whether it will survive and function well in a given environment. Thus, part-whole structure determines the natural categories of existing genera. And it is what our perceptual and motor systems have evolved to recognize at the basic level. That is why we have tended over our evolutionary history to function optimally in our basic-level interactions.

Though the facts of basic-level categorization do not fit metaphysical realism, they do provide us with the basis for embodied realism, which is an improvement over metaphysical realism in that it provides a link between our ideas and the world, at least at the level that matters most for our survival. The facts of basic-level categorization also remind us that our bodies contribute to our sense of what is real.

We turn next to spatial-relations concepts. These too are embodied. They have to be, because they allow us to negotiate space, to function in it as well as to conceptualize it and talk about it.

Spatial-Relations Concepts

Spatial-relations concepts are at the heart of our conceptual system. They are what make sense of space for us. They characterize what spatial form is and define spatial inference. But they do not exist as entities in the external world. We do not see spatial relations the way we see physical objects.

We do not see nearness and farness. We see objects where they are and we attribute to them nearness and farness from some landmark. The relations *in front of* and *in back of* are imposed by us on space in a complex way. When you go *in the front* of a church, you find yourself *in the back* of it. Or take the concept *across*. Suppose you are to row across a round pond. If you row "straight across" it (at a 90-degree angle from the shore), you have certainly rowed across it. If you row at a 45-degree angle, it is not as clear. If you row at a 15-degree angle, certainly not. Here, what counts as *across* varies with the shape of the area crossed and the angle of crossing and is also a matter of degree. Spatial-relations concepts are not simple or straightforward, and they vary considerably from language to language.

We use spatial-relations concepts unconsciously, and we impose them via our perceptual and conceptual systems. We just automatically and unconsciously "perceive" one entity as *in*, *on*, or *across from* another entity. However, such perception depends on an enormous amount of automatic unconscious mental activity on our part. For example, to see a butterfly as *in* the garden, we have to project a nontrivial amount of imagistic structure onto a scene. We have to conceptualize the boundaries of the garden as a three-dimensional container with an interior that extends into the air. We also have to locate the butterfly as a figure (or *trajector*) relative to that conceptual container, which serves as a ground (or *landmark*). We perform such complex, though mundane, acts of imaginative perception during every moment of our waking lives.

Most spatial relations are complexes made up of elementary spatial relations. English *into* is a composite of the English elementary spatial relations *in* and *to*. English *on* in its central sense is a composite of *above, in contact with,* and *supported by*. Each of these is an elementary spatial relation. Elementary spatial relations have a further internal structure consisting of an *image schema, a profile,* and a *trajector-landmark structure.*

To see what these terms mean, let us take a simple example.

The Container Schema

English *in* is made up of a container schema (a bounded region in space), a profile that highlights the interior of the schema, and a structure that identifies the boundary of the interior as the landmark (LM) and the object overlapping with the interior as a trajector (TR). In "Sam is in the house," the house is the landmark (LM) relative to which Sam, the trajector (TR), is located.

Spatial relations also have built-in spatial "logics" by virtue of their image-schematic structures. Figure 3.1 illustrates the spatial logic built into the container schema:

- Given two containers, *A* and *B,* and an object, *X,* if *A* is *in B* and *X* is *in A,* then *X* is *in B.*

We don't have to perform a deductive operation to compute this. It is self-evident simply from the image in Figure 3.1.

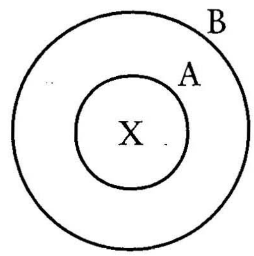

X is in A
A is in B
∴ X is in B

FIGURE 3.1 (Container Schema Logic)

A container schema has the following structure: an inside, a boundary, and an outside. This is a gestalt structure, in the sense that the parts make no sense without the whole. There is no inside without a boundary and an outside, no outside without a boundary and an inside, and no boundary without sides. The structure is topological in the sense that the boundary can be made larger, smaller, or distorted and still remain the boundary of a container schema.

A container schema, like any other image schema, is conceptual. Such a container schema can, however, be physically instantiated, either as a concrete object, like a room or a cup, or as bounded region in space, like a basketball court or a football field.

