2

A Brief Introduction to the Cognitive Science of the Embodied Mind

The Cognitive Unconscious

Perhaps the most fundamental, and initially the most startling, result in cognitive science is that most of our thought is unconscious—that is, fundamentally inaccessible to our direct, conscious introspection. Most everyday thinking occurs too fast and at too low a level in the mind to be thus accessible. Most cognition happens backstage. That includes mathematical cognition.

We all have systems of concepts that we use in thinking, but we cannot consciously inspect our conceptual inventory. We all draw conclusions instantly in conversation, but we cannot consciously look at each inference and our own inference-drawing mechanisms while we are in the act of inferring on a massive scale second by second. We all speak in a language that has a grammar, but we do not consciously put sentences together word by word, checking consciously that we are following the grammatical rules of our language. To us, it seems easy: We just talk, and listen, and draw inferences without effort. But what goes on in our minds behind the scenes is enormously complex and largely unavailable to us.

Perhaps the most startling realization of all is that we have unconscious memory. The very idea of an unconscious memory seems like a contradiction in terms, since we usually think of remembering as a conscious process. Yet hundreds of experimental studies have confirmed that we remember without being aware that we are remembering—that experiences we don't recall do in fact have a detectable and sometimes measurable effect on our behavior. (For an excellent overview, see Schacter, 1996.)

What has not been done so far is to extend the study of the cognitive unconscious to mathematical cognition—that is, the way we implicitly understand mathematics as we do it or talk about it. A large part of unconscious thought involves automatic, immediate, implicit rather than explicit understanding—making sense of things without having conscious access to the cognitive mechanisms by which you make sense of things. Ordinary everyday mathematical sense-making is not in the form of conscious proofs from axioms, nor is it always the result of explicit, conscious, goal-oriented instruction. Most of our everyday mathematical understanding takes place without our being able to explain exactly what we understood and how we understood it. Indeed, when we use the term "understanding" throughout this book, this automatic unconscious understanding is the kind of understanding we will be referring to, unless we say otherwise.

Therefore, this book is not about those areas of cognitive science concerned with conscious, goal-oriented mathematical cognition, like conscious approaches to problem solving or to constructing proofs. Though this book may have implications for those important fields, we will not discuss them here.

Our enterprise here is to study everyday mathematical understanding of this automatic unconscious sort and to ask a crucial question: How much of mathematical understanding makes use of the same kinds of conceptual mechanisms that are used in the understanding of ordinary, nonmathematical domains? Are the same cognitive mechanisms that we use to characterize ordinary ideas also used to characterize mathematical ideas?

We will argue that a great many cognitive mechanisms that are not specifically mathematical are used to characterize mathematical ideas. These include such ordinary cognitive mechanisms as those used for the following ordinary ideas: basic spatial relations, groupings, small quantities, motion, distributions of things in space, changes, bodily orientations, basic manipulations of objects (e.g., rotating and stretching), iterated actions, and so on.

To be more specific, we will suggest that:

- Conceptualizing the technical mathematical concept of a class makes use of the everyday concept of a collection of objects in a bounded region of space.
- Conceptualizing the technical mathematical concept of recursion makes use of the everyday concept of a repeated action.

- Conceptualizing the technical mathematical concept of complex arithmetic makes use of the everyday concept of rotation.
- Conceptualizing derivatives in calculus requires making use of such everyday concepts as motion, approaching a boundary, and so on.

From a nontechnical perspective, this should be obvious. But from the technical perspective of cognitive science, one must ask:

Exactly what everyday concepts and cognitive mechanisms are used in exactly what ways in the unconscious conceptualization of technical ideas in mathematics?

Mathematical idea analysis, as we will be developing it, depends crucially on the answers to this question. Mathematical ideas, as we shall see, are often grounded in everyday experience. Many mathematical ideas are ways of mathematicizing ordinary ideas, as when the idea of a derivative mathematicizes the ordinary idea of instantaneous change.

Since the cognitive science of mathematics is a new discipline, not much is known for sure right now about just how mathematical cognition works. Our job in this book is to explore how the general cognitive mechanisms used in everyday nonmathematical thought can create mathematical understanding and structure mathematical ideas.

Ordinary Cognition and Mathematical Cognition

As we saw in the previous chapter, it appears that all human beings are born with a capacity for subitizing very small numbers of objects and events and doing the simplest arithmetic—the arithmetic of very small numbers. Moreover, if Dehaene (1997) is right, the inferior parietal cortex, especially the angular gyrus, "plays a crucial role in the mental representation of numbers as quantities" (p. 189). In other words, there appears to be a part of the brain innately specialized for a sense of quantity—what Dehaene, following Tobias Dantzig, refers to as "the number sense."

But there is a lot more to mathematics than the arithmetic of very small numbers. Trigonometry and calculus are very far from "three minus one equals two." Even realizing that zero is a number and that negative numbers are numbers took centuries of sophisticated development. Extending numbers to the rationals, the reals, the imaginaries, and the hyperreals requires an enormous cognitive apparatus and goes well beyond what babies and animals, and even a normal adult without instruction, can do. The remainder of this book will be concerned with the embodied cognitive capacities that allow one to go from innate basic numerical abilities to a deep and rich understanding of, say, college-level mathematics.

From the work we have done to date, it appears that such advanced mathematical abilities are not independent of the cognitive apparatus used outside mathematics. Rather, it appears that the cognitive structure of advanced mathematics makes use of the kind of conceptual apparatus that is the stuff of ordinary everyday thought. This chapter presents prominent examples of the kinds of everyday conceptual mechanisms that are central to mathematics—especially advanced mathematics—as it is embodied in human beings. The mechanisms we will be discussing are (a) image schemas, (b) aspectual schemas, (c) conceptual metaphor, and (d) conceptual blends.

Spatial Relations Concepts and Image Schemas

Every language has a system of spatial relations, though they differ radically from language to language. In English we have prepositions like *in, on, through, above,* and so on. Other languages have substantially different systems. However, research in cognitive linguistics has shown that spatial relations in a given language decompose into conceptual primitives called *image schemas*, and these conceptual primitives appear to be universal.

