

1

Introduction

Concepts are the glue that holds our mental world together. When we walk into a room, try a new restaurant, go to the supermarket to buy groceries, meet a doctor, or read a story, we must rely on our concepts of the world to help us understand what is happening. We seldom eat the same tomato twice, and we often encounter novel objects, people, and situations. Fortunately, even novel things are usually similar to things we already know, often exemplifying a category that we are familiar with. Although I've never seen this particular tomato before, it is probably like other tomatoes I have eaten and so is edible. If we have formed a concept (a mental representation) corresponding to that category (the class of objects in the world), then the concept will help us understand and respond appropriately to a new entity in that category. Concepts are a kind of mental glue, then, in that they tie our past experiences to our present interactions with the world, and because the concepts themselves are connected to our larger knowledge structures.

Our concepts embody much of our knowledge of the world, telling us what things there are and what properties they have. It may not seem to be a great intellectual achievement to identify a bulldog or to know what to do with a tomato, but imagine what our lives would be like without such a conceptual ability (Smith and Medin 1981, p. 1). We might know the things we had experienced in the past—a particular chair, our bed, the breakfast we had today, our science teacher, etc.—but when we encountered new exemplars of these categories, we would be at a loss. When going into a new room and seeing a new chair, we would have to study it from scratch, attempt to determine whether it is alive or dead, what its function is, whether it will hurt us, or how it might help us. Instead, of course, we may not even consciously think “chair,” but simply identify the object's category and plop down into it. By using our concept of chairs, we immediately draw the inference that it is appropriate to sit on this object, even if we have never seen anyone sit on it before. At a new

restaurant, we read names of dishes such “gnocchi,” “jerk chicken,” and “pad thai” and feel we can decide which one we would prefer to eat, even though we have never had that exact meal, or even an example of that kind of meal at this restaurant. The speed and ease with which we identify objects as chairs or draw inferences about jerk chicken (too hot to have for lunch) can mislead us about how complex this process is and how much information we may have stored about everyday categories. The psychology of concepts is like other areas of psychology, in which a phenomenologically simple cognitive process, like understanding speech or walking, turns out to be maddeningly complex. Much of the excitement in the field arises from this complexity, as a topic that seemed to be fairly straightforward in 1960 has turned out to be a much deeper and richer scientific problem than researchers had expected.

The mental glue provided by concepts applies not only to the familiar categories of objects, like chairs and tomatoes, but also to a number of other domains that are of interest to psychologists, such as social and person categories, emotions, linguistic entities, events and actions, and artistic styles. For example, if we meet a new, highly talkative person and begin to suspect that he or she is a bore or instead a sociopath, our behaviors toward the person will differ accordingly. If told by someone else that the person is a lawyer or instead a priest, our behaviors will again differ. We rely on such categories to direct our behavior, sometimes despite more reliable information directly observed about the person.

The psychology of concepts cannot by itself provide a full explanation of the concepts of all the different domains that psychologists are interested in. This book will not explore the psychology of concepts of persons, musical forms, numbers, physical motions, and political systems. The details of each of these must be discovered by the specific disciplines that study them; to fully understand people’s musical concepts will require much research into the psychology of music, rather than being predictable solely from what we know of concepts per se. Nonetheless, the general processes of concept learning and representation may well be found in each of these domains. For example, I would be quite surprised if concepts of musical forms did not follow a prototype structure (chapter 2), did not have a preferred level of categorization (chapter 7), and did not show differences depending on expertise or knowledge (chapter 6). Spelling out what categories people have of musical forms, what levels of abstraction there are, and what knowledge influences the concepts is primarily part of the psychology of music rather than the psychology of concepts. But once the basic elements of musical concepts are identified, the concepts will likely be found to follow the principles identified in other domains. In

short, the psychology of concepts has much to offer other fields of psychology and cognitive science more generally.

Similarly, concepts are ubiquitous across different populations and ages—it is hard to see how any intelligent creature could do without them. It used to be thought that infants and young children were lacking in true conceptual abilities, which had to be onerously acquired over the preschool years. However, more recent research has found basic conceptual abilities in infants only a few months old (chapter 9), and preschool children now appear to have sophisticated conceptual abilities, even if they are lacking much of the conceptual content that adults have (chapter 10).