Suppose the boundary of a container schema is physically instantiated in a concrete object, say, a box. A physical boundary can impose forceful and visual constraints: It can protect the container's contents, restrict their motion, and render them inaccessible to vision. It is important to distinguish a purely conceptual schema from a physically instantiated one; they have different properties.

Container schemas, like other image schemas, are cross-modal. We can impose a conceptual container schema on a visual scene. We can impose a container schema on something we hear, as when we conceptually separate out one part of a piece of music from another. We can also impose container schemas on our motor movements, as when a baseball coach breaks down a batter's swing into component parts and discusses what goes on "inside" each part.

The Source-Path-Goal Schema

As with a container schema, there is a spatial logic built into the source-path-goal schema (Figure 3.2). The source-path-goal schema has the following elements (or "roles"):

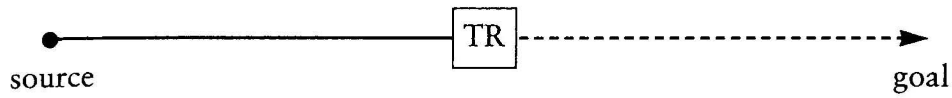

FIGURE 3.2 (Source-Path-Goal Schema)

A trajector that moves
A source location (the starting point)
A goal, that is, an intended destination of the trajector
A route from the source to the goal
The actual trajectory of motion
The position of the trajector at a given time
The direction of the trajector at that time
The actual final location of the trajector, which may or may not be the intended destination

Extensions of this schema are possible: a vehicle, the speed of motion, obstacles to motion, forces that move one along a trajectory, additional trajectors, and so on.

This schema is topological in the sense that a path can be expanded, shrunk, or deformed and still remain a path. Trajectories are imaginative insofar as they are not entities in the world; they are conceptualized as a linelike "trail" left by an object as it moves and projected forward in the direction of motion.

As with the container schema, we can form spatial relations from this schema by the addition of profiling (also called *highlighting*) and a trajector-landmark relation. The concept expressed by *to* profiles the goal and identifies it as the landmark relative to which the motion takes place. The concept expressed by *from* profiles the source, taking the source as the landmark relative to which the motion takes place.

The source-path-goal schema also has an internal spatial "logic" and built-in inferences:

- If you have traversed a route to a current location, you have been at all previous locations on that route.
- If you travel from *A* to *B* and from *B* to *C*, then you have traveled from *A* to *C*.

- If there is a direct route from *A* to *B* and you are moving along that route toward *B*, then you will keep getting closer to *B*.
- If *X* and *Y* are traveling along a direct route from *A* to *B* and *X* passes *Y*, then *X* is further from *A* and closer to *B* than *Y* is.
- If *X* and *Y* start from *A* at the same time moving along the same route toward *B* and if *X* moves faster than *Y*, then *X* will arrive at *B* before *Y*.

Our most fundamental knowledge of motion is characterized by the source-path-goal schema, and this logic is implicit in its structure. Many spatial-relations concepts are defined using this schema and depend for their meaning on its inherent spatial logic, for example, *toward, away, through,* and *along*.

Bodily Projections

Bodily projections are especially clear instances of the way our bodies shape conceptual structure. Consider examples such as *in front of* and *in back of*. The most central senses of these terms have to do with the body. We have inherent fronts and backs. We see from the front, normally move in the direction the front faces, and interact with objects and other people at our fronts. Our backs are opposite our fronts; we don't directly perceive our own backs, we normally don't move backwards, and we don't typically interact with objects and people at our backs.

We project fronts and backs onto objects. What we understand as the front of a stationary artifact, like a TV or a computer or a stove, is the side we normally interact with using our fronts. What we take to be the front of a moving object like a car is that part of the object that "faces" the direction in which it normally moves. We project fronts onto stationary objects without inherent fronts such as trees or rocks. English speakers project fronts onto such objects so the front faces the speaker. In other languages (e.g., Hausa), speakers project fronts onto such objects in the opposite direction, facing away from the speaker.

The concepts *front* and *back* are body-based. They make sense only for beings with fronts and backs. If all beings on this planet were uniform stationary spheres floating in some medium and perceiving equally in all directions, they would have no concepts of *front* or *back*. But we are not like this at all. Our bodies are symmetric in some ways and not in others. We have faces and move in the direction in which we see. Our bodies define a set of fundamental spatial orientations that we use not only in orienting ourselves, but in perceiving the relationship of one object to another.