For example, the English word *on*, in the sense used in "The book is *on* the desk," is a composite of three primitive image schemas:

- The *Above schema* (the book is *above* the desk)
- The *Contact schema* (the book is *in contact with* the desk)
- The *Support schema* (the book is *supported by* the desk)

The Above schema is orientational; it specifies an orientation in space relative to the gravitational pull one feels on one's body. The Contact schema is one of a number of topological schemas; it indicates the absence of a gap. The Support schema is force-dynamic in nature; it indicates the direction and nature of a force. In general, static image schemas fall into one of these categories: orientational, topological, and force-dynamic. In other languages, the primitives combine in different ways. Not all languages have a single concept like the English *on*. Even in a language as close as German, the *on* in *on the table* is rendered as *auf*, while the *on* in *on the wall* (which does not contain the Above schema) is translated as *an*.

A common image schema of great importance in mathematics is the *Container schema*, which occurs as the central part of the meaning of words like *in* and *out*. The Container schema has three parts: an Interior, a Boundary, and an

Exterior. This structure forms a gestalt, in the sense that the parts make no sense without the whole. There is no Interior without a Boundary and an Exterior, no Exterior without a Boundary and an Interior, and no Boundary without sides, in this case an Inside and an Outside. This 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.

To get schemas for the concepts In and Out, more must be added to the Container schema. The concept In requires that the Interior of the Container schema be "profiled"—that is, highlighted or activated in some way over the Exterior and Boundary. In addition, a figure/ground distinction must be added. For example, in a sentence like "The car is in the garage," the garage is the ground; that is, it is the landmark relative to which the car (the figure) is located. In cognitive linguistics, the ground in an image schema is called the *Landmark*, and the figure is called the *Trajector*. Thus, the In schema has the structure:

- Container schema, with Interior, Boundary, and Exterior
- Profiled: the Interior
- Landmark: the Interior

Image schemas have a special cognitive function: They are both *per*ceptual and *con*ceptual in nature. As such, they provide a bridge between language and reasoning on the one hand and vision on the other. Image schemas can fit visual perception, as when we see the milk as being *in* the glass. They can also be imposed on visual scenes, as when we see the bees swarming *in* the garden, where there is no physical container that the bees are in. Because spatial-relations terms in a given language name complex image schemas, image schemas are the link between language and spatial perception.

In addition, complex image schemas like *In* have built-in spatial "logics" by virtue of their image-schematic structures. Figure 2.1 illustrates the spatial logic built into the Container schema. In connection with this figure, consider the following two statements:

1. Given two Container schemas *A* and *B* and an object *X*, if *A* is *in B* and *X* is *in A*, then *X* is *in B*.
2. Given two Container schemas *A* and *B* and an object *Y*, if *A* is *in B* and *Y* is *outside of B*, then *Y* is *outside of A*.

We don't have to perform deductive operations to draw these conclusions. They are self-evident simply from the images in Figure 2.1. Because image

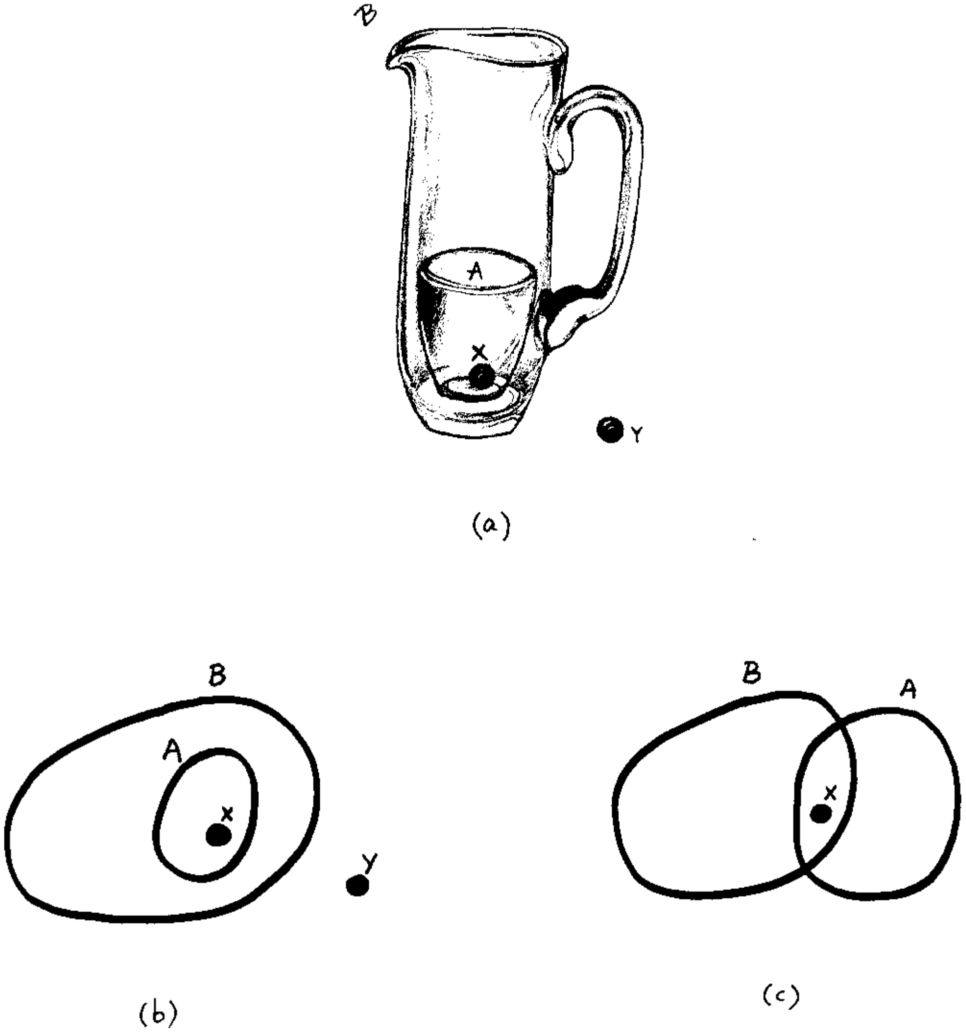

FIGURE 2.1 The logic of cognitive Container schemas. In (a), one cognitive Container schema, *A*, occurs inside another, *B*. By inspection, one can see that if *X* is in *A*, then *X* is in *B*. Similarly, if *Y* is outside *B*, then *Y* is outside *A*. We conceptualize physical containers in terms of cognitive containers, as shown in (b), which has the same logic as (a). However, conceptual containers, being part of the mind, can do what physical containers usually cannot—namely, form intersections, as in (c). In that case, an imagined entity *X* can be in two Container schemas *A* and *B* at once. Cognitive Container schemas are used not only in perception and imagination but also in conceptualization, as when we conceptualize bees as swarming *in* the garden. Container schemas are the cognitive structures that allow us to make sense of familiar Venn diagrams (see Figure 2.4).

