Preface

This book is about how we see and how we visualize. But it is equally about how we are easily misled by our everyday experience of these faculties. Galileo is said to have proclaimed (Galilei, 1610/1983; quoted in Slezak, 2002), "...if men had been born blind, philosophy would be more perfect, because it would lack many false assumptions that have been taken from the sense of sight." Many deep puzzles arise when we try to understand the nature of visual perception, visual imagery, or visual thinking. As we try to formulate scientific questions about these human capacities we immediately find ourselves being entranced by the *view from within*. This view, which the linguist Kenneth Pike (Pike, 1967) has referred to as the *emic* perspective (as opposed to the external or *etic* perspective), is both essential and perilous. As scientists we cannot ignore the contents of our conscious experience, because this is one of the principal ways of knowing what we see and what our thoughts are about. On the other hand, the contents of our conscious experience are also insidious, because they lead us to believe that we can see directly into our own minds and observe the causes of our cognitive processes.

Such traps are nothing new; psychology is used to being torn by the duality of mental life—its subjective and its objective (causal) side. Since people first began to think about the nature of mental processes, such as seeing, imagining, and reasoning, they have had to contend with the fact that knowing how these achievements appear to us on the inside often does us little good, and indeed often leads us in entirely the wrong direction, when we seek a scientific explanation. Of course we have the option of putting aside the quest for a scientific explanation and setting our goal toward finding a satisfying description in terms that are consonant with how seeing and imagining appear to us. This might be called a

phenomenological approach to understanding the workings of the mind, or the everyday folk understanding of vision.

There is nothing wrong with such a pursuit. Much popular psychology revels in it, as do a number of schools of philosophical inquiry (e.g., ordinary language philosophy, phenomenological philosophy). Yet in the long term, few of us would be satisfied with an analysis or a natural history of phenomenological regularities. One reason is that characterizing the systematic properties of how things seem to us does not allow us to connect with the natural sciences, to approach the goal of unifying psychology with biology, chemistry, and physics. It does not help us to answer the *how* and *why* questions: How does vision work? Why do things look the way they do? What happens when we think visually?

The problem with trying to understand vision and visual imagery is that on the one hand these phenomena are intimately familiar to us from the inside so it is difficult to objectify them, even though the processes involved are also too fast and too ephemeral to be observed introspectively. On the other hand, what we do observe is misleading because it is always the *world* as it appears to us that we see, not the real work that is being done by the mind in going from the proximal stimuli, generally optical patterns on the retina, to the familiar experience of seeing (or imagining) the world. The question: How do we see? appears very nearly nonsensical. Why, we see by just looking, and the reason that things look as they do to us is that this is the way that they actually are. It is only by objectifying the phenomena, by "making them strange" that we can turn the question into a puzzle that can be studied scientifically. One good way to turn the mysteries of vision and imagery into a puzzle is to ask what it would take for a computer to see or imagine. But this is not the only way, and indeed this way is often itself laden with preconceptions, as I will try to show throughout this book.

The title of this book is meant to be ambiguous. It means both that seeing and visualizing are different from thinking (and from each other), and that our intuitive views about seeing and visualizing rest largely on a grand illusion. The message of this book is that seeing is different from thinking and that to see is not, as it often seems to us, to create an inner replica of the world we are observing or thinking about or visualizing. But this is a long and not always intuitively compelling story. In fact, its counterintuitive nature is one reason it may be worth telling. When

things seem clearly a certain way it is often because we are subject to a general shared illusion. To stand outside this illusion requires a certain act of will and an open-minded and determined look at the evidence. But some things about vision and mental imagery are clear enough now that only deeply ingrained prejudices keep them from being the received view. These facts, which seem to me (if not to others) to be totally persuasive, are the topic of this book. If any of the claims appear radical it is not because they represent a leap into the dark caverns of speculation, but only that some ways of looking at the world are just too comfortable and too hard to dismiss. Consequently, what might be a straightforward story about how we see becomes a long journey into the data and theory, developed mostly over the past thirty years, as well into the conceptual issues that surround them.

The journey begins with the question of why things we see appear to us as they do—a question that has been asked countless times in each generation. In chapter 1, I describe experiments and demonstrations, as well as provide general arguments, to try to persuade you that when you look around, the impression you have that you are creating a large panoramic picture in your head is a total illusion. When you *see*, there is no intermediate *picture* of the world anywhere, your subjective impression notwithstanding; there is just the world and your interpretation of the incoming information (and there are also mechanisms for keeping the two in correspondence, which is the topic of chapters 4 and 5). There is, however, a representation of the visual world; and here the great J. J. Gibson was wrong in trying to develop a theory of direct perception, in which seeing is unmediated by representations and reasoning. But he was right that there is no pictorial object, no stable, spatially distributed, topographically organized representation that we would call a picture, at least not by the time the visual information becomes available to our cognitive mind.