Another way that concepts infiltrate our everyday life and thoughts is through communication. When we talk, we are attempting to communicate ideas about the objects, people, and events that take place around us. Since we understand those objects, people, and events through concepts, our word and sentence meanings must make contact with conceptual representations. Not surprisingly, it turns out that many properties of concepts are found in word meaning and use, suggesting that meanings are psychologically represented through the conceptual system (chapters 11 and 12).

There is a real temptation for researchers in the field of concepts to get carried away on the “everything is concepts” bandwagon that I have started rolling here. (If I were more melodramatic, I could spin a fairy tale in which a person who has no concepts starves while surrounded by tomatoes, because he or she had never seen *those particular* tomatoes before and so doesn’t know what to do with them.) Although in unguarded moments I do think that everything is concepts, that is not as restrictive a belief as you might think. Concepts may have a great variety of forms and contents, and this is part of what has made the field so complex. Across different people, levels of experience with the category, tasks, and domains, concepts may vary in a large number of ways. Although this is not itself a principle of the psychology of concepts, many examples of this variation will be seen throughout the book. Reconciling those differences, or at least understanding how different kinds of concepts are coordinated, is an important goal of this field, one that has not been fully accomplished yet.

The psychology of concepts, then, has the goal of understanding the representations that allow us to do all these things, most importantly, identifying objects and events as being in a certain category, drawing inferences about novel entities, and communicating about them. Although the field (and hence this book) concentrates on common object concepts, the principles involving concept formation and use are thought to be to some degree generalizable across different domains and settings.

How This Book Differs from the Book You Expected to Read

Most books and review articles that I know of on concepts have used an organization based on theories. The best example is the classic text by Smith and Medin (1981). After some preliminaries, it had chapters on the classical view of concepts, each of three probabilistic views, and then the exemplar view. Each chapter described the theory, presented evidence relevant to it, and then evaluated it. In 1981, this was an excellent organization. In 2001, however, this organization would have a number of problems. The first is that the field shows very little agreement on which theory of concepts is correct. One might have hoped that twenty years after Smith and Medin's review, the field would have sorted out many of the issues their book raised. However, there is as much, and perhaps more, dissension now as there was then. Focusing on theories, therefore, is not the best way to document the important progress that has been made in the psychology of concepts. Many interesting principles and generalizations have been discovered about concepts, and even if the field does not agree on the overarching theory that encompasses all of them, that does not deny that those discoveries have been a real advance.

The second reason not to organize this book around theories is that too many interesting questions do not fall easily into theoretical pigeonholes. Issues such as infant concept learning or conceptual combination are cohesive topics in their own rights, but they are difficult to parcel out to chapters on the classical theory or exemplar theory. If one were to divide up the parts of each topic that are most relevant to a given theory, one would have balkanized a previously coherent topic.

The third reason I have not followed the theoretical organization is that I am becoming increasingly uneasy about the particular theoretical disputes that have characterized the field. Much of the literature has compared exemplar and prototype theory (see chapter 3), but it seems fairly clear that both theories are wrong to a greater or lesser degree, and so focusing on them serves to reinforce this way of dividing up the field, when new ways of thinking may be required. In particular, I will suggest at the end of the book that a more inclusive approach may be needed, an approach that would be ill-served by organizing the book around distinguishing the different theories.

As a result, I have organized most of the book around phenomena or issues rather than theories. The book does begin theoretically, starting with the so-called classical theory of concepts and its downfall. The three main approaches to concepts are described in the next chapter, as a preparation for the discussions in individual

chapters. As a result, chapters 2 and 3 should be read first (except perhaps by those already familiar with the field), and after that, readers may graze on the subsequent chapters in any order. Each chapter describes a particular topic or set of related phenomena and evaluates the explanations of those phenomena. At the end of each chapter, I discuss the implications of those phenomena and their explanations for the three main theoretical approaches discussed in chapter 3. The final chapter revisits the main theories and makes proposals for future directions. Thus, theoretical issues are by no means ignored in the book. However, one discovery I have made in organizing the book this way is that no theory has a ready explanation for all of the findings even *within* each specific topic. This helps to point out goals for researchers who take a given theoretical approach, but it also serves to point out the limitation of organizing the field primarily by the theories.

This book is not a complete reference work of our knowledge of concepts. It is a selective review, and I sometimes do deviate from the majority of the field in my choices of what I think is most significant or interesting. My goal here has not been to evaluate everything that has been done but instead to explain the most basic and exciting findings in the field and to try to draw some conclusions about how to explain them. I have attempted to take the “long view” and not to necessarily include the hot topics of this moment, much less the abandoned hot topics of yesteryear. For the most part, then, I have focused on the topics about which enough work has been done to draw a conclusion, rather than on topics that are still unsettled and perhaps of unclear future interest.