schemas have spatial logics built into their imagistic structure, they can function as spatial concepts and be used directly in spatial reasoning. Reasoning about space seems to be done directly in spatial terms, using image schemas rather than symbols, as in mathematical proofs and deductions in symbolic logic.

Ideas do not float abstractly in the world. Ideas can be created only by, and instantiated only in, brains. Particular ideas have to be generated by neural structures in brains, and in order for that to happen, exactly the right kind of neural processes must take place in the brain's neural circuitry. Given that image schemas are conceptual in nature—that is, they constitute ideas with a structure of a very special kind—they must arise through neural circuitry of a very special kind.

Terry Regier (1996) has used the techniques of structured connectionism to build a computational neural model of a number of image schemas, as part of a neural simulation of the learning of spatial-relations terms in various languages. The research involved in Regier's simulation makes certain things clear. First, topographic maps of the visual field are needed in order to link cognition to vision. Second, a visual "filling-in" mechanism (Ramachandran & Gregory, 1991), in which activation spreads from outside to inside in a map of the visual field, will, in combination with other neural structures required, yield the topological properties of the Container schema. Third, orientation-sensitive cell assemblies found in the visual cortex are employed by orientational schemas. Fourth, map comparisons, requiring neural connections across maps, are needed. Such map-comparison structures are the locus of the relationship between the Trajector and the Landmark. Whatever changes are made in future models of spatial-relations concepts, it appears that at least these features will be needed.

Here is the importance of this for embodied mathematics: The concept of containment is central to much of mathematics. Closed sets of points are conceptualized as containers, as are bounded intervals, geometric figures, and so on. The concept of orientation is equally central. It is used in notions like angles, direction of change (tangents to a curve), rotations, and so on. The concepts of containment and orientation are not special to mathematics but are used in thought and language generally. Like any other concepts, these arise only via neural mechanisms in the right kind of neural circuitry. It is of special interest that the neural circuitry we have evolved for other purposes is an inherent part of mathematics, which suggests that embodied mathematics does not exist independently of other embodied concepts used in everyday life. Instead, mathematics makes use of our adaptive capacities—our ability to adapt other cognitive mechanisms for mathematical purposes.

Incidentally, the visual system of the brain, where such neural mechanisms as orientational cell assemblies reside, is not restricted to vision. It is also the locus of mental imagery. Mental imagery experiments, using fMRI techniques, have shown that much of the visual system, down to the primary visual cortex, is active when we create mental imagery without visual input. The brain's visual system is also active when we dream (Hobson, 1988, 1994). Moreover, congenitally blind people, most of whom have the visual system of the brain intact, can perform visual imagery experiments perfectly well, with basically the same results as sighted subjects, though a bit slower (Marmor & Zaback, 1976; Carpenter & Eisenberg, 1978; Zimler & Keenan, 1983; Kerr, 1983). In short, one should not think of the visual system as operating purely on visual input. Thus, it makes neurological sense that structures in the visual system can be used for conceptual purposes, even by the congenitally blind.

Moreover, the visual system is linked to the motor system, via the prefrontal cortex (Rizzolatti, Fadiga, Gallese, & Fogassi, 1996; Gallese, Fadiga, Fogassi, & Rizzolatti, 1996). Via this connection, motor schemas can be used to trace out image schemas with the hands and other parts of the body. For example, you can use your hands to trace out a seen or imagined container, and correspondingly you can visualize the structure of something whose shape you trace out with your hands in the dark. Thus, congenitally blind people can get "visual" image-schematic information from touch. Image schemas are kinesthetic, going beyond mere seeing alone, even though they use neural structures in the visual system. They can serve general conceptual purposes and are especially well suited for a role in mathematical thought.

There are many image schemas that characterize concepts important for mathematics: centrality, contact, closeness, balance, straightness, and many, many more. Image schemas and their logics are essential to mathematical reasoning.

Motor Control and Mathematical Ideas

One might think that nothing could be further from mathematical ideas than motor control, the neural system that governs how we move our bodies. But certain recent discoveries about the relation between motor control and the human conceptual system suggest that our neural motor-control systems may be centrally involved in mathematical thought. Those discoveries have been made in the field of structured connectionist neural modeling.

Building on work by David Bailey (1997), Srini Narayanan (1997) has observed that neural motor-control programs all have the same superstructure:

- *Readiness:* Before you can perform a bodily action, certain conditions of readiness have to be met (e.g., you may have to reorient your body, stop doing something else, rest for a moment, and so on).

- *Starting up:* You have to do whatever is involved in beginning the process (e.g., to lift a cup, you first have to reach for it and grasp it).
- *The main process:* Then you begin the main process.
- *Possible interruption and resumption:* While you engage in the main process, you have an option to stop, and if you do stop, you may or may not resume.
- *Iteration or continuing:* When you have done the main process, you can repeat or continue it.
- *Purpose:* If the action was done to achieve some purpose, you check to see if you have succeeded.
- *Completion:* You then do what is needed to complete the action.
- *Final state:* At this point, you are in the final state, where there are results and consequences of the action.

This might look superficially like a flow diagram used in classical computer science. But Narayanan's model of motor-control systems differs in many significant respects: It operates in real time, is highly resource- and context-dependent, has no central controller or clock, and can operate concurrently with other processes, accepting information from them and providing information to them. According to the model, these are all necessary properties for the smooth function of a neural motor-control system.

