Preface

The research program I describe in this book began quite inadvertently, during an experiment on the organization of "semantic memory," when I noticed that two people in a row had judged as false the statement "A flea can bite." I made it a habit to talk to the subjects after each session and to review with them any peculiar responses. When asked about this strange mistake (I was convinced that everyone, even people who had never owned a dog or cat, knew fleas could bite!), one subject said something like "I looked for a mouth, but couldn't find one." The other subject said "I looked, but couldn't see any teeth." Looked? See? What is this, I wondered, this talk about visual imagery? Imagery did not seem to fit in at all with the semantic net models I had been thinking about, and it was not clear how factors that should affect the organization of such a net would affect using imagery. Thus, I asked all my subjects whether they had tended to use mental imagery in deciding whether the statements presented to them were true or false, and I graphed the results separately for people who claimed to use imagery and for those who did not. Amazingly, the patterns were wildly different. For those who professed not to use imagery, decision times decreased with increasing association strength between a noun and a property, as I had expected (if association strength reflects "distance" between two representations in some kind of net). But for those who used imagery, times were not affected by association strength at all! (Years later I found out why: In this earlier experiment less strongly associated properties also tended to be larger, and larger properties are more easily seen on images—as is discussed in chapter 7). In the next experiment, Keith Nelson and I tested two groups of subjects, asking one group to use images in evaluating statements and giving the other group no instructions about strategy. These results again showed different effects of association strength depending on whether imagery was used. It was not clear to me where to go from there, but three other fortuitous

events came together and ended up defining the present research program.

At the same time I was measuring people's reaction times and asking about the bases for their decisions, I was pursuing my long-standing interest in developmental psychology. I read Bruner, Olver, and Greenfield's book, in which the authors claimed that young children rely on imagery in their thinking more than adults do. This was an interesting idea but seemed hopelessly intractable given our dearth of knowledge about imagery. To track the bear, one needs to know what its footprints look like; without knowing what it means to "use imagery," how can one know what would be the behavioral consequences of using imagery in one's thinking? In chapter 10 I discuss the literature on this "representational-development" hypothesis and argue that the present enterprise not only has clarified the nature of the question being asked but has provided the beginnings of an answer to it.

The second major root of the present project was a single line in Gordon Bower's paper, "Mental Imagery and Associative Learning," which was later published in a highly revised form in 1972. In the early version, Bower was describing some of his introspections: "Frequently only part of the scene is 'focussed' as figure, the rest being indeterminant ground until we fill in or generate another part of the scene, much as we move our eye over actual scenes." (This line was altered beyond recognition [p. 57] in the final version.) When I read this sentence, an experiment immediately sprang to mind: If Bower's observations were true, it should require more time to "see" parts of images that were not in focus at the time of query than parts that were in focus. I was struck by the idea that if images were in some sense structurally like pictures, then they should depict spatial extent and more time should be required to traverse greater distances across them. Oddly enough, when I designed and conducted my first scanning experiment (described in chapter 3, section 2.1), I was only vaguely aware of the stunning work on mental rotation that Roger Shepard and my fellow students Jacky Metzler and Lynn Cooper were conducting right down the hall.

After publishing the scanning experiment in 1973 I lost direction; it was not clear to me where research on imagery, either mine or that of the field in general, was leading. The catalyst that truly launched the present project was the idea that images are like displays a computer can generate on a cathode-ray tube. On this view, images are "surface representations" that are generated from more abstract underlying representations. Although this idea was inspired in part by Chomsky's work in linguistics, it was primarily a simple extension of a class assignment in my first computer programming course, taught by Jerry Feldman. We were asked to write a program to draw various geometric figures using a set of shapes. The shapes themselves were generated by sets of procedures, and the fig-

ures were composed by printing out different shapes in various combinations and relative positions. It gradually dawned on me that I had the beginnings of a model of human mental imagery in my class project.

Thus, the earlier work provided a set of techniques for studying imagery, primarily based on simple extensions of the kinds of sentence verification tasks used to study semantic memory. Further, my early interest in semantic memory shaped the kinds of questions I wanted to answer, questions about how images serve as repositories of information, about how we extract information from images, and about the situations in which it is advantageous or necessary to use mental imagery. The cathode-ray tube metaphor provided a way of thinking about these questions and the kinds of answers one might find. This metaphor also guided and directed the research strategy that has structured the project to date.

I have dedicated this book to my parents, and I mean this in the broadest sense of the term. I cannot thank my biological parents, S. Duke and Rhoda Kosslyn, enough for all they have done to shape me and hence my work. I also am at a loss to extend adequate thanks to my "intellectual parents." In particular, I owe more than I could ever express to John P. Seward, who employed me in his animal lab during my undergraduate years and showed me what it is to be a scientist and a truly civilized human being. Edward Sadalla also was a source of intellectual stimulation during my undergraduate days. During graduate school Gordon Bower, Herbert Clark, and Edward Smith all spent generous amounts of time with me and much of my education occurred during informal interactions with these gentlemen, as well as with other members of the Stanford University community, notably Susan Haviland, Dan Osherson, John Anderson, Keith Holyoak, Arnold Glass, and Eve Clark. Many of my current ideas are direct outgrowths of discussions and arguments with these people.

I also wish to extend my deepest appreciation to George E. Smith, who meticulously read an early version of the book and devoted one evening a week to me for months on end. George asked hard questions and helped me struggle with the answers. I am also indebted to R. Duncan Luce for carefully reading the entire manuscript and offering deep and penetrating comments. Duncan pointed out conceptual, stylistic, and organizational problems, and deserves credit for helping make the book more coherent and readable. Keith Holyoak, Steven Pinker, and Steven Shwartz performed similar services, and did not spare me when they spotted a factual or a conceptual error. Discussions with Steven Pinker proved especially helpful when I was beset by conceptual confusions, and I wish to thank him for his time, energy, and insights. Conversations with Jerry Fodor also proved very stimulating, and chapter 11 is written around some of his illuminating ideas. Eric Wanner, my editor and friend, more than once forced me to see the error of my ways and is responsible for improving the

book greatly over its earlier form. Eric helped me figure out how to put first things first, and I am much appreciative. Camille Smith proved patient and tolerant in editing the manuscript.

I also want to thank all of my collaborators, who are cited in the appropriate places in the text. Without their help, I could not have accomplished a tenth of what is described herein. I am especially grateful to my research assistants Brian Reiser, Martha Farah, and Sharon Fliegel (who have since gone on to bigger and better things), who came through when it counted. Bill Shebar and Roger Wallach helped me dig out references and forced me to think about them. Steven Shwartz, my collaborator in the simulation project, deserves much of the credit for the elegance of the model. Willa Rouder, Debbie Lathrop, Angela O'Neill, and Loretta Czernis suffered through typing the manuscript into the computer and altering it too often to count. Scott Bradner was absolutely invaluable in getting the book off of my computer tapes and onto these pages.

Finally, I want to thank the institutions that nurtured and supported my work. The National Science Foundation, primarily through the good offices of Joseph Young, has been magnanimous in providing me with support (Grants BNS 76–16987, BNS 77–21782, and BNS 79–12418). Stanford University, Johns Hopkins University, and Harvard University have generously given me space and support. Further, in each university I have found the kind of intellectual and emotional environment conducive to the kind of research I do and the kind of life I live.

S.M.K.

Cambridge, Massachusetts
1980