

Categories and Concepts

Edward E. Smith and Douglas L. Medin

Contents

1. Introduction	1
2. Preliminary Issues	7
3. The Classical View	22
4. The Probabilistic View: Featural Approach	61
5. The Probabilistic View: Dimensional Approach	102
6. The Probabilistic View: Holistic Approach	130
7. The Exemplar View	143
8. Summary and Implications	162
Notes	185
References	191
Index	201

Harvard University Press
Cambridge, Massachusetts
London, England
1981

1 | Introduction

WITHOUT CONCEPTS, mental life would be chaotic. If we perceived each entity as unique, we would be overwhelmed by the sheer diversity of what we experience and unable to remember more than a minute fraction of what we encounter. And if each individual entity needed a distinct name, our language would be staggeringly complex and communication virtually impossible. Fortunately, though, we do not perceive, remember, and talk about each object and event as unique, but rather as an instance of a class or concept that we already know something about. When entering a new room, we experience one particular object as a member of the class of chairs, another as an instance of desks, and so on. Concepts thus give our world stability. They capture the notion that many objects or events are alike in some important respects, and hence can be thought about and responded to in ways we have already mastered. Concepts also allow us to go beyond the information given; for once we have assigned an entity to a class on the basis of its perceptible attributes, we can then infer some of its nonperceptible attributes. Having used perceptible properties like color and shape to decide an object is an apple, we can infer the object has a core that is currently invisible but that will make its presence known as soon as we bite into it. In short, concepts are critical for perceiving, remembering, talking, and thinking about objects and events in the world.

The point of this discussion is that a great deal in psychology hinges on how people acquire and use concepts, which in turn depends on the structure of concepts. And lately many psychologists seem to be changing their views about such structure. Until recently the dominant position – which we will call the *classical* view – held that all instances of a concept shared common properties, and that these common properties were necessary and sufficient to define the

concept. This view, which dates back to Aristotle, has always had its critics, but in the past decade the criticisms have become more frequent and intense, and new views have emerged. Perhaps the most prominent of these assumes that instances of a concept vary in the degree to which they share certain properties, and consequently vary in the degree to which they represent the concept. This view has sometimes gone by the name of *prototype*, but since this label has been used to mean many different things, we prefer to call it the *probabilistic* view. Another emerging view, which offers an even more extreme departure from the classical one, holds that there is no single representation of an entire class or concept, but only specific representations of the class's exemplars. This we call the *exemplar* view.

The development of alternatives to the classical view has been accompanied by a bustle of research activity. Indeed, the activity has been so frenzied that we think it is time to take a hard look at what has been learned. This book attempts to provide that look. It offers a systematic analysis of the three views of concepts, the processing models they have generated, and the empirical findings that the views and models endeavor to capture. In particular, we will emphasize the major problems that are responsible for the recent shift away from the classical view, show how the newer views take account of these problems, and indicate some of the costs incurred by such an accounting. But before beginning such a survey, it will be useful to provide some introductory examples of the three views.

To illustrate the classical view, we can consider the geometric concept of squares. Suppose that people in general represented this concept in terms of four properties: (1) closed figure, (2) four sides, (3) sides equal (in length), and (4) angles equal. Since these four properties, or criteria, would be applied to any object whose squareness is at issue, we have a unitary description of the concept "square." Moreover, the four properties that make up this concept are precisely those that any square must have. Roughly, then, to have a classical-view concept is to have a unitary description of all class members, where this description specifies the properties that every member must have.

Generalizing this approach, let us try to come up with a plausible classical-view concept that people might have of a cup. Such a concept might consist of the following five properties: (1) concrete object, (2) concave, (3) can hold liquids, (4) has a handle, and (5) can be used to drink hot liquid out of. These properties offer a unitary description of cups, but are all the properties true of everything people would call a cup? Properties 1-3 seem to be, but 4 and 5

are debatable. The teacups commonly used in Chinese restaurants typically do not have handles, yet they are still cups; and one can certainly imagine a poorly manufactured cup which conducts heat and so is useless for drinking hot liquids, but which we would still call a cup. But if we give up the last two properties, we are left with properties 1-3, which are true of some non-cups—bowls, for example—and so do not uniquely describe cups. Considerations like these led Labov (1973) to argue that people's concept of a cup does not conform to the classical view because the properties that make up the concept are not common to all members. Then what kind of view would capture people's concept of cup? According to Labov and others, one needs a view that posits a unitary description of cups, but where the properties in this description are true of most though not all members. This is the probabilistic view. Given such a view, some instances of the class are going to have more of the critical properties than others, and those that do will seem more representative of the concept.

To illustrate a third possible concept, let us consider the class of all psychiatric patients who have suicidal tendencies. One could try to construct a classical-view concept that clinicians might have for this class, but years of futile effort suggests that this cannot be done. What about a probabilistic concept—can it capture a clinician's concept of the suicide-prone? No doubt a probabilistic concept would work better, but even it might fail to capture the knowledge that a clinician uses in deciding that a particular patient is suicidal. For a probabilistic concept provides a single description of all people with suicidal tendencies, yet clinicians may sometimes decide that a patient is suicidal by comparing him to other specific patients known to be suicidal (this suggestion is taken from Tversky and Kahneman, 1973). That is, the class of people with suicidal tendencies may be represented not by a single description, but rather by separate descriptions for various patients known to be members of the class. This corresponds to the exemplar view of concepts.

The examples given above suggest that the various views of concepts can be partly understood in terms of two fundamental questions: (1) Is there a single or unitary description for all members of the class? and (2) Are the properties specified in a unitary description true of all members of the class? The classical view says yes to both questions; the probabilistic view says yes to the first but no to the second; and the exemplar view says no to the first question, thereby making the second one irrelevant. These distinctions, which we will refine and augment later in the book, are summarized in Figure 1.

Though our principal concern in this book is with the three




Figure 1 Three views of concepts

views, there are important preliminary matters to be considered first. Some of these have to do with what we consider a concept to be in the first place; others concern the nature of the properties that are used to describe concepts. These preliminary issues will be discussed in Chapter 2. In Chapter 3 we discuss the classical view. After describing it in detail, we review numerous arguments and experimental findings that have been offered as evidence against it. We will emphasize that the classical view is a theory about representations, and that to evaluate the view against experimental findings one must add processing assumptions to it, thereby converting a theory about concepts into a model about categorization. In Chapters 4, 5, and 6 we move on to the probabilistic view. These three chapters correspond to the method of describing concepts in the probabilistic view—by qualitative features, quantitative dimensions, or holistic patterns—for the three modes of description have somewhat different consequences. Again we will be concerned not only with the view itself, but with processing models that have been generated from it and with their ability to account for empirical findings. Chapter 7 deals with the exemplar view, again with an eye toward the specific models that instantiate the view and the findings they purport to explain. Finally, Chapter 8 summarizes our main conclusions and raises some unexplored issues.

What exactly are the empirical phenomena with which we intend to measure the views and models? Our major source of phenomena deals with *categorization*. Specifically, we will be concerned primarily with data on how adults use natural concepts with one-word names to classify things. Such “things” might be pictures of objects, say of a dog or a cup, and subjects might be asked if a pictured object belongs in a particular category. Or the things to be categorized might be denoted by words, as when subjects are asked, “Is a raisin a fruit?” and timed as they make their decisions.

(Our reasons for so emphasizing the categorization aspect of concepts will be discussed further at the beginning of the next chapter.)

In almost all cases we will be concerned with object concepts—animals, plants, human artifacts, and so on. Our main reason for choosing this domain is that it has been the most extensively studied in the last decade of experimental research. There is also an ancillary benefit of working with this domain—namely, it is a particularly interesting test case for the three views. That is, had we chosen geometric concepts, like “square,” as a target domain, we might have prejudiced things in favor of the classical view; for we know that mathematicians have constructed classical-view descriptions of these concepts and have at least partially succeeded in inculcating these descriptions into many people’s conceptual lives. Similarly, had we chosen as our domain abstract concepts, such as “love” or “brilliance,” we might have prejudiced the case against the classical view; for no mathematician or metaphysician has come even close to constructing a classical-view description of such concepts. Thus natural objects and human artifacts offer an in-between case—between concepts that any schoolboy can define and concepts that no scholar can grapple with.)

There is another rationale for concentrating on object concepts.¹ There are two questions we can ask about any potential psychological concept: (1) Can it be given a classical-view definition in *any* language? and (2) Can it be given a classical-view definition in the language of the mind—that is, do people have a mental representation that corresponds to a classical-view definition of the concept? While question 2 is clearly a psychological one, question 1 is more the province of philosophy and logic. However, question 1 has important implications for a psychological analysis of concepts, for a *no* answer to it virtually necessitates a *no* answer to question 2. (We hedge with “virtually” just in case all things in the world turn out not to have classical-view definitions but our minds impose classical-view structure on them.) It is for this reason that we ignore abstract concepts like love and brilliance: since there is good reason to doubt they can be classically defined in any language, to study whether they can be classically defined psychologically may be beside the point. In contrast, the object concepts that we will focus on seem more likely to be definable in some language, that is, more likely to provide a *yes* answer to question 1, thereby making question 2 a worthwhile topic of study.