When we perceive a cat as being in front of a car or behind a tree, the spatial relationships *in front of* and *behind*, between cat and car or between cat and tree, are not objectively there in the world. The spatial relation is not an entity in our visual field. The cat is behind the tree or in front of the car only relative to our capacity to project fronts and backs onto cars and trees and to impose relations onto visual scenes relative to such projections. In this way, perceiving the cat as being behind the tree requires an imaginative projection based on our embodied nature.

Compared to certain other languages, English is relatively impoverished in its use of bodily projections to conceptualize spatial relations. By contrast, languages of the Otomonguean family, such as Mixtec, use bodily projections as their primary means of characterizing spatial relations (A1, Brugman 1985).

For example, in Mixtec, there is no unitary concept or word corresponding to English *on*. The range of cases covered by English *on* is instead described by using body-part projections. Suppose you want to say "He is on top of the hill." You say the equivalent of "He is located head hill." If you want to say "I was on the roof of the house," you say the Mixtec equivalent of "I was located animal-back house," in which an animal back, being canonically oriented horizontally, is projected onto the house. If you want to say "I am sitting on the branch of the tree," you say the equivalent of "I am sitting arm tree."

One way in which languages differ is that, while some have mainly body-centered relations like *in front of*, others have mainly externally based relations, like *to the north of*, and still others have mixed systems (A8, Levinson 1992–present).

Other Image Schemas and Elements of Spatial Relations

The study of spatial-relations concepts within cognitive linguistics has revealed that there is a relatively small collection of primitive image schemas that structure systems of spatial relations in the world's languages. Here are some examples, without the full detail given above: part-whole, center-periphery, link, cycle, iteration, contact, adjacency, forced motion (e.g., pushing, pulling, propelling), support, balance, straight-curved, and near-far. Orientations also used in the spatial-relations systems of the world's languages include vertical orientation, horizontal orientation, and front-back orientation. (For a fuller discussion see A4, Lakoff 1987, case study 2; A1, Johnson 1987; A8, Talmy 1983; and B2, Regier 1996.)

One of the important discoveries of cognitive science is that the conceptual systems used in the world's languages make use of a relatively small number of basic image schemas, though the range of complex spatial relations that can be built out of these schemas is very large. As we shall see when we get to the discussion of conceptual metaphor, the spatial logics of these body-based image schemas are among the sources of the forms of logic used in abstract reason.

The Embodied Nature of Spatial-Relations Concepts

Spatial-relations concepts are embodied in various ways. Bodily projections are obviously based on the human body. Concepts like *front* and *back* and those in Mixtec arise from the body, depend on the body, and would not exist if we did not have the kinds of bodies we have. The same is true of fundamental force-dynamic schemas: pushing, pulling, propelling, supporting, and balance. We comprehend these through the use of our body parts and our ability to move them, especially our arms, hands, and legs.

Other image schemas are also comprehended through the body. Our bodies are containers that take in air and nutrients and emit wastes. We constantly orient our bodies with respect to containers—rooms, beds, buildings. We spend an inordinate amount of time putting things in and taking things out of containers. We also project abstract containers onto areas in space, as when we understand a swarm of bees as being *in* the garden. Similarly, every time we see something move, or move ourselves, we comprehend that movement in terms of a source-path-goal schema and reason accordingly.

These forms of embodiment arise from the way we schematize our own bodies and things we interact with daily (C2, Gallagher 1995). We will refer to this as *phenomenological embodiment.* But there is also *neural embodiment,* as we saw in the case of color. Neural embodiment characterizes the neural mechanisms that give rise to concepts—for example, the neural circuitry connected to the color cones that brings color into existence and characterizes the structure of color categories. These neural mechanisms explain why color categories have many of the phenomenological properties they have.

We do not yet know the exact neural mechanisms that give rise to spatial-relations concepts, but a beginning has been made. A computational neural model has been constructed that characterizes certain image schemas neurally, explains why they should exist, and accounts for their topological and orientational properties. Let us now turn to this research.

The Neural Modeling of Spatial and Motor Concepts

As we mentioned above, much of the Western philosophical tradition assumes a form of faculty psychology, according to which we have a faculty of reason separate from our faculties of perception and bodily movement. Concepts and the forms of reason based on them are assumed to be purely part of the faculty of reason. Perception may inform reason, and movement may be a consequence of reason, but in the tradition no aspect of perception or movement is *part* of reason.