One might think the motor-control system would have nothing whatever to do with concepts, especially abstract concepts of the sort expressed in the grammars of languages around the world. But Narayanan has observed that this general motor-control schema has the same structure as what linguists have called *aspect*—the general structuring of events. Everything that we perceive or think of as an action or event is conceptualized as having that structure. We reason about events and actions in general using such a structure. And languages throughout the world all have means of encoding such a structure in their grammars. What Narayanan's work tells us is that *the same neural structure used in the control of complex motor schemas can also be used to reason about events and actions* (Narayanan, 1997).

We will call such a structure an *Aspect schema*.

One of the most remarkable of Narayanan's results is that exactly the same general neural control system modeled in his work can carry out a complex bodily movement when providing input to muscles, or carry out a rational inference when input to the muscles is inhibited. What this means is that neural control systems for bodily motions have the same characteristics needed for rational inference in the domain of aspect—that is, the structure of events.

Among the logical entailments of the aspectual system are two inferential patterns important for mathematics:

- The stage characterizing the completion of a process is further along relative to the process than any stage within the process itself.
- There is no point in a process further along than the completion stage of that process.

These fairly obvious inferences, as we shall see in Chapter 8 on infinity, take on considerable importance for mathematics.

Verbs in the languages of the world have inherent aspectual structure, which can be modified by various syntactic and morphological means. What is called *imperfective aspect* focuses on the internal structure of the main process. *Perfective aspect* conceptualizes the event as a whole, not looking at the internal structure of the process, and typically focusing on the completion of the action. Some verbs are inherently imperfective, like *breathe* or *live*. The iterative activity of breathing and the continuous activity of living—as we conceptualize them—do not have completions that are part of the concept. Just as the neural motor-control mechanism governing breathing does not have a completion (stopping, as in holding one's breath, is quite different from completion), so the concept is without a notion of completion. Death follows living but is *not* the *completion* of living, at least in our culture. Death is conceptualized, rather, as the cutting-off of life, as when a child is killed in an auto accident: Death follows life, but life is not completed. And you can say, "I have lived" without meaning that your life has been completed. Thus, an inherently imperfective concept is one that is conceptualized as being open-ended—as not having a completion.

There are two ways in which processes that have completions can be conceptualized: The completion may be either (1) internal to the process or (2) external to the process. This is not a matter of how the natural world really works but of how we conceptualize it and structure it through language. Take an example of case 1: If you jump, there are stages of jumping—namely, taking off, moving through the air, and landing. Landing completes the process of jumping. The completion, landing, is conceptualized as part of the jumping, as *internal* to what "jump" means. There is a minimally contrasting case that exemplifies case 2: flying. In the everyday concept of flying, as with birds and planes, landing is part of the conceptual frame. Landing follows flying and is a completion of flying. But *landing* is not conceptualized as part of *flying*. *Landing* is a completion of *flying* but it is *external* to *flying*. The distinction between an internal completion, as in *jump,* and an external completion, as in *fly*, is crucial in aspect.

Aspectual ideas occur throughout mathematics. A rotation through a certain number of degrees, for example, is conceptualized as a process with a starting point and an ending point. The original notion of continuity for a function was conceptualized in terms of a continuous process of motion—one without intermediate ending points. The very idea of an algorithmic process of calculation involves a starting point, a process that may or may not be iterative, and a well-defined completion. As we shall see in Chapters 8 through 11, all notions of infinity and infinitesimals use aspectual concepts.

The Source-Path-Goal Schema

Every language includes ways of expressing spatial sources (e.g., "from") and goals (e.g., "to," "toward") and paths intermediate between them (e.g., "along," "through," "across"). These notions do not occur isolated from one another but, rather, are part of a larger whole, the *Source-Path-Goal schema*. It is the principal image schema concerned with motion, and it has the following elements (or *roles*):

- 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: the speed of motion, the trail left by the thing moving, 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 or shrunk or deformed and still remain a path. As in the case of the Container schema, we can form spatial relations from this schema by the addition of profiling 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 (see Figure 2.2):

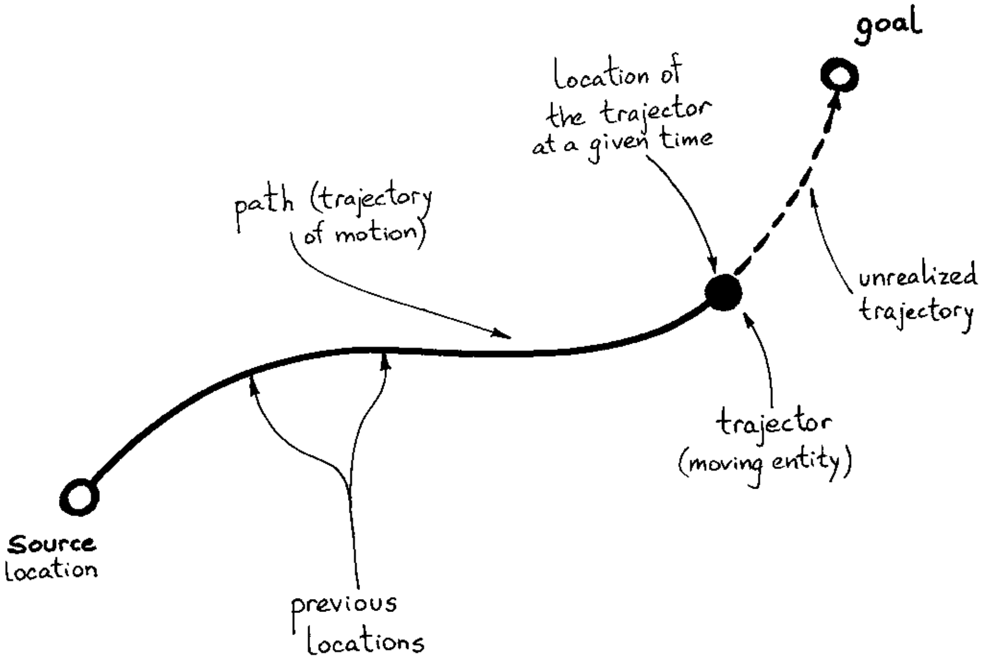

FIGURE 2.2 The Source-Path-Goal schema. We conceptualize linear motion using a conceptual schema in which there is a moving entity (called a trajector), a source of motion, a trajectory of motion (called a path), and a goal with an unrealized trajectory approaching that goal. There is a logic inherent in the structure of the schema. For example, if you are at a given location on a path, you have been at all previous locations on that path.