In addition to work on natural concepts, we will also consider some studies of artificial concepts. By an artificial concept we mean an equivalence class constructed for a particular experiment, say a class of schematic faces that tend to have similar properties. There

are two reasons why such artificial classes are particularly useful for testing proposals about concepts. First, if natural concepts are hypothesized to have a particular structure, one can build this structure into an artificial class and see if people use this class in the same way they use a natural one; if they do, we have added support for the hypothesized structure of natural concepts. Second, the use of artificial concepts allows the experimenter to have precise control over the experiences or instances that the learner is exposed to in acquiring the concept. There are also drawbacks to using artificial concepts, which we will discuss later in the book. But first we need some discussion of what a concept is all about.

2 | Preliminary Issues

Functions of a Concept

THOUGH ONE'S SPECIFIC NOTION of a concept depends on which of the views one endorses, there are some aspects or functions of a concept that seem generally agreed upon. Most would agree that people use concepts both to provide a taxonomy of things in the world and to express relations between classes in that taxonomy (Woods, 1981). The taxonomic function can itself be split into the following two generally agreed upon component functions:

1. *Categorization.* This function involves determining that a specific instance is a member of a concept (for example, this particular creature is a guppie) or that one particular concept is a subset of another (for example, guppies are fish).
2. *Conceptual combination.* This function is responsible for enlarging the taxonomy by combining existent concepts into novel ones (for example, the concepts pet and fish can be combined into the conjunction pet-fish).

Similarly, the notion of using concepts to express relations can be subdivided into the following two component functions:

1. *Constructing propositional representations.* This function is at the heart of language understanding. The concepts denoted by the words in a sentence are mapped into a representation of the proposition expressed by that sentence (for example, the sentence "fish are friendly" is mapped into a proposition in which the concept of "friendly" – or a token of it – is predicated of the concept "fish").

2. *Interrogating propositional representations.* Here a relation between concepts is typically used as a basis for drawing certain inferences from representations (for example, given that one knows that guppies are fish and that fish have scales, one can infer that guppies have scales).

Though this fourfold breakdown has its rough edges, it gives some idea of the various research areas that bear on the nature of concepts. It also clearly suggests that the last three functions of concepts involve combinatorial procedures that play no role in the categorization function.

In this book we will focus almost exclusively on what we have called the categorization function. The main reason for so limiting our inquiry is that most of the critical work in psychology that bears on a shift from the classical to the probabilistic and exemplar views is categorization research. True, one can find analyses of constructing and interrogating representations that explicitly opt for the probabilistic or exemplar views over the classical one (for example, Rips, 1975; Clark and Clark, 1977; Collins, 1978), but the rationale usually given for this move is based on categorization studies. This is not to say that things have to be this way. Indeed, it is quite possible that research on conceptual combinations and on the construction and interrogation of propositional representations will ultimately be essential for constraining a view of concepts; we will return to this possibility at the end of the book. But until then, we are mainly concerned with what research on categorization can tell us about the nature of concepts.

Categorization and Inference

To say that concepts have a categorization function is to acknowledge that concepts are essentially pattern-recognition devices, which means that concepts are used to classify novel entities and to draw inferences about such entities. To have a concept of *X* is to know something about the properties of entities that belong to the class of *X*, and such properties can be used to *categorize* novel objects. Conversely, if you know nothing about a novel object but are told it is an instance of *X*, you can *infer* that the object has all or many of *X*'s properties; that is, you can "run the categorization device in reverse."

Consider the concept of a hat. Let us assume that you have two properties for this concept: (1) it has an aperture that is the size of a human head, and (2) it was manufactured with the intent of a human wearing it. If we give you a novel object and ask, "Is it a hat?" presumably you would try to determine if it has properties 1

and 2. While you can check property 1 by means of perceptual tests, you need more than that to establish whether property 2 applies. Hence, classification may require recourse to nonperceptual information, and when we say that a concept is a pattern recognition device we do not necessarily mean that it uses only perceptual information. Now suppose that instead of asking you to classify the novel object, we hide it from view, tell you that it is a hat, and ask you to tell us something about it. Presumably you would say something about an aperture that was head-sized, and something like "you're supposed to wear it." These statements would be inferences, activated by gaining access to your concept via the word that denotes it, and something like this may happen whenever you hear the word.

Our example neatly separates classification from inference, but this is an oversimplification; most classification situations involve some inferences as well. If you have a prior reason to believe an object is a hat (it is on a hat rack, for example), you might infer that it has a head-sized opening and then perform only minimal perceptual checks to confirm your inference. More generally, where context suggests that an unexamined object belongs to a particular concept, inferences may be drawn about that object's properties, and such inferences will reduce the effort that need be put into classification. A similar phenomenon can occur even without context. Having determined that an object has some perceptual properties of a hat, you might tentatively assume that the object is a hat and then infer the less perceptual properties. This is like the apple example we used earlier—having determined that an object is small, round, and red, you assume that it is an apple and infer that it has a core.

While all three views of concepts—classical, probabilistic, and exemplar—would acknowledge that concepts have the twin functions of categorization and inference, the views differ with respect to the confidence one can have in the accomplishments of these functions. According to the classical view, aside from perceptual limitations one can have complete confidence in any categorization that considers all relevant properties of the concept, for such properties offer sufficient conditions for class membership. The case is different with the probabilistic and exemplar views. In these views, one can never be certain about a categorization, for categorization is always probabilistic when there are no sufficient conditions. A similar story holds for inferences. In the classical view, the inferences are deductions because the inferred properties are necessary ones (if we tell you "it's a square," you can deduce that it has four sides); in the other two views, many of the inferred properties will be probabilistic inferences because they do not conform to nec-

essary properties. These distinctions will become clearer when we consider the views in detail.

Stability of Concepts

A good deal of work presupposes that concepts are relatively stable mental representations. They are stable in two ways: within an individual and across individuals. A concept is stable within an individual to the extent that once a person has a concept, then, except for early developmental changes or physiological ones, that person will always have the same set of properties in mind. A concept is stable across individuals to the extent that when any two people have the same concept, they have the identical sets of properties in mind. If concepts are indeed stable in these two respects, it seems reasonable to think of a concept as a very bounded unit of knowledge, something of which we can say, "He has it completely, she has it partially, and this child doesn't have it at all." This kind of thinking seems to lie behind much of the literature on conceptual and semantic development, where researchers talk of a child's acquiring and mastering concepts.

It turns out that this way of thinking is view-dependent. Believing in the stability of concepts is more consistent with the classical view than with the other two approaches. To illustrate, if one believes that concepts are represented by exemplars rather than by necessary and sufficient conditions, then frequent experience with a new exemplar, say a derby hat, can alter one's concept of hats, which constitutes a breakdown in within-individual stability. Similarly, if one person has intensive experience with a particular exemplar while another person does not, they may end up with very different concepts, a breakdown in across-individual stability.

Because of this view-dependence, we caution the reader not to prejudge issues of concept stability. If in the next chapter we sound as if we are treating concepts as stable, bounded entities, it is because we are describing the classical view, which thinks of concepts in this way. If in Chapter 7 we sometimes sound as if we think concepts are capable of constantly changing with individual experience, again it is the view under consideration — this time the exemplar view — that is dictating our description.

Kinds of Properties for Concepts

COMPONENT VERSUS HOLISTIC PROPERTIES

In talking about the knowledge contained in a concept, we have thus far used the term *property* with no attempt to specify exactly what we mean. In our attempt to unpack some of the meaning of

this term, we will rely heavily on distinctions previously drawn by Garner (1978) and Tversky (1977).

Following Garner (1978), we first distinguish between *component* and *holistic* properties of an object concept. A component property is, roughly, one that helps to describe an object but does not usually constitute a complete description of the object. Some examples should help to get this across. For one's concept of a car, the component properties might include the following: having wheels, having a motor, the average shape of a car, and the fact that its major function is transportation. Note that while some of these components refer to parts of the object (like wheels), others depict global aspects of the entire object (shape), or the purpose or function of the object (transportation). In contrast, a holistic property offers a complete description of the object. For example, your concept of a car might be represented by some sort of template of an ideal car.

Intuitive though it may be, there is a problem with this distinction — namely, how do we distinguish global component properties from holistic ones? What, for example, is the difference between a car's shape when it is treated as a global component and when it is given by a holistic template? We think there are two answers. First, usually when we talk about components there will be more than one of them, and we assume that each component is processed as a unit. This contrasts with a holistic property, like that given in a template, where one property suffices to represent the object class and that property is the only unit that gets processed. Second, when a global property like shape is treated as a component, we intend it as an abstraction from the object, not as a point-for-point isomorphism with the object. In contrast, when shape is given by a template, we often intend it more as a point-for-point isomorphism. Thus a component version of a property is usually more abstract than the corresponding holistic version.¹

Rough as it is, we think the distinction between component and holistic properties is a useful one, and we employ it in what follows.

