Preface and Acknowledgments

This book is based in part on the Jean Nicod Lectures that I delivered in Paris in May–June 2004. The temporal gap between the lectures and the publication is not entirely due to my slow typing, but arose from the need to assimilate the rather wide range of publications that are relevant to the thesis I am presenting. The thesis, it turns out, is one that I have been gestating over many years, and hints of it occur in fragmentary form in a number of my publications. Many of these are reports of experimental work carried out with graduate students over the years, whose contribution is much appreciated.

The thesis rests on a growing appreciation of an idea I first learned from David Marr, who refers to it as the *principle of natural constraints*. The mind has been tuned over its evolutionary history so that it carries out certain functions in a modular fashion, without regard for what an organism knows or believes or desires, but because it is in its nature, or as I more often put it, because of its architecture. So far this is an innocent enough idea that fits many different schools of psychology (and in fact is a familiar part of J. J. Gibson's direct realism theory, though used there to very different ends). The particular constraint I am interested in here takes the form of a mechanism that allows the modular perceptual system to do things that many philosophers have said (correctly) can only be done by using the sophisticated machinery of concepts and the logical machinery of induction, deduction, and what Charles Sanders Peirce called abduction. The mechanism includes the capacity to select individual things in one's field of view, to reidentify each of them under certain conditions as the same individual thing that was seen before, and to keep track of their enduring individuality despite radical changes in their properties. I claim that so long as we are in the kind of world we live in there are mechanisms that allow the visual system to do these things without using the heavy equipment of concepts, identity, and tenses (which are needed for other tasks).

For example, this is a world in which most surfaces that we see are surfaces of physical objects, so that most of the texture elements we see move coherently as the object moves; almost all elements nearby on the proximal image are at the same distance from the viewer; and, when objects disappear, they frequently reappear nearby, and often with a particular pattern of occlusion and disocclusion at the edges of the occluding opaque surfaces, and so on. Identifying things as ones we have seen before and keeping track of them as being the same individual objects over time is at the heart of my research, which has shown that we are very good at doing this in a way that does not use encoded properties (or the conceptual category) of the things that are tracked and reidentified. This mechanism is important to us because if it were not for the existence of such nonconceptual processes, our concepts would not be grounded in experience and thus would not have the meaning that they do.

I have proposed that the capacity to individuate and track several independently moving things is accomplished by a mechanism in the early vision module that I have called *FINSTs* (I call them ''Fingers of INSTantiation'' because they were initially viewed as a mechanism for instantiating or binding the arguments of visual predicates to objects in the world). This primitive nonconceptual mechanism functions to identify, reidentify, and track distal objects. It is an ability that we exercise every waking minute, and it has also been understood to be fundamental to the way we see and understand the world.

I came upon these ideas in quite a different context, initially when I (along with my colleague Edward Elcock) attempted to develop a computer system for reasoning about diagrams, and later when I was carrying out experimental research on vision, visual attention, and mental imagery. This may seem like a circuitous route, but it has turned out that all these endeavors involve the same puzzles, which I later discovered were also the puzzles that preoccupy many philosophers: how concepts are grounded in experience; how we manage to encode and represent properties of the world when there are so many of them; why we feel that we are conscious of seeing an enormous number of things but are unable to report most of them, while at the same time a great deal of information of which we are not conscious can affect our behavior. These puzzles appear in their most striking form in discussions of two related problems: What are the properties of mental images that allows them to function in thought, and how do certain kinds of thoughts—thoughts about spatial layouts—manage to display properties very similar to those of perceived space?