- 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.

The Source-Path-Goal schema is ubiquitous in mathematical thought. The very notion of a directed graph (see Chapter 7), for example, is an instance of the Source-Path-Goal schema. Functions in the Cartesian plane are often conceptualized in terms of motion along a path—as when a function is described as "going up," "reaching" a maximum, and "going down" again.

One of the most important manifestations of the Source-Path-Goal schema in natural language is what Len Talmy (1996, 2000) has called *fictive motion*. In one form of fictive motion, a line is thought of in terms of motion tracing that line, as in sentences like "The road *runs* through the woods" or "The fence *goes* up the hill." In mathematics, this occurs when we think of two lines "*meeting* at a point" or the graph of a function as "*reaching* a minimum at zero."

Conceptual Composition

Since image schemas are conceptual in nature, they can form complex composites. For example, the word "into" has a meaning—the *Into schema*—that is the composite of an *In schema* and a *To schema.* The meaning of "out of" is the composite of an *Out schema* and a *From schema.* These are illustrated in Figure 2.3. Formally, they can be represented in terms of correspondences between elements of the schemas that are part of the composite.

The following notations indicate composite structures.

The Into schema

- The In schema: A Container schema, with the Interior profiled and taken as Landmark
- The To schema: A Source-Path-Goal schema, with the Goal profiled and taken as Landmark
- Correspondences: (Interior; Goal) and (Exterior; Source)

The Out-of schema

- The Out schema: A Container schema, with the Exterior profiled and taken as Landmark
- The From schema: A Source-Path-Goal schema, with the Source profiled and taken as Landmark
- Correspondences: (Interior; Source) and (Exterior; Goal)

Conceptual Metaphor

Metaphor, long thought to be just a figure of speech, has recently been shown to be a central process in everyday thought. Metaphor is not a mere embellishment; it is the basic means by which abstract thought is made possible. One of the principal results in cognitive science is that abstract concepts are typically understood, via metaphor, in terms of more concrete concepts. This phenome-

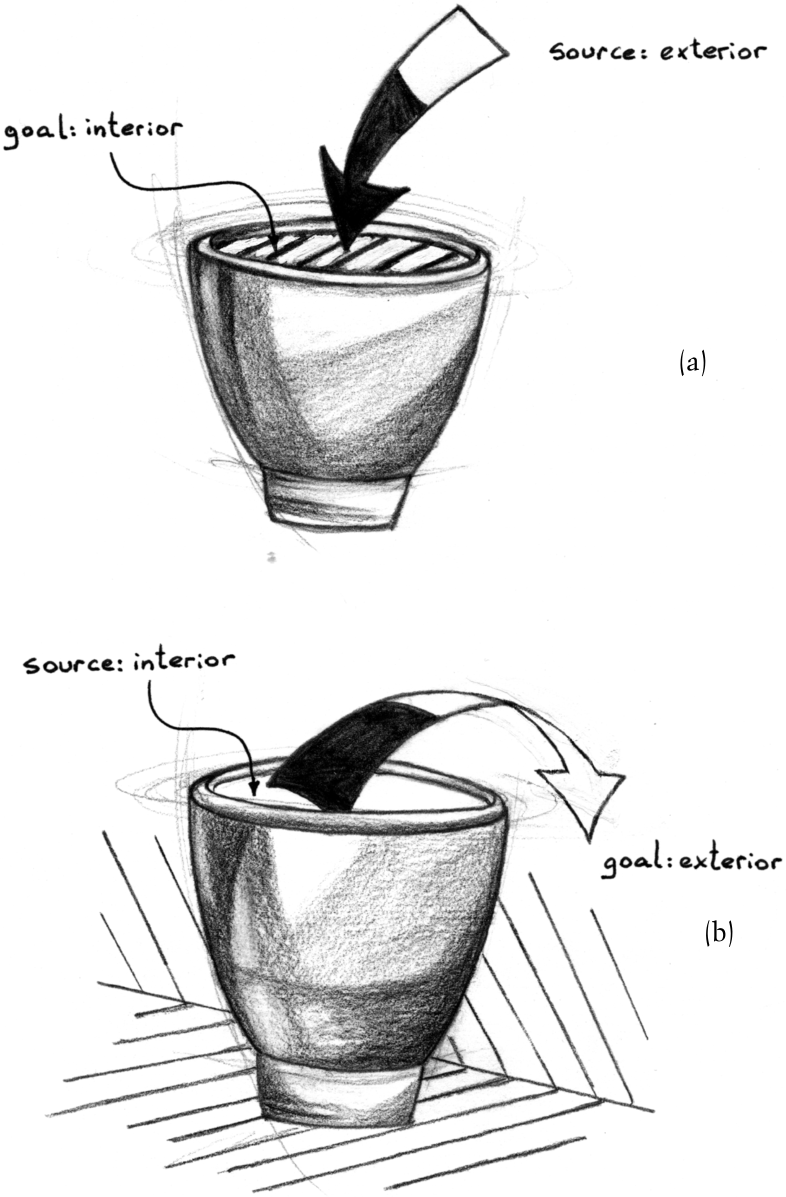

FIGURE 2.3 Conceptual composition of schemas. The English expressions "into" and "out of" have composite meanings. "In" profiles the interior of a Container schema, while "out" profiles the exterior. "To" profiles the goal of the Source-Path-Goal schema, while "from" profiles the source. With "into" (a), the interior is the goal and the exterior is the source. With "out of" (b), the interior is the source and the exterior is the goal.

non has been studied scientifically for more than two decades and is in general as well established as any result in cognitive science (though particular details of analyses are open to further investigation). One of the major results is that metaphorical mappings are systematic and not arbitrary.