DIMENSIONS AND FEATURES

If we decide to represent object concepts in terms of components, we have a choice of how to characterize these components — either by quantitative components, called *dimensions*, or by qualitative components, called *features*. To illustrate this distinction, let us consider one's possible concepts of weapons. You could represent all weapons in terms of a few dimensions, like degree of potential damage, with, say, penknife near one end and atomic bomb near

the other. Alternatively, you could represent each weapon by a set of features: for knife, such features might include (1) sharp, (2) has a handle, and (3) metallic. The key difference is that dimensions naturally capture quantitative variations, while features indicate qualitative ones. Thus if two concepts differ with respect to a particular dimension, one concept must have more of that dimension (a higher value) than the other; for example, one weapon is more damaging than the other. But if two concepts differ with respect to a feature, then one concept has "it" while the other does not. Or as Garner (1978) puts it, a feature is a component that either exists or does not exist.²

Each way of describing components appears to have a natural limitation: dimensions cannot handle qualitative variations while features will not work with quantitative variations. But these appearances are somewhat misleading. For featural descriptions can clearly be extended to cover quantitative variations, while a limiting case of what we mean by a dimension is extendable to qualitative variations. These ideas require some amplification.

Consider first the extendability of featural descriptions. As Atkinson and Estes (1963) pointed out some time ago, the values for any dimension may be expressed by a set of nested features. It is simplest to illustrate with the dimension of line length. Suppose that the shortest length you can detect is one millimeter; we can represent the length of this standard line by a single feature, call it F_1 . Suppose the next shortest line you can detect as different from our standard is two millimeters long; that is, a just-noticeable-difference, or jnd, is one millimeter. We represent that jnd by another feature, F_2 , and represent the second line by the feature set F_1 and F_2 . Hence the feature representing the length of the first line is nested within the feature set representing the length of the second line. By continuing in this way—representing each successive jnd by a new feature, and nesting the feature sets of shorter lines in those of longer ones—we can represent any value of the dimension of length by a set of nested features. The same procedure will work for any quantitative dimension. And the resulting feature representation captures the important intuition that the closer two objects are on a supposed dimension, the more features they share and the fewer features there are to distinguish between them. So feature descriptions can clearly cover quantitative variations, and can do so in a manner that captures intuitions about quantitative similarity in terms of common and distinctive features (Tversky, 1977). (We note, though, that some implications of this extension of feature descriptions may prove incorrect; for example, since the descriptions of entities at the upper end of a continuum need to contain more features than those at the lower end, the former entities would be expected to be more complex.)

The case for the extendability of dimensions is more precarious. Note first that while dimensions are typically quantitative, they need not be continuous; discrete dimensions are still dimensions (Garner, 1978). To return to our weapon example, your dimension of potential damage may only include five discriminable values, but it is still quantitative as long as you can reliably judge one value to be greater than another. This raises the possibility of a dimension with only two values, where one value indicates some positive amount and the other, or zero value, indicates the absence of the dimension. Such a limiting case of a dimension captures something of what we mean by a feature, with the positive dimension value mapping onto the presence of a feature and the zero value mapping onto the absence of the feature. There are two problems with this extendability argument. First, a zero value on a dimension may not capture what people mean by the absence of a feature; thus it is one thing to say that the feature of "red" does not apply to the concept of truth, and quite another to say that truth has a zero value of redness. Second, our limiting case of a dimension may no longer constitute a "dimension" in any rigorous sense. Following Beals, Krantz, and Tversky (1968), it seems reasonable to insist that for something to be called a dimension it must have the property of *betweenness*—roughly that there exist some triple of values along the dimension such that one value appears (psychologically) to be between the others. Given this constraint, binary-valued contrasts are not dimensions, and true featural variations cannot be captured by dimensional descriptions. The upshot is that only a weak notion of dimensions can be fully extended to cover featural variations, and in this respect dimensional descriptions may be less powerful and more constrained than their featural counterparts.

Let us now return to differences between dimensions and features as they are typically used. In addition to the quantitative versus qualitative distinction, there is another difference in how dimensions and features are employed to represent objects. In describing a set of related objects in terms of dimensions, a person (1) typically uses a small number of dimensions and (2) assumes that every object has some value on each one. For example, a dimensional description of weapons might involve only three to five dimensions and assume that each weapon can be placed on each dimension. Feature descriptions, in contrast, are likely (1) to use many features and (2) to include features for some objects that are not applicable to others. One weapon, say a knife, might have dozens of features corresponding to the size, shape, and structure of its blade and handle, while another weapon, say a grenade, might be described in terms of fewer features that need not correspond to any of those used to describe the previous instance (what feature of a grenade

can possibly correspond to the pointedness of a knife's blade?). The feature approach again seems less constrained than the dimensional one.³

Regardless of whether dimensions or features are used to describe concepts, the goal is the same: to make apparent the relations between concepts. There are numerous ways this can be done. With the dimensional approach, each concept may be represented by a list of the values it takes on a set of dimensions, and the relation between any pair of concepts is determined by the proximity of their values along all dimensions; alternatively, the concepts may be represented as points in a multidimensional space and the relation between concepts given by the distance between the corresponding points. With the featural approach, concepts are often represented by feature lists, and the relation between concepts is given in terms of common and distinctive features. Figure 2 illustrates the latter approach. It contains (partial) featural descriptions of the concepts of robin, chicken, collie, and daisy. The features are of various types—animate, feathered, furry, and so on—but they are all potentially useful in describing relations between concepts. For example, Figure 2 makes it apparent why robin and chicken are more similar than robin and collie: robin and chicken have more common features and fewer distinctive ones than robin and collie, and the similarity of any two items should increase with the number of shared features and decrease with the number of distinctive ones (Tversky, 1977). An alternative to simple feature lists is to depict not only the features of a concept but relations between them as well—for example, to indicate explicitly that the features “animate” and “feathered” are correlated. When this is done many of the features are interconnected, and a network (rather than a list) of features results.

Though all the alternatives discussed above can be used to represent concepts and their interrelations, the feature list approach has

	<u>Robin</u>	<u>Chicken</u>	<u>Collie</u>	<u>Daisy</u>
F ₁	animate	animate	animate	inanimate
F ₂	feathered	feathered	furry	stem
F ₃	flies	pecks	brown-grey	white
	—	—	—	—

Figure 2 Some featural descriptions

been the most extensively used, particularly for classical-view representations of concepts. For this reason, the rest of this chapter focuses on features per se. Later, when we take up the probabilistic view, we will return to other kinds of descriptions.

Contents of Features

CONSTRAINTS ON FEATURES

If we are going to describe concepts in terms of features, we need some guidelines as to what can count as a feature. If any property, no matter how arbitrary or complex, could serve as a feature, we would soon be faced with all kinds of difficulties. For example, one could have a “concept” that consisted of the pseudofeatures “saw all of Tuesday Weld’s movies in one week” and “prefers Bloody Marys that contain two parts Worcestershire sauce for one part Tabasco.” Such a “concept” seems totally unnatural and points to the need to constrain features.

One constraint follows directly from the main purpose of features. If features make apparent relations between concepts, then a property is a useful feature to the extent that it reveals many relations between concepts. The feature of being male, for example, is useful because it brings out the relation between various classes of animate beings, like that between son and father, or between boy and colt. In contrast, neither of the pseudofeatures mentioned above would reveal many relations between concepts.

The above constraint can be strengthened. A set of features should not only make relations apparent but ideally should exhaust all potential relations between the concepts of interest. To illustrate, suppose someone postulated that young-male served as a single feature in describing concepts of animate beings. Although this would make apparent the relation between boy and colt (both contain the feature of young-male), it would not exhaust all potential relations of interest; for example, it would not show the relation between boy and girl. To capture the latter relation we obviously need to decompose young-male into the component features of (1) young (in the sense of a quality, not a quantity) and (2) male, where feature 1 can now be used to show the relation between boy and girl. To ask for features that exhaust all possible relations is to ask for features that are not themselves decomposable, that is, for primitives or all-or-none features. This is a difficult constraint to meet (we never really invoke it), but it is surely powerful since it unequivocally excludes a pseudofeature like “saw all of Tuesday Weld’s movies in one week.”⁴

As a second constraint we seek features with some generality, in

the sense that a feature should apply to many concepts within a domain rather than to a few. (Of course, we would not want the feature to apply to every concept in the domain, for then it would have no discriminative value.) A kind of corollary of this generality constraint is that the number of features needed to describe a conceptual domain should be small relative to the number of concepts in that domain; that is, all other things being equal, the greater the applicability of a set of features, the fewer the features needed to characterize a domain. Although this constraint of generality rests on rough intuitions, it again seems sufficient to rule out properties like "saw all of Tuesday Weld's movies in one week." (Note that even though we seek features that are as general as possible, they still may be more numerous than the dimensions needed to describe the same conceptual domain.)