Terminology and Typography

In general, I try to use the word *concepts* to talk about mental representations of classes of things, and *categories* to talk about the classes themselves. However, in both everyday speech and the literature in this field, it is often hard to keep track of which of these one is talking about, because the two go together. That is, whatever my concept is, there is a category of things that would be described by it. Thus, when talking about one, I am usually implying a corresponding statement about the other. Writers in this field often say things like “four-year-olds have a category of animals,” meaning “four-year-olds have formed a concept that picks out the category of animals.” However, being too fussy about saying *concept* and *category* leads to long-winded or repetitious prose (like my example) with little advantage in clarity. When it is important to distinguish the mental representation from the

category itself, I will be careful. In other situations, I will leave it to the reader to make the appropriate translation (e.g., from “having a category” to “having a concept that corresponds to a particular category”).

Many writers also use typographic conventions to distinguish concepts (or categories) from the actual things. Obviously, a concept of a dog is not the same as a dog, and some authors indicate this by printing concept names in italics or small capitals. I have gone this route in the past, but I decided not to do so here. This issue is related to an underlying assumption about the relation of concepts and categories, so I will briefly explain why I have come to this decision.

In many cases, it is simply obvious whether one is talking about the concept or the object, and so the typographical convention is not necessary. In other cases, though, the ambiguity between the concept and the thing is intentional. Indeed, I am not always sure that authors get it right when they decide to capitalize the word. For example, suppose you learn that dogs bark. Have you learned that dogs bark or that DOG (the concept) has the property “barks”? From one perspective, you have learned about a property in the world, and so you have learned something about actual, lower-case dogs. From another perspective, you have changed your concept, and so you have modified your mental representation, that is, DOG. In fact, you have probably learned both. But when one follows such a distinction, choosing one typography implies that the other one is not intended. That is, if I were to say something about dogs, I would not intend you to understand this to be true of the concept of dogs, because I didn’t write DOG. But often I, in fact, *would* intend to be making the statement about both: The parallelism between concepts and categories means that when you learn something about dogs, your concept of dogs has also changed, and so constantly making this distinction leads to false implications. So, I do not use a separate typography for concepts and real things, but instead simply say “the concept of dog” or the like when I want specifically to discuss concepts.

I follow standard practice in linguistics by italicizing cited words, as in “the word *dog* has one syllable.” Things that people say are quoted.

A Note to Students

I have attempted to write this book at a level that an advanced undergraduate or beginning graduate student could understand. Although I assume general knowledge about experimental psychology and some familiarity with cognitive psychology, I have tried to start from scratch when it comes to concepts, so that little or no

knowledge is assumed other than what has already been covered in the book. In fact, I have erred on the side of redundancy, so that chapters would be more self-contained. If you are not very familiar with the “traditional” literature of concepts (i.e., from the 1970s), then the first few chapters will be necessary reading. Even seasoned concepts researchers may be surprised to find themselves interested in these chapters, as it is often the most basic phenomena that are the most difficult to explain.

A student picking up this book might wonder whether there are still interesting questions to be asked in the psychology of concepts. Have the basic questions been answered, and are we only spelling out the details now? The book’s conclusion is one answer to this question. Before getting there, however, I think I can safely say that the answer is “no.” There are still important and surprising discoveries being made in the field. Many of these are coming through attempts to look at real-world concepts in greater detail, in which the findings are sometimes quite different from what would be expected from studies with artificial categories. Related to this is the benefit that concept research is receiving from connections to other areas of cognitive science such as psycholinguistics, reasoning, anthropology, neuropsychology, and problem-solving. Working out how concepts are influenced by and in turn influence processes in these other domains has been a fruitful and in some cases surprising enterprise. However, before a new researcher can engage in such cross-discipline interactions, he or she must understand the basic issues and findings in the field, and this is what the next few chapters discuss.