Affection, for example, is understood in terms of physical warmth, as in sentences like "She *warmed* up to me," "You've been *cold* to me all day," "He gave me an *icy* stare," "They haven't yet *broken the ice*." As can be seen by this example, the metaphor is not a matter of words, but of conceptual structure. The words are all different (*warm, cold, icy, ice*), but the conceptual relationship is the same in all cases: Affection is conceptualized in terms of warmth and disaffection in terms of cold.

This is hardly an isolated example:

- Importance is conceptualized in terms of size, as in "This is a big issue," "He's a giant in the meatpacking business," and "It's a small matter; we can ignore it."
- Similarity is conceptualized in terms of physical closeness, as in "These colors are very close," "Our opinions on politics are light-years apart," "We may not agree, but our views are in the same ballpark," and "Over the years, our tastes have diverged."
- Difficulties are conceptualized as burdens, as in "I'm weighed down by responsibilities," "I've got a light load this semester," and "He's overburdened."
- Organizational structure is conceptualized as physical structure, as in "The theory is full of holes," "The fabric of this society is unraveling," "His proposed plan is really tight; everything fits together very well."

Hundreds of such conceptual metaphors have been studied in detail. They are extremely common in everyday thought and language (see Lakoff & Johnson, 1980, 1999; Grady, 1998; Núñez, 1999). On the whole, they are used unconsciously, effortlessly, and automatically in everyday discourse; that is, they are part of the cognitive unconscious. Many arise naturally from correlations in our commonplace experience, especially our experience as children. Affection correlates with warmth in the experience of most children. The things that are important in their lives tend to be big, like their parents, their homes, and so on. Things that are similar tend to occur close together: trees, flowers, dishes, clouds. Carrying something heavy makes it difficult to move and to perform other activities. When we examine a complex physical object with an internal structure, we can perceive an organization in it. Not all conceptual metaphors arise in this way, but most of the basic ones do.

Such correlations in experience are special cases of the phenomenon of *conflation* (see C. Johnson, 1997). Conflation is part of embodied cognition. It is the simultaneous activation of two distinct areas of our brains, each concerned with distinct aspects of our experience, like the physical experience of warmth and the emotional experience of affection. In a conflation, the two kinds of experience occur inseparably. The coactivation of two or more parts of the brain generates a single complex experience—an experience of affection-with-warmth, say, or an experience of difficulty-with-a-physical-burden. It is via such conflations that neural links across domains are developed—links that often result in conceptual metaphor, in which one domain is conceptualized in terms of the other.

Each such conceptual metaphor has the same structure. Each is a unidirectional mapping from entities in one conceptual domain to corresponding entities in another conceptual domain. As such, conceptual metaphors are part of our system of thought. Their primary function is to allow us to reason about relatively abstract domains using the inferential structure of relatively concrete domains. The structure of image schemas is preserved by conceptual metaphorical mappings. In metaphor, conceptual cross-domain mapping is primary; metaphorical language is secondary, deriving from the conceptual mapping. Many words for source-domain concepts also apply to corresponding target-domain concepts. When words for source-domain concepts do apply to corresponding target concepts, they do so systematically, not haphazardly.

To see how the inferential structure of a concrete source domain gives structure to an abstract target domain, consider the common conceptual metaphor that States Are Locations, as in such expressions as "I'm in a depression," "He's close to hysteria; don't push him over the edge," and "I finally came out of my funk." The source domain concerns bounded regions in physical space. The target domain is about the subjective experience of being in a state.

STATES ARE LOCATIONS

Source Domain SPACE		*Target Domain* STATES
Bounded Regions in Space	→	States

Here is an example of how the patterns of inference of the source domain are carried over to the target domain.

If you're in a *bounded region*, you're not out of that *bounded region*.	→	If you're in a *state*, you're not out of that *state*.
If you're out of a *bounded region*, you're not in that *bounded region*.	→	If you're out of a *state,* you're not in that *state*.
If you're deep in a *bounded region*, you are far from being out of that *bounded region*.	→	If you're deep in a *state*, you are far from being out of that *state*.
If you are on the edge of a *bounded region*, you are close to being in that *bounded region.*	→	If you are on the edge of a *state*, you are close to being in that *state*.

Throughout this book we will use the common convention that *names* of metaphorical mappings are given in the form "A Is B," as in "States Are Bounded Regions in Space." It is important to distinguish between such names for metaphorical mappings and the metaphorical mappings themselves, which are given in the form "B → A," as in "Bounded Regions in Space → States." Here the source domain is to the left of the arrow and the target domain is to the right.

An enormous amount of our everyday abstract reasoning arises through such metaphorical cross-domain mappings. Indeed, much of what is often called logical inference is in fact spatial inference mapped onto an abstract logical domain. Consider the logic of the Container schema. There is a commonplace metaphor, Categories Are Containers, through which we understand a category as being a bounded region in space and members of the category as being objects inside that bounded region. The metaphorical mapping is stated as follows:

Categories Are Containers

Source Domain		*Target Domain*
Containers		Categories
Bounded regions in space	→	Categories
Objects inside the bounded regions	→	Category members
One bounded region inside another	→	A subcategory of a larger category

Suppose we apply this mapping to the two inference patterns mentioned above that characterize the spatial logic of the Container schema, as follows:

Source Domain CONTAINER SCHEMA INFERENCES		*Target Domain* CATEGORY INFERENCES
Excluded Middle		*Excluded Middle*
Every object *X* is either in *Container schema A* or out of *Container schema A*.	→	Every entity *X* is either in *category A* or out of *category A*.
Modus Ponens		*Modus Ponens*
Given two *Container schemas A* and *B* and an object *X*, if *A* is in *B* and *X* is in *A*, then *X* is in *B*.	→	Given two *categories A* and *B* and an entity *X*, if *A* is in *B* and *X* is in *A*, then *X* is in *B*.
Hypothetical Syllogism		*Hypothetical Syllogism*
Given three *Container schemas A*, *B*, and *C*, if *A* is in *B* and *B* is in *C*, then *A* is in *C*.	→	Given three *categories A*, *B* and *C*, if *A* is in *B* and *B* is in *C*, then *A* is in *C*.
Modus Tollens		*Modus Tollens*
Given two *Container schemas A* and *B* and an object *Y*, if *A* is in *B* and *Y* is outside *B*, then *Y* is outside *A*.	→	Given two *categories A* and *B* and an entity *Y*, if *A* is in *B* and *Y* is outside *B*, then *Y* is outside *A*.