The two constraints discussed above deal with structural aspects of concepts. Consider now a processing constraint: the features posited should serve as the inputs for categorization processes. Thus, for the feature of male to be a useful one, it must not only reveal relations between concepts in an economical fashion (our two structural constraints), but it must also be used by people in reaching decisions about categorization. Though this constraint may carry little force with those interested in a purely structural approach to concepts (many linguists and philosophers), it is essential for psychologists interested in the processing of concepts. Indeed, it must take precedence over other constraints, as we will readily accept a nonprimitive and nongeneral property as a feature if there is convincing evidence that it is used in categorization.

It is instructive to illustrate these three constraints in relation to recent work on the perception of letters. This work essentially treats each letter of the alphabet as a separate concept and asks what features should be posited to define the concept. Introspection as well as early physiological evidence (for example, Hubel and Wiesel, 1962) suggested that the features may be line segments, angles, and curves. But exactly which lines, angles, and curves are posited as features depends on the three constraints. Just about everyone would agree that the features must account for why, under limited viewing conditions, a particular letter (for example, *P*) is confused with some letters (*B*, *R*) but not with others (*X*, *V*; see Gibson, 1969; Rumelhart and Siple, 1974; Townsend and Ashby, 1976). Since the probability of confusing a particular pair of letters is assumed to be a measure of the similarity of that pair, accounting for patterns of confusions amounts to making similarity relations apparent, which is our first constraint. Furthermore, all serious proposals in this domain posit features that apply to many letters —

a vertical line is a feature of *B*, *D*, *E*, *F*, and so on — and this is in keeping with our second constraint. Finally, researchers have tried to show that these features are the inputs to recognition processes by demonstrating, for example, that an estimate of the exact features extracted from a letter can be used to predict the accuracy, confidence, and speed with which a letter can be categorized (see Rumelhart and Siple, 1974; Townsend and Ashby, 1976). Failure to meet this third constraint would usually make researchers discard their featural definitions.

This example can also be used to address a fundamental issue about features, one that arises whenever any component properties are posited. The issue, in a nutshell, is that there are no a priori means for distinguishing between a concept and a feature, for any feature can itself be treated as a concept. In the letter perception work, for example, the concept for the letter *E* might include the feature of a vertical line, but a vertical line is also a possible concept, so what reasons are there for calling it a "feature" and reserving the term "concept" for *E*? Or to put it another way, do not all the problems we face in explaining how *E* is represented, accessed, and processed also apply to how a vertical line is represented, accessed, and processed? If so, explaining *E*s in terms of vertical lines is simply begging the question, and the business of seeking constraints on features is beside the point.

There is a way out of this conceptual morass, we think, that is nicely demonstrated by the research on letter perception. In essence the solution is this: true, there may be no conceptual reasons for distinguishing between concepts and features, but there may well be empirical reasons for making this distinction. Thus, when it comes to explaining results in letter perception, there are many empirical reasons for assuming that a vertical line is a feature while the letter *E* is not. To restate a previously mentioned finding, treating a vertical line as a feature helps to explain why some letter pairs are more likely than others to be confused under limited viewing conditions; for example, *E* and *F* are more likely to be confused than *E* and *V* because the former pair shares a vertical line. No such account of confusions is forthcoming if one treats *E* and *F* as features here. This is just one example among many of how treating components such as vertical lines as features has led to an account of empirical findings about letter perception that might have remained a mystery had researchers not reduced the letters to their component properties.

These empirical successes in distinguishing visual features from visual concepts provides a rationale for thinking that the same kind of distinction can be supported with semantic concepts and their

features. Indeed, we think there is already some evidence that describing concepts in terms of their component features will lead to accounts of empirical phenomena that would not be forthcoming otherwise.

One important caveat must be added to the foregoing discussion. If empirical results determine whether we should assign an entity the role of feature or that of concept, then a major change in the set of results to be accounted for can lead to a change in role assignment. To illustrate, if we change the empirical results of interest from findings on letter perception to findings on word perception, we may want to change the status of *E* from concept to feature. We may, for example, need to assume that *E* is a feature of word concepts in order to explain why certain words, such as *bed* and *ted*, are likely to be confused under limited viewing conditions. When such a change in role assignments occurs, it may indicate we are dealing with a system that has several levels, so that what is a concept at level n is best thought of as a feature at level $n + 1$.

PERCEPTIBILITY OF FEATURES

Perceptual versus Abstract Features

Even with the constraints mentioned above, the features of object concepts can vary a great deal. One particular source of variation, the perceptibility of features, will be of great concern to us. At one extreme are true perceptual features—those corresponding to outputs from the perceptual system; examples might include the presence of a line or the presence of some degree of curvature. At the other extreme are abstract features that have minimal connection to perceptual experience; as an example, a feature of the concept “clothing” might be “manufactured with the intent of human usage.” And then there are a myriad of in-between cases, including the familiar features used with animal concepts like “animate” and “male,” and functional features like “used for transportation.”⁵

This perceptibility distinction raises a number of issues. The first is whether the features of object concepts should be restricted to more perceptual features, that is, whether perceptibility should be another constraint on features. This constraint has frequently been tried, particularly in conjunction with the classical view, for there is a long tradition in psychology to define stimulus objects in physical terms. This move always seems to end in failure, however. The most recent effort, and it is a valiant one, can be found in the work of Miller and Johnson-Laird (1976). After systematically reviewing what is known about the possible outputs of the human perceptual system, they attempt to use only these outputs in their feature de-

scription of certain object concepts (such as “table”). They convincingly argue that this enterprise is doomed.

Miller and Johnson-Laird first try to represent the concept of table by perceptual features like the following: (1) is connected and rigid, (2) has a flat and horizontal top, and (3) has vertical legs. They then spell out two serious drawbacks to this approach. First, some of these features are not true of all things we call tables (for example, many drafting tables do not have a horizontal top); this seems to be the case no matter how one alters the list of perceptual features. Second, if concept descriptions are confined to perceptual features, how can one ever explain that the same object may sometimes be treated as an instance of the concept “kitchen table” and other times as an instance of the concept “work table”? The solution to these problems, according to Miller and Johnson-Laird, is to include functional features in concept descriptions. For example, in the concept “table” if we replace the perceptual features “has a flat and horizontal top” with the functional feature “has a top capable of support,” we can see how drafting tables can be included in the general concept of tables. And once we accept functional features, there is no problem in explaining how the same object can sometimes be subcategorized a kitchen table and sometimes a work table; the more specific concepts of kitchen table and work table have different functional features, and sometimes use of the object will match one of these functions, while other times its use will match the other function. Given arguments like this, we take it as a starting point that the features of an object concept can contain abstract and functional features as well as perceptual ones.

But now we have a new problem. How can one decide that a visually presented object is an instance of a particular concept, since features of the object must be true perceptual ones whereas some features of the concept may be abstract? The only solution is that one must have some knowledge that is capable of mediating between the features at the two levels; that is, to determine whether an abstract feature is perceptually *instantiated* in an object, one must have recourse to ancillary knowledge about the relation between abstract and perceptual features.

Let us illustrate with the concept of boy. Following traditional linguistic wisdom, we assume that people may represent this concept by three abstract features: human, male, and young (relative to some criterion). How can one use these features to determine that some particular person (a clothed one) is in fact a boy? To establish that the person is male, one would presumably check things like body proportions, dress, and hairstyle, which are more perceptual features. (Even these features are more abstract than

outputs of the perceptual system, but we will ignore this complication in what follows.) To establish that the person is young, one would use perceptual features like height and weight. Hence, both male and young are instantiated by more perceptual features, and categorization therefore requires ancillary knowledge about which perceptual features instantiate which abstract ones.

The Core versus Identification Procedure of a Concept

In view of the foregoing discussion, let us distinguish between the *core* of the concept and the ancillary knowledge, which we will call the *identification procedure*. (A similar distinction is used by Miller and Johnson-Laird, 1976, and by Woods, 1981). The core of the concept "boy" would include the abstract features of human, male, and young, while the identification procedure would include perceptual features like height and weight as well as an indication that they instantiate the feature of young. The features of the core are primarily responsible for revealing certain relations between concepts, like that between boy and girl (both young and human) or that between boy and colt (both young and male), while the features of the identification procedure are used for categorizing real-world objects. To the extent that the features of a core are perceptual, there is no need for a separate identification procedure; for a geometric concept like "square," the core contains perceptual features like "four-sided" which can be used for categorizing real-world objects, thereby making an identification procedure superfluous. Perceptual features in the core, then, shrink the requisite identification procedure. There is also another way in which the core takes psychological precedence over the identification procedure: abstract features in the core determine the nature of the perceptual features in the identification procedure. To return to the boy example, being male (a core feature) is partly a matter of hormones, and these hormones produce secondary sex characteristics like body proportions (a feature in the identification procedure). From this perspective, if one wants to understand the nature of concepts, it is best to go for the core.⁶

The distinction discussed above is clearly related to Frege's classic distinction (1892) between *sense* and *reference*: the sense of a concept is given by its relation to other concepts, while the reference of a concept is given by its relation to objects and events in the world. Thus Frege's sense corresponds to our core, and Frege's reference to our identification procedure. Just as we have assumed that the core dictates the contents of the identification procedure, so Frege postulated that the sense of a concept determines its reference. We prefer our terms to Frege's because his distinction was intended as a

philosophical rather than a psychological one, and hence it may be wrong to identify his notion of reference with the features people actually use in categorizing real-world objects.