In this book I examine a number of critical functions of early vision (the part of vision that is informationally encapsulated from the rest of the mind) in the light of the FINST mechanism. Chapter 1 looks at the nature of the problems that FINST are intended to solve, from our initial encounter with the problem of incrementally building a representation over time as various aspects are noticed, to the deep problem of grounding concepts in sensory information. This chapter also introduces FINST theory in terms of a number of experiments that illustrate their nature as pointers to things in the perceived world. It also offers the suggestion that FINSTs serve to provide what philosophers have called *demonstrative reference* or *demonstrative identification*. Chapter 2 focuses on a particular function that FINSTs serve, namely, the nonconceptual *tracking* of individual things that move and change their properties. Since tracking is one of the critical aspects of our commerce with the world, these experiments serve as concrete examples of the role that FINSTs play in this process. It also provides a basis for a number of additional properties of this mind–world connection: it shows that things can be tracked as unidentified things with an enduring numerical identity (where by ''unidentified'' I mean they are not represented in terms of any conceptual category or in terms of distinctive properties). The notion of tracking also links this work to some findings in cognitive development where it has been shown that very young infants (under six months of age) are able to track things that briefly disappear and are able to anticipate how many hidden things there are (as long as there are three or fewer).

Chapter 3 looks more closely at another of the functions that FINSTs perform, that of selecting things through something like attention (FINSTs are not exactly the same as focal attention and I discuss the differences). This chapter raises some of the controversial aspects of attentional selection. It argues that selecting is nonconceptual and does not depend on the prior encoding of any properties of the things selected—including their location. I argue that the reason this seems unintuitive is that we fail to distinguish between the various roles that the properties of things play in this process. Properties are involved in picking out things to which FINSTs are assigned, and they are involved in determining whether things can be tracked, but they need not be encoded and used in the process of maintaining the identity of the things that are tracked. I spend time in chapter 3 distinguishing between causes and codes and suggest that we should be conservative in describing certain mind–world connections as representations. This brings us to an important function that FINST selection plays, solving what has

been called the *binding problem*: the problem of encoding certain sets of properties as being conjoined, as being properties of the same thing, as opposed to being properties that merely occur simultaneously in the scene. Whereas much of the psychological and philosophical literature sees the binding problem as being solved in terms of the collocation of properties, my proposal is that properties are considered conjoined if they are properties of the same FINSTed thing. This brings us to a point where we can say roughly what FINSTs attach to—what it is that I have been calling ''things'' (or even FINGs, to indicate that they are interdefined with FINSTs, as those things that FINSTs select and refer to)—they attach to what, in our sort of world, typically turn out to be individual visible physical objects. I discuss the frequently cited notion of nonconceptual representation and suggest how this idea is closely related to the story I am telling about FINSTs. In chapter 3 I propose that the only nonconceptual representational content we have is that secured by FINST indexes.

The view that only properties of FINSTed things get represented puts me in conflict with those who appeal to the richness of conscious experience in defending nonconceptual representation. For this reason I devote chapter 4 to a discussion of the role of conscious contents in the process of connecting mind and world. What I claim is that the contents of conscious experience are only one source of evidence for mental contents, and not even a very reliable one. I claim that there is no level of representation that corresponds specifically to the content of conscious experience and, therefore, that equating nonconceptual representation with the content of conscious experience is a mistake. The discussion of the content of conscious experience brings the topic to the nature of the mental images that we experience when reasoning about concrete sensory appearances, about spatial layouts, or when we use spatial mental model models in reasoning. In chapter 4 I focus primarily on what we can make of the contents of conscious experience, and I use theories of mental imagery as the example of how we are misled when we view conscious content as a type of representation.