The point here is that the logic of Container schemas is an embodied spatial logic that arises from the neural characterization of Container schemas. The excluded middle, modus ponens, hypothetical syllogism, and modus tollens of classical categories are metaphorical applications of that spatial logic, since the Categories Are Containers metaphor, like conceptual metaphors in general, preserves the inferential structure of the source domain.

Moreover, there are important entailments of the Categories Are Containers metaphor:

The overlap of the interiors of two bounded regions	→	The conjunction of two categories
The totality of the interiors of two bounded regions	→	The disjunction of two categories

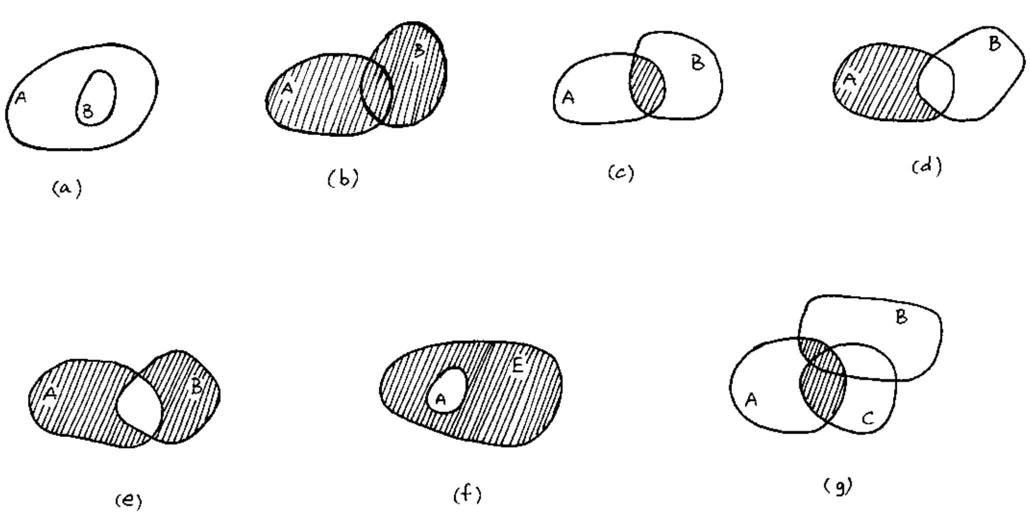

FIGURE 2.4 Venn diagrams. Here is a common set of Venn diagrams of the sort one finds in texts on classes and sets, which are typically conceptualized metaphorically as containers and derive their logics from the logic of conceptual Container schemas. When one "visualizes" classes and sets in this way, one is using cognitive Container schemas in the visualization. The diagrams depict various mathematical ideas: (a) the relation $B \subseteq A$; (b) $A \cup B$; (c) $A \cap B$; (d) the difference $A \backslash B$; (e) the symmetric difference $A \Delta B$; (f) the complement $C_E A$; and (g) $A \cap (B \cup C)$, which equals $(A \cap B) \cup (A \cap C)$.

In short, given the spatial logic of Container schemas, the Categories Are Containers metaphor yields an everyday version of what we might call folk Boolean logic, with intersections and unions. That is why the Venn diagrams of Boolean logic look so natural to us (see Figure 2.4), although there are differences between folk Boolean logic and technical Boolean logic, which will be discussed in Chapter 6. Folk Boolean logic, which is *con*ceptual, arises from a *per*ceptual mechanism—the capacity for perceiving the world in terms of contained structures.

From the perspective of the embodied mind, spatial logic is primary and the abstract logic of categories is secondarily derived from it via conceptual metaphor. This, of course, is the very opposite of what formal mathematical logic suggests. It should not be surprising, therefore, that embodied mathematics will look very different from disembodied formal mathematics.

Metaphors That Introduce Elements

Conceptual metaphors do not just map preexisting elements of the source domain onto preexisting elements of the target domain. They can also *introduce* *new elements* into the target domain. Consider, for example, the concept of love. There is a common metaphor in the contemporary Western world in which Love Is a Partnership. Here is the mapping.

LOVE IS A PARTNERSHIP

Source Domain BUSINESS		*Target Domain* LOVE
Partners	→	Lovers
Partnership	→	Love relationship
Wealth	→	Well-being
Profits from the business	→	"Profits" from the love relationship
Work for the business	→	"Work" put into the relationship
Sharing of work for the business	→	Sharing of "work" put into the relationship
Sharing of profits from the business	→	Sharing of "profits" from the relationship

Love need not always be conceptualized via this metaphor as a partnership. Romeo and Juliet's love was not a partnership, nor was Tristan and Isolde's. Similarly, love in many cultures around the world is not conceptualized in terms of business—and it need not be so conceptualized for individual cases in the Western world. But this is a common metaphorical way of understanding love—so common that it is sometimes taken as literal. For example, sentences like "I'm putting all the work into this relationship and you're getting everything out of it," "It was hard work, but worth it," and "The relationship was so unrewarding that it wasn't worth the effort" are so commonplace in discussions of love relationships that they are rarely noticed as metaphorical at all.

From the perspective of this book, there is an extremely important feature of this metaphor: It *introduces* elements into the target domain that are not inherent to the target domain. It is not inherent in love-in-itself that there be "work in the relationship," "profits (increases in well-being) from the relationship," and a "sharing of relationship work and profits." Romeo and Juliet would have been aghast at such ideas. These ideas are elements introduced into the target domain by the Love Is a Partnership metaphor, and they don't exist there without it.

The fact that metaphors can introduce elements into a target domain is extremely important for mathematics, as we shall see later.