Consider now how we can use our core versus identification distinction in evaluating experimental studies of concepts. As mentioned earlier, in some experiments subjects are asked to decide whether statements about subset relations (for example, "A raisin is a fruit") are true or false. In this kind of *semantic* task, subjects can answer just by consulting the cores of the relevant concepts (do the core features of the concept fruit match those of the concept raisin?). They might also consider the identification procedures of the two concepts, but there is no necessity to do so. In other experiments, however, subjects are asked to decide whether a pictured object is an instance of a target concept (for example, "Is this specific object a fruit?"). In such a *perceptual* task, subjects must consult the identification procedure of the target concept unless the core consists entirely of perceptual features. Since the views of concepts we are interested in, particularly the classical view, are mainly concerned with concept cores, we will emphasize experimental studies that either use a semantic task or use a perceptual one where there is good reason to believe the core consists entirely of perceptual features. Later, at the end of our discussion of the classical view, we will have occasion to question the wisdom of our decision.

3 | The Classical View

THE CLASSICAL VIEW is a psychological theory about how concepts are represented in humans and other species. In philosophy, the origins of this view go back to Aristotle, while in experimental psychology the view can be traced to Hull's 1920 monograph on concept attainment. In assembling our rendition of the classical view, however, we have relied mainly on contemporary sources. These sources include philosophically oriented studies of language (for example, Katz, 1972, 1977; Fodor, 1975); linguistic studies (Lyons, 1968; Bierwisch, 1970; Bolinger, 1975); psycholinguistics (Fodor, Bever, and Garrett, 1974; Miller and Johnson-Laird, 1976; Anglin, 1977; Clark and Clark, 1977); and psychological studies of concept attainment (Bruner, Goodnow, and Austin, 1956; Bourne, 1966; Hunt, Marin, and Stone, 1966).

While we think we have captured some common assumptions in these various sources, we are less sure that we have been faithful to the spirit of these works. For instance, the psychological studies of Bruner, Goodnow, and Austin (1956) were more concerned with the strategies people use in determining the relevant features of concepts than with supporting the classical view. Indeed, these authors even devoted one chapter of their influential book to concepts structured according to the probabilistic view. Still, the bulk of their effort employed artificial concepts structured according to the classical view, and there is no guarantee that the strategies that Bruner, Goodnow, and Austin turned up will be easily extendable to other views of concepts. Similar caveats apply to many of the other sources.

In terms of distinctions drawn earlier, we will be concerned here exclusively with feature descriptions, since all the sources given above (as well as many not listed) have analyzed concepts in terms of features. Also, it seems that practitioners of the classical view

have been primarily interested in characterizing the core of concepts, not their identification procedures, and as mentioned earlier, our treatment of the view will focus on the core. Finally, a word about the role of process models is in order. The classical view is a proposal about representations, not about processes. Once we have described the representational assumptions that make up the classical view and the criticisms of these assumptions, we will give some consideration to process models that can be generated from the view.

Representational Assumptions

SUMMARY REPRESENTATIONS

The first assumption is as follows: The representation of a concept is a summary description of an entire class, rather than a set of descriptions of various subsets or exemplars of that class. To illustrate, in representing the concept of bird we would not list separate descriptions for different species (like robin and chicken) or for specific instances (like our pet canary Fluffy), but rather would give a summary representation for all birds. As Rosch (1978) has emphasized, condensing a concept into a single summary greatly reduces the amount of information we need to store.

This notion of a summary representation is sufficiently important that it is worth specifying some explicit criteria for it. A summary representation, then, (1) is often the result of an abstraction process, (2) need not correspond to a possible specific instance, and (3) applies to all possible test instances. Thus: (1) one's summary representation for fruit is often based on induction from specific instances (as well as on facts one has been told about fruits in general); (2) the representation might contain fewer features than would be found in the representation of any possible instance; and (3) whenever one is asked whether or not a test item designates an instance or subset of fruit, the same summary representation is always retrieved and examined.

NECESSARY AND SUFFICIENT FEATURES

The heart of the classical view is contained in its second assumption: The features that represent a concept are (1) singly necessary and (2) jointly sufficient to define that concept. For a feature to be singly necessary, every instance of the concept must have it; for a set of features to be jointly sufficient, every entity having that set must be an instance of the concept. It is convenient to illustrate with a geometric concept—squares again. Recall that the concept of square may be represented by some in terms of the following

features: closed figure, four sides, sides equal in length, and equal angles. Being a closed figure is a necessary condition, since any square must have this feature; the same is true of the features of having four sides, the sides being equal, and the angles being equal; and these four features are jointly sufficient, since any entity that is a closed figure, has four sides equal in length, and has equal angles must be a square. We will sometimes refer to such necessary and sufficient features as *defining* ones.

Many scholars who have written about the classical view have emphasized that this assumption is about necessity or essentialism, not probability (see Cassirer, 1923; Katz, 1972). It is not just that all squares happen to have four sides, but rather that having four sides is essential to being a square. Or take another example: the defining features of bachelor—male and unmarried—are not only true of all bachelors (which is merely a statement about conditional probabilities), but are essential conditions for being a bachelor. To appreciate this distinction between probability and essentialism, suppose that the feature of “not wearing wedding bands” is also true of all bachelors. While this feature has the same conditional probability as being unmarried, only the latter would be essential.

It is important to note that this assumption about defining features implies that natural concepts are never disjunctive. To illustrate, let us consider first a *totally disjunctive* concept, which says that an instance either has features F_1, F_2, F_3, F_4 or features F'_1, F'_2, F'_3, F'_4 ; that is, two instances need have no features in common. This means there are no necessary features, which violates the classical view's assumption about defining features. Now consider a *partially disjunctive* concept, which says that an instance either has features F_1, F_2, F_3, F_4 or F_1, F_2, F_3, F'_4 ; that is, any two instances must have some features in common (F_1 - F_3), but other features may differ (F_4 versus F'_4). This means there is no set of necessary features that are jointly sufficient: F_1 - F_3 are necessary but not jointly sufficient, while either F_1 - F_4 or F_1 - F'_4 are sufficient but include at least one nonnecessary feature (F_4 or F'_4). This too violates the assumption about defining features.

NESTING OF FEATURES IN SUBSET RELATIONS

The final representational assumption of interest is as follows: If concept X is a subset of concept Y, the defining features of Y are nested in those of X. It is again convenient to illustrate the assumption with geometric concepts. Suppose that people represent the concept of quadrilateral by two features: closed figure and four-sided. These two features are the ones we have included in our previous example of the concept square, along with the features of

equal sides and equal angles. Hence a square is a subset of quadrilateral, and the defining features of quadrilateral are nested in those of square. Similarly, the defining features of bird (for example, animate and feathered) are nested in those of robin, since robins are a subset of birds. Of course the more specific concept—square or robin—must also include some defining features that are not shared by its superset; for example, robin must contain some features that distinguish it from other birds. This guarantees that the representation of a concept cannot be a realizable instance, since the concept must contain fewer features than any of its instances.

Although this nesting assumption is a common one among advocates of the classical view, some would not buy it wholesale. Fodor (1975) in particular questions the assumption, and suggests instead that related concepts may be defined by different sets of features. For example, the feature of bird that specifies “feathered” may not be identical to any specific feature of chicken. If we accept this possibility, the classical view is considerably weakened in the claims it makes about concepts. Given this, for the time being we opt for the version of the view that includes assumption 3.

SUMMARY

The three assumptions of the classical view are summarized in Table 1. Although they are not the only assumptions used by proponents of the classical view, they are the modal ones. Indeed, they are presupposed by most of the significant psychological work done on artificial concepts from 1920 to 1970 (for reviews, see Bruner, Goodnow, and Austin, 1956; Bourne, 1966; Bourne, Dominowski, and Loftus, 1979).