Chapter 2 goes on to make the argument that although the visual system is "smart"—honed by many millennia of evolutionary shaping—it is nonetheless essentially ignorant. The visual system (or at least the part of it I will focus on, the so-called *early-vision* system) does what it was designed to do with neither help nor hindrance from the cognizing mind. How it can do as well as it does in the face of the inherent uncertainty of the incoming visual data is one of the questions that has been

explored most successfully in recent decades; it is one of the outstanding achievements of contemporary cognitive science. According to the thesis developed in chapters 2 and 3, there are two primary ways in which the mind (or, as I will say, the cognitive system) affects visual processing. One is that it is able to control where it concentrates its efforts or its limited resources. It does this by determining where to look as well as where to focus attention. The second way is by considering the product of visual processing, together with everything else the whole organism knows and believes, to figure out what is actually in the visual scene. Chapter 3 continues this discussion and goes on to describe some of the recent findings from neuroscience and psychophysics concerning the nature of this largely automatic early vision system.

Chapters 4 and 5 take a second look at the idea introduced in these last two chapters—that focal attention plays an important role in connecting vision and cognition—and suggests that there must be a mechanism, closely related to attention, that also plays a crucial role in connecting visual representations with things in the world. This special connection, which I refer to as visual indexing, allows parts of visual representations to be bound to parts of the visual world so that they can refer to these parts directly, rather than in a way that is mediated by an encoding of the properties that these things have. Visual indexes serve much like what is called demonstrative reference, the sort of reference that we might make in language when we refer to *this* or to *that* without regard to what the *this* and *that* are or what properties they may have. We *pick out* and individuate primitive visual objects as a precursor to focusing attention on them and encoding their properties. Such visual indexes play an important role in attaching mental particulars to things, and they also play a role in allowing mental images to inherit some spatial properties from perceived space. The ideas introduced in chapter 5 fill in some of the missing aspects of the sort of symbolic representations of percepts introduced in chapter 1 and also explain some of the alleged spatial properties of mental images discussed in subsequent chapters.

Finally, the last three chapters build on the ideas introduced in chapters 1 through 5 to make sense of some of the puzzling aspects of mental imagery. The message of these last chapters is that, notwithstanding what it feels like to visualize or to examine a mental image in one's mind's eye,

imagining and visualizing are a form of reasoning. And, as with all forms of reasoning, we have little conscious access to the real machinery of the mind, the machinery that encodes our thoughts and that transforms them as we reason, as we draw inferences, as we search our memory, and as we understand the evidence of our senses—including that which arrives in the form of language. Because our intuitions and our introspections are either silent or seriously misleading concerning what our thoughts are like and how they are processed, we need to impose some constraints on our speculations on the subject. In the last three chapters, and especially in chapter 8, I consider these constraints and, drawing on work by Jerry Fodor (1975) as well as our joint work (Fodor and Pylyshyn, 1988), I argue that any form of reasoning, including reasoning by visualizing, must meet the constraints of productivity and systematicity. This leads to the inevitable conclusion that reasoning with mental imagery or reasoning by visualizing or "visual thinking" requires a combinatorial system—a language of thought—that itself is not in any sense "pictorial." Although this conclusion may appear to fly in the face of an enormous amount of evidence collected over the past thirty years showing that thinking using mental images is more like seeing pictures than it is like reasoning with an "inner dialogue," I argue in detail in chapters 6, 7, and 8 that the evidence does not support the assumption that visualizing is special in the way that current theorists have assumed it to be. The data, I will argue, have been widely and seriously misunderstood.

Even if the reader agrees with the general thrust of this discussion, there will no doubt remain many areas of legitimate discomfort. Some of this discomfort concerns the apparent discrepancy between what it feels like to see and to visualize and the nature of vision according to the class of symbolic computational theories that have been developed in cognitive science. This discrepancy is addressed in various places in the book, but in the end it cannot be dismissed, simply because of our deep-seated commitment to what Dan Dennett has called the "Cartesian Theater" view of the mind, a view that is deeply embedded in our psyche. This "explanatory gap" was around long before cognitive science and the computational view of mind (in fact it was the subject of considerable debate between Nicolas Malebranche and Antoine Arnauld in the seventeenth century; see Slezak, 2002). In addition to this deeply held worldview, there are also other aspects of mental imagery where the

discomfort owes more to the way we conceive of certain problems than it does to how the contents of our conscious experience strike us. I consider one class of these in chapter 8, where I discuss the close connection that may exist between imagery, imagination, and creativity. I will argue that this connection, although real enough, does not favor one view of the nature of mental imagery over another. It rests, rather, on our ignorance of the nature of creativity and on the puzzle of where new thoughts and ideas come from.

Although the book devotes considerable space to the salubrious task of clarifying many misconceptions that pervade the study of vision and visualization, the principal goal is to provide a new analysis of the nature of visual processing and of mental imagery. The analysis rests on both empirical evidence, some of which comes from the author's laboratory, and on reinterpretations of well-known findings. This book is in many ways a continuation of the investigation of cognition that began with my *Computation and Cognition* (Pylyshyn, 1984a), but it focuses on problems within one of the most highly developed areas in cognitive science, visual perception, and traces the relation between the study of vision, the study of mental imagery, and the study of cognition more generally.