Evidence

Over the past two decades, an enormous range of empirical evidence has been collected that supports this view of conceptual metaphor. The evidence comes from various sources:

- generalizations over polysemy (cases where the same word has multiple systematically related meanings)
- generalizations over inference patterns (cases where source and target domains have corresponding inference patterns)
- novel cases (new examples of conventional mappings, as in poetry, song, advertisements, and so on) (see Lakoff & Turner, 1989)
- psychological experiments (see Gibbs, 1994)
- historical semantic change (see Sweetser, 1990)
- spontaneous gesture (see McNeill, 1992)
- American Sign Language (see Taub, 1997)
- child language development (see C. Johnson, 1997)
- discourse coherence (see Narayanan, 1997)
- cross-linguistic studies.

For a thorough discussion of such evidence, see Lakoff and Johnson, 1999.

Sophisticated Mathematical Ideas

Sophisticated mathematics, as we have pointed out, is a lot more than just basic arithmetic. Mathematics extends the use of numbers to many other ideas, for example, the numerical study of angles (trigonometry), the numerical study of change (calculus), the numerical study of geometrical forms (analytic geometry), and so on. We will argue, in our discussion of all these topics and more, that conceptual metaphor is the central cognitive mechanism of extension from basic arithmetic to such sophisticated applications of number. Moreover, we will argue that a sophisticated understanding of arithmetic itself requires conceptual metaphors using nonnumerical mathematical source domains (e.g., geometry and set theory). We will argue further that conceptual metaphor is also the principal cognitive mechanism in the attempt to provide set-theoretical foundations for mathematics and in the understanding of set theory itself.

Finally, it should become clear in the course of this discussion that much of the "abstraction" of higher mathematics is a consequence of the systematic layering of metaphor upon metaphor, often over the course of centuries.

Each metaphorical layer, as we shall see, carries inferential structure systematically from source domains to target domains—systematic structure that gets lost in the layers unless they are revealed by detailed metaphorical analysis. A good part of this book is concerned with such metaphorical decomposition of sophisticated mathematical concepts. Because this kind of study has never been done before, we will not be able to offer the extensive forms of evidence that have been found in decades of studies of conceptual metaphor in everyday language and thought. For this reason, we will limit our study to cases that are relatively straightforward—cases where the distinctness of the source and target domains is clear, where the correspondences across the domains have been well established, and where the inferential structures are obvious.

Conceptual Blends

A *conceptual blend* is the conceptual combination of two distinct cognitive structures with fixed correspondences between them. In mathematics, a simple case is the unit circle, in which a circle is superimposed on the Cartesian plane with the following fixed correspondences: (a) The center of the circle is the origin (0,0), and (b) the radius of the circle is 1. This blend has entailments that follow from these correspondences, together with the inferential structure of *both domains*. For example, the unit circle crosses the x-axis at (1,0) and (–1,0), and it crosses the y-axis at (0,1) and (0,–1). The result is more than just a circle. It is a circle that has a fixed position in the plane and whose circumference is a length commensurate with the numbers on the x- and y-axes. A circle in the Euclidean plane, where there are no axes and no numbers, would not have these properties.

When the fixed correspondences in a conceptual blend are given by a metaphor, we call it a *metaphorical blend*. An example we will discuss extensively below is the Number-Line Blend, which uses the correspondences established by the metaphor Numbers Are Points on a Line. In the blend, new entities are created—namely, *number-points*, entities that are at once numbers and points on a line (see Fauconnier 1997; Turner & Fauconnier, 1995; Fauconnier & Turner, 1998). Blends, metaphorical and nonmetaphorical, occur throughout mathematics.

Many of the most important ideas in mathematics are metaphorical conceptual blends. As will become clear in the case-study chapters, understanding mathematics requires the mastering of extensive networks of metaphorical blends.

Symbolization

As we have noted, there is a critical distinction to be made among mathematical concepts, the written mathematical symbols for those concepts, and the words for the concepts. The words (e.g., "eighty-five" or "quatre-vingt-cinq") are part of some natural language, not mathematics proper.

In embodied mathematics, mathematical symbols, like 27, π, or $e^{\pi i}$, are meaningful by virtue of the mathematical concepts that they attach to. Those mathematical concepts are given in cognitive terms (e.g., image schemas; imagined geometrical shapes; metaphorical structures, like the number line; and so on), and those cognitive structures will ultimately require a neural account of how the brain creates them on the basis of neural structure and bodily and social experience. To understand a mathematical symbol is to associate it with a concept—something meaningful in human cognition that is ultimately grounded in experience and created via neural mechanisms.

As Stanislas Dehaene observed in the case of Mr. M—and as many of us experienced in grade school—numerical calculation may be performed with or without genuine understanding. Mr. M could remember his multiplication tables, but they were essentially meaningless to him.

The meaning of mathematical symbols is not in the symbols alone and how they can be manipulated by rule. Nor is the meaning of symbols in the interpretation of the symbols in terms of set-theoretical models that are themselves uninterpreted. Ultimately, mathematical meaning is like everyday meaning. It is part of embodied cognition.

This has important consequences for the teaching of mathematics. Rote learning and drill is not enough. It leaves out understanding. Similarly, deriving theorems from formal axioms via purely formal rules of proof is not enough. It, too, can leave out understanding. The point is not to be able to prove *that* $e^{\pi i} = -1$ but, rather, to be able to prove it knowing what $e^{\pi i}$ means, and knowing *why* $e^{\pi i} = -1$ on the basis of what $e^{\pi i}$ means, not just on the basis of the formal proof. In short, what is required is an adequate mathematical idea analysis to show *why* $e^{\pi i} = -1$ given our understanding of the ideas involved.

Euler's equation, $e^{\pi i} + 1 = 0$, ties together many of the most central ideas in classical mathematics. Yet on the surface it involves only numbers: e, π, i, 1, and 0. To show how this equation ties together *ideas,* we must have a theory of mathematical ideas and a theory of how they are mathematicized in terms of numbers.

Our interest, of course, goes beyond just $e^{\pi i}$ as such. Indeed, we are also interested in the all-too-common conception that mathematics is about calculation and about formal proofs from formal axioms and definitions and not about ideas and understanding. From the perspective of embodied mathematics, ideas and understanding are what mathematics is centrally about.