TABLE 1 THREE ASSUMPTIONS OF THE CLASSICAL VIEW

- | |
|--|
| 1. Summary representation |
| 2. Necessary and sufficient (defining) features |
| 3. Nesting of concept's defining features in its subsets |

One last point: the three assumptions say nothing about possible relations between features—that is, the features are treated as if they were independent. This treatment may be adequate for certain semantic domains, called *paradigms*; an example would be kinship concepts, like mother, father, son, and daughter. Here the features—sex and age—seem to combine as independent entities. However, the idea of independent features does not fit with other

semantic domains, called *taxonomies*; an example would be animal concepts, like robin, bird, animal, and organism. Here the features are clearly related; for example, the feature "living" (defining for organism) is implied by the feature "animate" (defining for animal). Though such relations are not mentioned in our three assumptions, we do not mean to exclude them from the classical view. Rather, we are trying to keep the assumptions down to a minimum, agreed-upon set. Even this small set will soon be shown to contain a great deal of debatable content.

General Criticisms of the Classical View

Throughout the years there have been various general criticisms of the assumptions of the classical view. In what follows we consider four criticisms that seem especially widespread, along with possible rebuttals.

FUNCTIONAL FEATURES

Some have argued as follows:

1. The classical view deals only with structural features—fixed properties (of varying perceptibility) that describe an entity in isolation, like the handle or concavity of a cup—and prohibits functional features, like the fact that a cup is used to hold something.
2. But for many concepts, particularly those corresponding to human artifacts like cups and chairs, the defining features are functional ones.
3. Therefore, the classical view cannot handle all concepts.

Cassirer (1923) put forth this argument some time ago, and Nelson (1974) and Anglin (1977) have recently reiterated it and suggested that it is devastating to the classical view.

Given our earlier discussion about the need to consider abstract, functional features in concept cores, it should come as no surprise that we think this argument is based on a faulty premise, namely premise 1. Nothing in our three assumptions excludes functional features. A functional feature, such as the fact that a cup can hold liquid, can be used in a summary description of an entire class, can be singly necessary and part of a jointly sufficient set, and can be nested in other feature sets. Furthermore, none of our constraints on features is inconsistent with functional features. A feature like "holdability" can bring out relations between concepts (for example, between cup and bowl), can apply to many different classes, and can be used as an input to categorization processes. In short,

our rendition of the classical view is as hospitable to functional features as it is to structural ones.

Why, then, do so many psychologists think the classical view should be restricted to structural features? No doubt because such features are generally more perceptual than functional features. But then why do psychologists think the classical view should be restricted to perceptual features? We discussed one answer to this in the previous section—perceptual features greatly simplify the analysis of how people categorize perceptual objects. Another reason perceptual features have proved so attractive to psychologists is that such features are very easy to manipulate in experimental studies of concept attainment and utilization.

When Hull started his experimental study of the classical view in 1920, he used novel visual forms that were composed of multiple features. This allowed him to control precisely which features occurred in all instances of a concept, that is, which features were necessary. Had he used more abstract features, like functional ones, Hull would have had either to give his subjects real manipulable objects and let them discover the function (a messy task at best), or to give them pictures of objects that instantiated the function to varying degrees (which again is a relatively uncontrolled paradigm, though very likely a more ecologically valid one). Hull's emphasis on easily manipulable perceptual features proved so attractive that more than two generations of experimental psychologists have bought it, thereby making it seem that perceptual features are part and parcel of the classical view. It is only during the last fifteen years, with the influence of nonexperimental disciplines like linguistics and philosophy on psychology, that psychologists have begun to realize that some concepts may have functional features at their core.

DISJUNCTIVE CONCEPTS

A more powerful argument against the classical view is the following:

1. The classical view excludes disjunctive concepts.
2. But many concepts are clearly disjunctive, like that of a strike in baseball (which can be either a *called* or a *swinging* strike).
3. Therefore, the classical view cannot handle all concepts.

Certainly we agree with premise 1, for we noted earlier that the assumption of defining features excludes disjunctive concepts. Premise 2, though, is debatable. Specifically, how widespread are

disjunctive concepts? Unfortunately, there is nothing resembling a clear-cut answer to this question. If we rely on intuitions (our own and those published by semanticists) and restrict ourselves to concepts about naturally occurring objects (flora and fauna), we can think of no obvious disjunctive concepts. Disjunctive concepts, then, may be rare, restricted to man-made concoctions (like a baseball strike), and constitute special cases that should not obscure the general conjunctive nature of concepts.

This reasoning may be too facile, however. There are alternatives to intuitive analyses of concepts, and at least one of these suggests that disjunctive concepts may be quite widespread. Rosch and her colleagues (1976) asked people to list the features of concepts, where the concepts varied in their level of inclusiveness (for example, kitchen chair, chair, and furniture). Their data suggest that the more inclusive or superordinate concepts may be disjunctive. For superordinate concepts like animal, plant, vehicle, furniture, clothing, and tool, people list few if any features; for concepts that are one level less inclusive, like bird, flower, truck, chair, hat, and hammer (what Rosch and colleagues call the *basic level*), people list a substantial number of features. This finding suggests that superordinate concepts are often disjunctive (and that basic-level concepts are the most inclusive level at which conjunctive concepts appear).

There is, however, an alternative interpretation of the data compiled by Rosch and her associates, one that can save the classical view from a plethora of disjunctive concepts. The features that people listed may well have been part of the identification procedure, not the core. But then why should identification procedures be disjunctive only for superordinate concepts? The reason is very likely that the cores of superordinate concepts contain abstract features (remember "intended to be worn by a human"), and such features can only be instantiated disjunctively at the perceptual level. Under this interpretation, concept cores are as conjunctive as the classical view claims they are, and those who mistakenly think otherwise have confused the identification procedure with the core. To illustrate further the flavor of this argument, let us consider the concept "extreme." Some might deem it disjunctive because it implies one pole or the other, but this may be an aspect of the identification procedure, not the core, where the latter may mean "a value far from the central tendency." Another example is the concept of split personality: doesn't this mean personality X or personality Y, but not both (an exclusive disjunction)? Perhaps it does at the level of an identification procedure, but the concept core may mean "manifests different personalities," which is not inherently disjunc-

tive. (Anisfeld, 1968, has made a similar argument using the notions of sense and reference.)

The upshot is that we have no firm evidence, intuitive or otherwise, about the prevalence of disjunctive concepts. Without such evidence, it is difficult to say how damaging the disjunctive-concepts argument is to the classical view.

UNCLEAR CASES

A third argument against the classical view (see, for example, Hampton, 1979) takes the following form:

1. The classical view assumes that if concept X is a subset of concept Y, the defining features of Y are nested in those of X.
2. Given this, judgments about whether one concept is a subset of another should be clear-cut, since one merely has to compare defining features.
3. In fact, it is often unclear whether one concept is a subset of another. People disagree with one another about a particular subset relation, and the same person may even change his mind when asked the same question on different occasions (see McCloskey and Glucksberg, 1978). The classical view has no way of accounting for such unclear cases.

The weak part of this argument is premise 2, since a nesting of defining features does not guarantee that judgments about subset relations will be clear-cut. We can think of at least two reasons why this is so, and it is best to illustrate them by a specific example.

When asked, "Is a tomato a fruit?" many people, even college-educated ones, are unsure of whether this particular subset relation holds. One simple reason they may be uncertain is that their concepts of tomato and fruit may be faulty or incomplete—that is, they are missing some defining features of fruit and consequently cannot tell whether or not a tomato is a fruit. To put it more generally, the classical view does not stipulate that every adult has mastered every familiar concept; rather, it allows for the possibility that many of us are walking around with incomplete concepts, just as long as whatever features we do have are at least necessary ones. (Such incomplete concepts could not be *too* incomplete, however, since adults obviously do a good job of using their concepts in dealing with their environment.) A second way to reconcile the classical view with unclear cases is to assume that some concepts have two definitions, a common and a technical one (Glass and Holyoak, 1975). Thus one might be unsure about what concept a tomato belongs to because a tomato meets the technical definition of a fruit

(for example, it has seeds) but the common definition of a vegetable (it plays a particular role in meals).

SPECIFYING THE DEFINING FEATURES OF CONCEPTS

Of all the arguments against the classical view, the best-known one goes as follows:

1. The heart of the classical view is its assumption that every concept has a set of necessary and sufficient features.
2. Decades of analysis have failed to turn up the defining features of many concepts.
3. Therefore, many concepts simply do not have defining features.

It was essentially this argument that Wittgenstein (1953) pursued in his well-known critique of a classical-view approach to natural concepts. One of Wittgenstein's most famous examples was that of the concept of games, and we can use it to illustrate the flavor of his argument. What is a necessary feature of the concept of games? It cannot be competition between teams, or even the stipulation that there must be at least two individuals involved, for *solitaire* is a game that has neither feature. Similarly, a game cannot be defined as something that must have a winner, for the child's game of *ring-around-a-rosy* has no such feature. Or let us try a more abstract feature—say that anything is a game if it provides amusement or diversion. Football is clearly a game, but it is doubtful that professional football players consider their Sunday endeavors as amusing or diverting. And even if they do, and if amusement is a necessary feature of a game, that alone cannot be sufficient, for whistling can also be an amusement and no one would consider it a game. This is the kind of analysis that led Wittgenstein to his disillusionment with the classical view.