Consequently, there is assumed to be an absolute dichotomy between *per*ception and *con*ception. While *per*ception has always been accepted as bodily in nature, just as movement is, *con*ception—the formation and use of concepts—has traditionally been seen as purely mental and wholly separate from and independent of our abilities to perceive and move.

We have already begun to get intimations that this picture is false. We have seen that basic-level concepts depend on motor movement, gestalt perception, and mental imagery, which is carried out in the visual system of the brain. We have seen that color is anything but purely mental, that our color concepts are intimately shaped not merely by perception as a faculty of mind but by such physical parts of our bodies as color cones and neural circuitry. And we have seen that spatial-relations concepts like *front* and *back* are not characterized by some abstract, disembodied mental capacity but rather in terms of bodily orientation. In these cases, the body is not merely somehow involved in conceptualization but is shaping its very nature.

Embodiment Not as Realization but as Shaping

What is the view that the mind is disembodied? It is the view that the contents of mind, the actual concepts, are not crucially shaped or given any significant inferential content by the body. It is the view that concepts are formal in nature and arise from the mind's capacity to generate formal structure in such a way as to derive further, inferred, formal structures. Advocates of the disembodied mind will, of course, say that conceptual structure must have a neural *realization* in the brain, which just *happens* to reside in a body. But they deny that anything about the body is essential for characterizing what concepts are.

The claim that the mind is embodied is, therefore, far more than the simple-minded claim that the body is needed if we are to think. Advocates of the disembodied-mind position agree with *that*. Our claim is, rather, that the very properties of concepts are created as a result of the way the brain and body are structured and the way they function in interpersonal relations and in the physical world.

The embodied-mind hypothesis therefore radically undercuts the *per*ception/*con*ception distinction. In an embodied mind, it is conceivable that the same neural system engaged in *per*ception (or in bodily movement) plays a central role in *con*ception. That is, the very mechanisms responsible for perception, movements, and object manipulation could be responsible for conceptualization and reasoning. Indeed, in recent neural modeling research, models of perceptual mechanisms and motor schemas can actually do *conceptual* work in language learning and in reasoning. This is a startling result. It flies in the face of time-honored philosophical theories of faculty psychology and their recent reincarnation in strong modularity theories of mind and language, each of which insists on a separation of the mechanisms for perception and conception.

Neural Modeling as an Existence Proof for the Embodiment of Mind

As yet, we do not have any strong neurophysiological evidence, say from PET scan or functional MRI results, that the same neural mechanisms used in perception and movement are also used in abstract reasoning. What we do have is an existence proof that this is possible and good reasons to believe that it is plausible. The existence proof comes from the field of neural modeling, and it comes in the following form. A neural model of a perceptual or motor mechanism is constructed, and that very same mechanism is used for conceptual tasks as well. The conceptual tasks are of two sorts: (1) learning the structure of a semantic field of lexical items so as to get the relationships among the lexical items correct and (2) performing abstract inferences.

These models are existence proofs in the sense that they show that neural structures that can carry out sensorimotor functions in the brain can in principle do both jobs at once—the job of perception or motor control, on the one hand, and the job of conceptualizing, categorizing, and reasoning, on the other.

What is particularly impressive about these models is that they are computational. The field of computational neuroscience is concerned not merely with *where* the neural computations are done but with *how*, that is, with precise neural computational mechanisms that perform sensorimotor operations and that carry out conceptualizing, categorizing, reasoning, and language learning. Each of the models we will discuss does such jobs in detail.

Models have been constructed for three kinds of concepts:

1. Spatial-relations concepts, for example, those named by English words like *in, on, over, through,* and *under.*

2. Concepts of bodily movement, represented by verbs like *grasp, pull, lift, tap,* and *punch.*
3. Concepts indicating the structure of actions or events (what linguists call *aspectual concepts*) like *starting, stopping, resuming, continuing, finishing,* including those indicated grammatically as in process (in English *is/are* plus the verb stem plus *-ing: is running*) or completed (*has/have* plus the verb stem plus *-ed: has lifted*).

Since these concepts are about what the body does, namely, perceive and move, one would expect that what the body actually does should shape these concepts. In particular:

- Since spatial-relations concepts are about space, it should not be surprising if our capacities for vision and negotiating space are used in constituting spatial-relations concepts and their logics.
- Since concepts of bodily movement are about motor actions, it should not be surprising if our motor schemas and parameters of bodily movement structure those concepts and their logics.
- Since moving the body is our most common form of action, it should not be surprising if the general structure of control schemas for bodily movements should be used to characterize aspectual structure, the structure we find in actions and events in general.