It is not until chapter 5 that I focus directly on the problem of the representation of space in active working memory (as opposed to long-term memory). Here I review some of the proposals on how some mental representations manage to exhibit certain sorts of spatial properties. Most of these proposals hypothesize some internal constraints on representations of spatial layouts. The most widely accepted and intuitive proposal is that the spatial character of representations of space arises because the representations themselves are instantiated by spatial layouts in the brain—what I

call *neural layouts*. I review this proposal as well as some related ones that attempt to retain the benefits of neural layouts without assuming a spatial arrangement in the brain. Of these, the one referred to as *functional space* fails to have any explanatory value, and the other, which I call the *spatial architecture proposal*, reduces essentially to the literal space alternative. In this discussion I distinguish representations involved in long-term memory from representations I refer to as *active spatial representations* (ASPARs), which are constructed both by vision and by reasoning that relies on imagined geometrical of spatial layouts. I then list what I take to be some of the conditions that an ASPAR must meet, which include its capacity to represent magnitudes, its spatial configurational stability, its amodal nature, and its intimate connection with the motor system. The latter brings us to an overview of Poincaré's proposal.

I finish chapter 5 by presenting what might be seen as a fairly radical and speculative proposal for an externalist theory of spatial representation in ASPARs (i.e., in spatial reasoning). The hypothesis, which arises from the ideas about FINSTs that I discuss in the first part of the book, is what I call the *index projection hypothesis*. This proposal claims that in constructing a spatial representation from approximate, partial, and qualitative information stored in long-term memory, we pick out things in the concurrently perceived world using FINSTs and associate imagined objects with them (we think of the imagined objects as being located at the sensed objects). This allows us to use the perceptual system to draw inferences by pattern recognition rather than logical reasoning. I then generalize the projection hypothesis to nonvisual modalities such as proprioception, which requires that I deal with the multiple frame of reference problem (as opposed to a unitary allocentric frame of reference). In this task, coordinate transformation mechanisms, which are plentiful in the brain (especially in parietal cortex, as well as in superior colliculus and premotor cortical areas), play a central role.

Throughout this essay I try to draw morals for a number of philosophical issues such as whether there are nonconceptual representations, how concepts are grounded in perception, and how the mind deals with spatial properties. Clearly this palette is more than can be dealt with adequately in one book. Yet there are some clear themes that run through these puzzles, especially when they are considered against the background of experimental findings in psychophysics, cognitive science, cognitive development, and neuroscience. So this is my attempt to find a way through that forest by focusing on one or two sunny spots where I think progress has been made in the past two decades.

In this pursuit I must express my gratitude to the Institut Jean-Nicod, who generously invited me to give these lectures, the Centre Nationale de la Recherche Scientifique (CNRS), L'École des Hautes Études en Sciences Sociales (EHESS), and L'École Normale Supérieure (ENS), who funded the Jean Nicod Prize and provided space for me the following year as I worked on this book. In particular I wish to thank François Récanati, who chairs the Nicod Prize committee and looked after the logistics of my visit; Pierre Jacob, who directs the Institut Jean-Nicod; and the many people from the Institute who made my stay memorable, especially Roberto Casati, Jérôme Dokic, Élizabeth Pacherie, Joelle Proust, and Dan Sperber. Michel Imbert and Sylvia Duchacek-Imbert were most welcoming and helped make our stay pleasant and memorable.

I also acknowledge the help I received, intended or not, from discussions with my friend Jerry Fodor and with friends and colleagues Charles R. Gallistel (who helped to educate me on the subject of navigation), Georges Rey (who carefully read an earlier version of the manuscript and tried his best to keep me from making some embarrassing philosophical slips and misusing some philosophical terms of art), Joseph Levine and Jonathan Cohen (who provided helpful comments on an earlier draft), Alan Leslie and Lila Gleitman (with whom I have taught some of this material in joint courses), Susan Carey, Ned Block, and Luca Bonatti. I am also grateful to the participants of the conference on Spatial Frames of Reference held in Paris in November 2005, including organizers Jérôme Dokic and Élizabeth Pacherie, and participants with whom I had a chance to try out the ideas in chapter 5: Jean-René Duhamel, Yves Rossetti, Charles Spence, Barbara Tversky, Yann Coello, Paolo Bartolomeo, and Sylvie Chokron.

Finally, I wish to thank my editor, Judy Feldmann, who displayed an uncanny knack for finding words for what I meant instead of what I wrote.