Although this argument clearly has merit, it is by no means ironclad, for its conclusion—that many concepts do not have defining features—is based on a lack of progress by the classical view. When Wittgenstein—or anyone else—asserts: "There are no defining features of concept X," it is equivalent to asserting: "No one has yet determined the defining features of concept X," since both assertions would be refuted by a cogent proposal of these features. Moreover, one could claim that part of the reason progress has been so slow is that we have been looking for the wrong kind of defining features—perceptual ones that are likely to be part of an identification procedure—when we should have been seeking abstract, relational, or functional features that may well make up the core of many concepts.

Thus the Wittgenstein argument is nothing like a principled disproof of the classical view; it is instead an empirical argument about the observed rate of progress of a theoretical approach to concepts. Once this is appreciated, one can acknowledge that the argument certainly has force (like any excellent empirical argument) but that it deals no death blow to the classical view.¹

A NOTE ON SCIENTIFIC CONCEPTS

It is worth pointing out that the last two criticisms of the classical view of psychological concepts—unclear cases and failure to specify defining features—have also been raised as criticisms of the classical view when it is used as a metatheory of scientific concepts. That is, in addition to its use as a psychological theory, the classical view has also served as a metatheoretical prescription of what scientific concepts should look like, and here it has run into problems similar to those we just described.

There are numerous unclear cases for classically defined biological concepts. For example, there is no uniform agreement among biologists as to whether *Euglena*, a mobile organism that manufactures chlorophyll, should be classified as an animal or a plant. Cases like this are occurring with sufficient frequency to lead scientists to question the validity of the classical view for biological classification (see Sokal, 1974).

Similarly, there has been substantial difficulty in specifying the defining features of biological species, at least in terms of structural features (Sokal, 1974; Simpson, 1961). Toward the end of the eighteenth century Linnaeus proposed that any biological species can be characterized by three kinds of features: (1) features that comprise the *essence* of the species, which are features that every member of the species must have and that correspond to what we have called defining features; (2) features called *properties*, which are common to all members of the species but are not part of the essence; and (3) features called *accidents*, which characterize some but not all members of a species. According to Linnaeus, only features comprising the essence should be used in classification. This classical-view approach has had great influence, but it now seems problematic as a guide to biological classification. For one thing, taxonomists have generally been unable to distinguish features comprising the essence from those called properties. For another, taxonomists have found that the so-called accidents, features not true of every species member, are sometimes genetically based and important for understanding and defining the species.

These developments in biological classification are relevant to a psychology of concepts. Recall that in the Introduction we noted that there was little hope for classically defined mental representa-

tions if there was little evidence that such concepts could be given a classical definition in some language. The most likely place to look for classical definitions of flora and fauna is the language of biology, and to the extent that the classical view fails here, it will likely fail as a psychological theory as well.

SUMMARY

Table 2 summarizes the four general criticisms we have discussed. How badly do they damage the classical view? In answering this, we must distinguish between an *in principle* criticism—one that shows that the view could never handle a particular problem—and an *empirical* criticism—one that shows that specific embodiments of the view have thus far failed to handle a particular problem. All four criticisms seem to be mainly empirical ones.

TABLE 2 FOUR GENERAL CRITICISMS OF THE CLASSICAL VIEW

- | |
|---|
| 1. Exclusion of functional features |
| 2. Existence of disjunctive concepts |
| 3. Existence of unclear cases |
| 4. Failure to specify defining features |

The first criticism—that the classical view excludes functional features—is clearly about typical applications of the view in psychology, and not about what this view can accomplish in principle. We showed that the assumptions of the classical view are as compatible with functional features as they are with structural ones. How successfully one can use the view with functional features, however, remains an open question. The work of Miller and Johnson-Laird (1976) at least suggests that one can develop a classical-view model of categorization that employs functional features.

The second criticism—that the view excludes disjunctive concepts—comes closest to offering an in-principle argument against the classical view. If some natural concepts are clearly shown to be disjunctive, they simply fall outside the domain of the classical view. Such convincing demonstrations, though, have been rare.

With regard to the third criticism, we argued that unclear cases are not necessarily inconsistent with a classical-view description of natural concepts because people may have incomplete, or multiple, definitions of a concept. Again, the criticism is hardly a proof against the classical view.

Finally, there is the criticism that the classical view has made lit-

tle progress in specifying defining features. In discussing this criticism, we emphasized its empirical nature—it is a statement about what has happened so far, not about what can happen. Still, as an empirical criticism, it is one of the strongest arguments against the classical view.

We emphasize the empirical nature of these criticisms because we wish to dispel the popular notion that the classical view has been proved wrong by a priori arguments and consequently that no empirical work is needed. A more correct reading of the situation is this: serious empirical criticisms have been raised against the classical view of natural concepts—serious enough to make us have grave reservations about the view, but not serious enough to say that the view should be discarded at this point.

Experimental Criticisms of the Classical View

Though the general criticisms discussed above are telling, they are not the only reasons why psychologists are currently forsaking the classical view in droves. There are other criticisms of this view that stem from experimental findings about how people use natural concepts, such as how they decide that apples are fruit. Before delving into these findings, we would like to interject a cautionary note. Since the findings deal with how people use concepts, they reflect categorization processes as well as concept representations. This means that we cannot go directly from the findings to claims about how concepts are represented; instead, we must interpret these findings in terms of both representations and processes—in short, in terms of models. This point has been missed in a good deal of recent research on natural concepts, where it has often been assumed that categorization data directly inform us about the nature of concepts. The best way to document the need for a model in interpreting categorization effects is to consider some results of interest and then show that their implications for the classical view depend on the specific model used to instantiate this view. This is the procedure we adopt in the following discussion.

SIMPLE TYPICALITY EFFECTS

Experimental Results

Of all the experimental findings used as evidence against the classical view, perhaps the best known are the effects of *typicality* (also called *prototypicality*). The most critical result is that items judged to be typical members of a concept can be categorized more efficiently than items judged to be less typical. The details of this result are as follows: People find it a natural task to rate the various

TABLE 3. TYPICALITY RATINGS FOR BIRD AND MAMMAL INSTANCES

Instance	Rating	Instance	Rating
Robin	3.00 ^a	Deer	2.83
Sparrow	3.00	Horse	2.76
Bluejay	2.92	Goat	2.75
Parakeet	2.83	Cat	2.67
Pigeon	2.83	Dog	2.67
Eagle	2.75	Lion	2.67
Cardinal	2.67	Cow	2.58
Hawk	2.67	Bear	2.58
Parrot	2.58	Rabbit	2.58
Chicken	2.00	Sheep	2.58
Duck	2.00	Mouse	2.25
Goose	2.00	Pig	2.17

Source: After Rips, Shoben, and Smith (1973).

a. Higher numbers indicate greater typicality.

subsets or members of a concept with respect to how typical or representative each one is of a concept.² Such ratings were first reported by Rips, Shoben, and Smith (1973) and by Rosch (1973). Table 3 presents the ratings of Rips, Shoben, and Smith for the concepts of birds and mammals. As can be seen, robin and sparrow are considered typical birds, hawk and eagle less typical, and chicken and penguin atypical. Ratings like these have now been obtained for many noun categories in English, and have been shown to be highly reliable across raters (Rosch, 1973) and to be relatively uncorrelated with frequency or familiarity (Mervis, Catlin, and Rosch, 1976).

What is most important about these ratings is that they predict how efficiently people can categorize the various members of a concept in a semantic categorization task. One variant of this task is illustrated in Figure 3. On each trial, the subject is given the name of a target concept, like bird, followed by a test item; the subject decides whether the test item names a subset or member of the target concept, like robin or chicken, or a nonmember, like dog or pencil. The main data of interest are the times for correct categorizations. When the test item, or probe, in fact names a member of the target concept, categorization times decrease with the typicality of the probe (see Rips, Shoben, and Smith, 1973; Rosch, 1973). For example, when bird is the target concept, test items corresponding to robin and sparrow are categorized more quickly than those corresponding to eagle and hawk, which in turn are categorized faster

	Target Concept	Test Item	Correct Response
Trial n	Bird	Robin	Yes
Trial $n + 1$	Fruit	Cup	No
Trial $n + 2$	Bird	Chicken	Yes

Figure 3 One variant of semantic categorization task (other variants require verification of sentences of form "A robin is a bird" or "All apples are fruit")

than the probes of chicken and goose. Furthermore, to the extent that there is any variation in the accuracy of these categorizations, error rates also decrease with the typicality of the probe (see Rips, Shoben, and Smith, 1973). These effects are extremely reliable: they have been documented in more than 25 experiments that have used many different variants of the semantic categorization task (see Smith, 1978, for a partial review).