Some students may have more basic questions of the sort that one is reluctant to ask one’s teacher or adviser. I find these illustrated in a letter written by an amateur student of astronomy to observers at the Mount Wilson Laboratory in 1933 (Simons 1993, p. 113):

Just a few lines to let you know that I am Interested in Astronomy. I have did quite a lot of reading on it and I am really interested in it. I have quite a bit of confidence in Materialism; I believe myself the whole Universe is substance. But what I would really like to know is will Astronomy get a person anywhere—is there any use in a person studying it. Will it put you in an unmentally condition?¹

These are questions I have asked myself about the psychology of concepts as well. As to whether studying concepts will “get a person anywhere,” it of course depends on where you want to get. I don’t think it will get you a high-paying job in the new economy, but I think it may help you to understand a basic function of the mind. If you want to understand social thinking or perception or cognitive development,

then learning about concepts will be a necessary part of your study. Another reason studying concepts may get you where you want to go is the incredible variety of research on concepts. It ranges from the mathematical models tested in artificial category-learning experiments to anthropological studies in rain forests to linguistic analyses of word and phrase meaning. It is possible to find your own niche within the study of concepts almost regardless of what your interests are. If you find one part of it (and of this book) too boring or too touchy-feely or too technical, you can wait a chapter or two and discover a topic and approach that are completely opposite.

The question of whether getting involved in this topic will put you in “an unmentally condition,” is a trickier one. Studying almost any question in psychology is bound to have its moments of mental unbalance, as past beliefs are called into question, and issues that were thought to be settled twenty years ago come back with renewed force. Overall, I believe that the risk is no greater here than in most areas of psychology. However, the risk of becoming unbalanced is clearly greatest in the prototype-exemplar theory debate, which shows little signs of abating after many years of controversy, so those who feel mentally vulnerable may wish to focus on other topics.

The writer of that letter concluded by saying:

But I know that the more you read up on it the more you get Interested. . . . Would you please give me some kind of basis to the Knowledge of astronomy?

My expectation is that the more you read up on concepts the more you will get interested in them as well, and the goal of this book is not so much to tell you all about concepts as to provide some kind of basis to your continuing acquisition of knowledge. That is, if you understand the sections on typicality and concept learning, for example, you should be able to pick up a contemporary paper on these topics and understand it. In serving that function, I hope that this book will have a shelf life that is longer than its function of describing recent discoveries, just as the Smith and Medin (1981) book could be used as background reading in the area for many years after its publication. The field has greatly expanded since 1981, covering many topics that did not really exist then; and much more has been learned about the topics that their book did cover. Therefore, this book is correspondingly bigger and more detailed than Smith and Medin’s. My guess is that in twenty more years, it will be impossible to write a single-authored book that covers the same ground, because there will be more research than any one author and volume can handle. That is exactly why I have written this book now.

Acknowledgments

Is there any section of a book more boring than the acknowledgments? After one has searched in vain for one's own name, there is little point in continuing to read. Nonetheless, I owe a great debt of gratitude to a number of people and institutions who have supported the writing of this book, and I only wish that I could force every reader to read each name and fully understand the thanks that I owe to each and every one of them. That being impossible, I am forced to be brief.

First, I need to thank the University of Illinois, and especially its Center for Advanced Studies, for providing me with a sabbatical leave on which I researched and wrote most of the first draft of the book. The Beckman Institute, through its administrative and secretarial support (in the person of the invaluable Linda May), also made the writing much easier for me. The City University, London, provided a supportive home for me during that sabbatical and also for a month in 2000 when I did much revision and polishing. I am especially grateful to my host at City, James Hampton, for selfless efforts on my behalf.

Second, I am grateful to the National Institute of Mental Health for supporting my research reported in this book, as well as part of the writing itself (grant MH41704). Other funding from the National Science Foundation (a grant on which Brian Ross was principal investigator) was helpful as well (grant SBR97-20304).

Third, a number of very busy people took time out from their other duties to critique and comment on various chapters. The book is much better as a result, and would have been even better if I had taken all their advice. These people include Renée Baillargeon, Paul Bloom, Seth Chin-Parker, Herbert Clark, John Coley, James Hampton, Evan Heit, Devorah Klein, Barbara Malt, Arthur Markman, Kristine Onishi, Paul Quinn, Bob Rehder, Brian Ross, Paula Schwanenflugel, Edward Smith, J. David Smith, Sandra Waxman, Edward Wisniewski, and Karen Wynn. James Hampton, Kristine Onishi, Brian Ross, and especially Arthur Markman deserve combat pay for reading multiple chapters.

Fourth, I thank my former colleagues and students at Illinois, where I conceived and did much of the work on this book, for helping to create a stimulating and rich intellectual environment in which to work and write. If I list one name, I would have to list dozens, so I will just thank you all.