These models suggest some things that make eminently good sense: The visual systems of our brains are used in characterizing spatial-relations concepts. Our actual motor schemas and motor synergies are involved in what verbs of motor movement mean. And the general form of motor control gives general form to all our actions and the events we perceive. The point is this: In such models, there is no absolute perceptual/conceptual distinction, that is, the conceptual system makes use of important parts of sensorimotor system that impose crucial conceptual structure.

The Three Models

The three models we are about to discuss are highly complex, and we can give only a very brief overview of them here. A more detailed discussion is found in the Appendix. (For full discussions of all the technical details, see B2, Regier 1996; Bailey 1997; and Narayanan 1997a, b.)

Regier's Model for Learning Spatial-Relations Terms

Terry Regier (B2, 1996) constructed a neural model for learning spatial-relations terms in the world's languages. Given a model of retinal input with geometric figures in various spatial configurations together with a linguistic description correctly describing the configuration in a given language, the neural model was to learn the system of spatial-relations concepts and terms so that it could correctly categorize and label novel configurations. It was to do this both in cases of static spatial configurations (e.g., *on*) and in cases involving motion (e.g., *onto*). The model learned using no negative evidence, that is, no incorrectly labeled cases, only correctly labeled ones.

Here is the idea behind the model: Though spatial-relations terms differ wildly across the world's languages, they have to categorize using structures found in the visual system of the brain. Spatial-relations concepts should therefore depend on neural structures found in the brain's visual system. Consequently, Regier's model is designed to make maximal use of the types of structures known to exist in the human visual system. Regier's major insights were, first, that topographic maps of the visual field should be instrumental in the computation of image schemas that have topological properties (e.g., the container schema); second, that orientation-sensitive cell assemblies should be able compute the orientational aspects of spatial concepts that rely on bodily orientation (e.g., *above*); third, that center-surround receptive fields should be crucial to characterizing concepts like *contact*; and finally that the "filling-in" architecture discovered by Ramachandran and Gregory (B1, 1991) should play a central role in characterizing the notion of *containment*.

The Regier model is simultaneously both perceptual and conceptual. By virtue of the way the perceptual mechanisms work, it accomplishes the conceptual task of categorizing spatial configurations adequately to fit the conceptual distinctions and contrasts among spatial-relations terms in natural languages. It thereby gives us some insight into how the neural structures in the brain that do perceptual work might be recruited to do conceptual work as well.

Bailey's Model for Learning Verbs of Hand Motion

David Bailey's model (B2, 1997) learns not only how to categorize and name hand motions in the world's languages but also how to use those verbs correctly to give orders to produce the corresponding hand motion in a computer model of the body. At the heart of Bailey's model are models of high-level motor-control schemas that operate dynamically in time to control motor synergies—subcortical neural circuits that act automatically to produce small, low-level movements. These synergies provide the parameters used by the motor-control schemas, called *X-schemas* (for *executing schemas*).

The idea behind the model is this: Verbs of hand action differ considerably around the world, categorizing actual hand actions in markedly different ways from language to language. Yet the categorization should depend on the actual motor schemas used in moving things with the hand and on parameters given by actual motor synergies. Thus, the actual motor mechanisms should also be doing the conceptual work of categorizing actions for the purpose of naming them. The success of the Bailey model suggests how neural circuitry used for motor control can be recruited for conceptual purposes.

Narayanan's Model of Motor Schemas, Linguistic Aspect, and Metaphor

Srini Narayanan (B2, 1997a, b), working with Bailey on modeling motor schemas, discovered that all motor schemas have the same high-level control structure:

Getting into a state of readiness
The initial state
The starting process
The main process (either instantaneous or prolonged)
An option to stop
An option to resume
An option to iterate or continue the main process
A check to see if a goal has been met
The finishing process
The final state

This should come as no surprise. Any high-level motor activity you undertake, from scratching your head to turning on a light switch to sipping a cup of tea, will have this structure. (It is actually more complex; for the sake of a brief presentation, we have simplified it a bit.) Narayanan then constructed a model of this control structure so that it could be structured separately from the individual special cases (e.g., lifting a cup). That permitted a great simplification in characterizing neural control structures.

Linguists should recognize this model immediately. It characterizes the semantic structure of events in general, what linguists call *aspect.* Any action one undertakes, whether a bodily movement or a more abstract undertaking like planning what to have for dinner, has such a structure. And each language has a linguistic means of highlighting aspects of such a structure. In English, for example, the present imperfect form of the verb (*is/are* plus the present stem of the verb plus *-ing*, as in *is walking*) focuses on the main process as it is happening.

Aspect—the general structure of events—has a conceptual structure and a logic. What Narayanan discovered was that exactly the same neural structure that can perform motor control also characterizes the conceptual structure of linguistic aspect, and the same neural mechanism that can control bodily movements can perform logical inferences about the structure of actions in general.

Narayanan devised an ingenious way to test whether his model of general high-level motor control could handle purely abstract inferences, inferences having nothing to do with bodily movement. He constructed a neural model of conceptual metaphor and then found cases in which body-based metaphors were used in an abstract domain, in this case, international economics. Prominent newspapers and journals use such metaphors every day in economic news reports; for example, "India loosened its stranglehold on business," "France fell into a recession and Germany pulled it out." Narayanan then showed that models of the motor schemas for physical actions can—under metaphoric projection—perform the appropriate abstract inferences about international economics.

The Body in the Mind

Each of these neural modeling studies constitutes an existence proof. Spatial-relations concepts *can* be represented and spatial-relations terms learned on the basis of neural perceptual apparatus in the brain's visual system (topographic maps of the visual field, orientation-sensitive cells, and so on). Concepts for hand motions *can* be represented and hand-motion terms learned on the basis of detailed models of high-level motor control and motor synergies. Aspectual concepts that characterize the structure of events *can* be adequately represented in terms of general motor-control schemas, and abstract reasoning using those schemas can be carried out using neural motor-control simulations. None of this proves that people actually use those parts of the brain involved in perception and motor control to do such reasoning, but it is in principle possible. At present, these systems that use neural models of motor-control schemas are the only ones capable of carrying out the given tasks.

Now that we know that there can be such a direct embodiment of reason, the question becomes an empirical one, to be settled in experimental neuroscience, not in the arena of philosophical argumentation. The evidence so far favors embodied cognition, and there are general reasons for believing that something like the embodied cognition theory will turn out to be true.

Brains tend to optimize on the basis of what they already have, to add only what is necessary. Over the course of evolution, newer parts of the brain have built on, taken input from, and used older parts of the brain. Is it really plausible that, if the sensorimotor system can be put to work in the service of reason, the brain would build a whole new system to duplicate what it could do already?

Regier has shown that the topological properties of spatial relations can be explained on the basis of the topological properties arising from applying center-surround receptive fields and Ramachandran's filling-in process to topographic maps of the visual field. Is it really plausible that the brain would develop another, nonvisual system with the same topological properties to reason about space, when we obviously already use vision to get around in space?

Narayanan has shown that the neural structure of motor control must already have all the capacities necessary to characterize aspect (the structure of events) and its logic. If the brain can reason about actions using the structure already present to perform actions, is it plausible that the brain would build another system to do the same thing? And if it did, is it plausible that it would take a significantly different neural form?

From a biological perspective, it is eminently plausible that reason has grown out of the sensory and motor systems and that it still uses those systems or structures developed from them. This explains why we have the kinds of concepts we have and why our concepts have the properties they have. It explains why our spatial-relations concepts should be topological and orientational. And it explains why our system for structuring and reasoning about events of all kinds should have the structure of a motor-control system.

It is only from a conservative philosophical position that one would want to believe in the old faculty psychology—in the idea that the human mind has nothing about it that animals share, that reason has nothing about it that smells of the body.

Philosophically, the embodiment of reason via the sensorimotor system is of great importance. It is a crucial part of the explanation of why it is possible for our concepts to fit so well with the way we function in the world. They fit so well because they have evolved from our sensorimotor systems, which have in turn evolved to allow us to function well in our physical environment. The embodiment of mind thus leads us to a philosophy of embodied realism. Our concepts cannot be a direct reflection of external, objective, mind-free reality because our sensorimotor system plays a crucial role in shaping them. On the other hand, it is the involvement of the sensorimotor system in the conceptual system that keeps the conceptual system very much in touch with the world.