Though most studies of typicality have been concerned with categorization times, Rosch and Mervis have demonstrated a host of other typicality effects. For instance, the typical members of a concept are the first ones learned by children, as judged by either a semantic categorization task (Rosch, 1973) or by how accurately children can sort pictured objects into taxonomic categories (Mervis, 1980). (The latter finding is not strictly relevant to issues about concept cores.) Further, the typical members of a concept are likely to be named first when subjects are asked to produce all members of a category (Mervis, Catlin, and Rosch, 1976). And typical members are also more likely to serve as cognitive reference points than are atypical members (Rosch, 1975); for example, people are more likely to say "An ellipse is almost a circle" (where *circle*, the more typical form, occurs in the reference position of the sentence) than "A circle is almost an ellipse" (where *ellipse*, the less typical form, occurs in the reference position). This list of effects could be extended (see, for example, Rosch, 1974, 1975, 1978), but it is adequate for our purposes.

What does all this have to do with the classical view? Simply this: typicality effects reveal that not all members of a concept are equal, or to put it more positively, that concepts possess an internal structure that favors typical members over less typical ones. The representational assumptions of the classical view, however, sug-

gest that all members of a concept are equal, since all members of concept X must have the defining features of X. At first glance, then, typicality effects seem incompatible with the classical view, a conclusion that has been drawn many times. A more thorough analysis is needed here, however—one that considers processes as well as representations.

A Classical-View Model for Simple Typicality Effects

Since we need to interpret categorization results in terms of process models, the real question for the classical view is this: Can a model based on this view (that is, one that incorporates its three representational assumptions) account for the typicality effects we have described? The answer is clearly yes, and we will describe such a model in this section. We caution the reader, however, that the model presented will run into problems when later confronted with other findings; our reason for describing the model here is to demonstrate how easy it is to come up with a classical-view model that can account for effects frequently claimed to be inconsistent with this view.

We call our proposal the *complexity model*. It assumes that concepts are represented just as the classical view says they are, and that in a categorization task these representations are processed by two sequentially ordered stages, the *access* and *comparison* stages. The stages operate as follows:

1. When given the target and probe concepts, the subject starts accessing the defining features of both concepts, with access order being random over trials.
2. As soon as any defining features are available, the subject compares those of the target concept to those of the probe. The subject responds affirmatively ("Yes, it's a member") only when every feature of the target has matched a feature of the probe, but can respond negatively ("No, it's not a member") as soon as any feature of the target mismatches a probe feature. This stage is limited in capacity, and therefore the time needed to compare probe and target concepts increases with the number of features in either concept.

Only one more assumption is needed, and it is the critical one: Typical members of a concept have fewer features than atypical ones; that is, typicality is in inverse measure of complexity. Figure 4 illustrates this idea. In the figure, features are assigned to concepts in accordance with the classical view; for example, the defining features of bird are contained in the defining features of robin and

Animal	Bird	Robin	Chicken
F ₁	F ₁	F ₁	F ₁
	F ₂	F ₂	F ₂
		F ₃	F ₄
			F ₅

Figure 4. Concept representations in a complexity model

chicken. In addition, chicken is assumed to contain more of its own defining features—those that distinguish it from other species of birds—than does robin. This is in keeping with our critical assumption.

These assumptions suffice to explain all simple typicality effects. Since atypical probes contain more features than typical ones, atypical probes will require longer comparison stages, and consequently they will be categorized more slowly. Furthermore, because atypical probes require more comparisons than typical ones, they are more likely to lead to an error (assuming each comparison has some fixed probability of being in error). The complexity model also provides a reasonable explanation of the other typicality effects reported by Rosch and Mervis. The fact that typical concepts are learned before atypical ones becomes just another example of simple concepts being mastered before complex ones (see Brown, 1973). And the fact that typical concepts serve as reference points may just indicate our preference for simple concepts as anchors. Then there is the question of how this model would interpret the typicality ratings themselves. The simplest possibility is that subjects rate the similarity of the probe to the target concept, with similarity (1) increasing with the number of features shared by probe and target and (2) decreasing with the number of probe features not present in the target (Tversky, 1977). Now, (1) is constant across all members of a target concept (each contains all the defining features of the target), but (2) must increase with the number of defining features in a probe; it follows that typical probes, which contain fewer of their own defining features, will be judged more similar to the target concept and hence rated more typical. The fact that typical items are more similar to their parent concepts also accounts for the one remaining

simple typicality effect – when given a concept and asked to produce its instances, a subject will name typical members first. The concept is essentially a memory probe, and research on memory retrieval indicates that items similar to the probe are retrieved first (see Tulving, 1974).

Despite what we have just discussed, we have no faith in the complexity model. For one thing, it is inconsistent with the finding that it takes no longer to respond to atypical than to typical probes when the probe is not a member of the target concept. For example, it takes no longer to disconfirm “a chicken is a fish” than “a robin is a fish,” even though chicken supposedly has more features than robin and consequently should require more comparison time (see Smith, Shoben, and Rips, 1974). Another problem for the complexity model comes from studies in which subjects are asked to list features of various members of a concept. Such studies have found either no difference in the number of features listed for typical versus atypical members, or that more features are listed for typical members (Ashcraft, 1978; Malt and Smith, 1981a). To the extent that the listed features correspond to the true core features of the instances involved, these findings contradict the complexity model’s assumption that atypical members have more features than typical ones. To reiterate, our point in presenting the complexity model was merely to show that simple typicality effects can readily be accounted for by a model based on the classical view.³

DETERMINANTS OF TYPICALITY: FAMILY RESEMBLANCE MEASURES

Experimental Results

In addition to research showing the effects of typicality on various measures of performance, there have been some studies that have tried to specify the determinants of typicality. The most important of these is Rosch and Mervis’s work on family resemblance (1975). We first present their results and then take up the question of whether these findings are incompatible with the classical view, as has often been claimed.

Some of Rosch and Mervis’s experiments used natural concepts. The subjects were asked to list features of various subsets of a superordinate concept, like those of furniture, where the subsets varied in typicality (table is typical, lamp atypical). Rosch and Mervis showed that the distribution of listed features could provide a basis for typicality. Their analysis is illustrated in Table 4. Each feature listed for a subset is weighted by the total number of subsets that it is listed for; then, for each subset, the weights of all of its

TABLE 4 FAMILY RESEMBLANCE ANALYSIS

Article of furniture	Listed features				Family resemblance measure
Chair	F ₁ (5) ^a	F ₂ (4)	F ₃ (3)	F ₄ (2)	14
Sofa	F ₁ (5)	F ₂ (4)	F ₃ (3)	F ₅ (2)	14
Cushion	F ₁ (5)	F ₂ (4)	F ₆ (1)	F ₇ (2)	12
Rug	F ₁ (5)	F ₃ (3)	F ₇ (2)	F ₈ (1)	11
Vase	F ₁ (5)	F ₅ (2)	F ₉ (1)	F ₁₀ (1)	9
Telephone	F ₂ (4)	F ₄ (2)	F ₁₁ (1)	F ₁₂ (1)	8

Source: After Rosch and Mervis (1975).

a. Numbers in parentheses indicate how often each feature occurs in set of instances.

features are summed, yielding a measure called *family resemblance*. In Rosch and Mervis’s study, these family resemblance measures were very highly correlated with typicality ratings of the subsets. In short, an item is a typical subset or member of a concept if it contains features shared by many other members of that same concept.

To back up this conclusion, Rosch and Mervis performed experiments with artificial concepts. In these experiments, subjects learned to assign visually presented letter strings to categories, with six strings belonging to one category and six to another. Table 5 illustrates two of the categories used. Each string can be treated as an instance of a concept, each letter in a string as a feature. Since no information but the letters was available to subjects, we may assume that the cores of the concepts were restricted to the letters. (This kind of assumption – that the core is restricted to the obvious perceptual features – is standard in studies of artificial concepts.) Note that the strings were constructed so that no feature (letter) was common to all instances of a concept, that is, there were no salient defining features for the concept. Though all strings in Table 5 contain the same number of letters, they vary with respect to their family resemblance scores; for example, a high family resemblance string like PHMQB contains letters that were usually shared by other strings in its category, while a low family resemblance string like JXPHM contains letters less likely to be shared by other strings. The results showed that the higher the family resemblance score of a string, the sooner it could be learned as a concept instance, the more quickly it could be categorized once learned, and the more typical it was rated of its concept. Thus an experimental manipulation in the distribution of features – that is, in family resemblance – produced many of the simple typicality effects found

TABLE 5 ARTIFICIAL CATEGORIES USED BY ROSCH AND MERVIS (1975)

Category	Letter string	Family resemblance measure
Category A	JXPHM	15
	XPHMQ	19
	PHMQB	21
	HMOBL	21
	MQBLF	19
	QBLFS	15
Category B	GVRTC	15
	VRTCS	19
	RTCSF	21
	TCSFL	21
	CSFLB	19
	SFLBQ	15

with natural concepts. This is solid evidence that the distribution of features is the cause of simple typicality effects. Rosch, Simpson, and Miller (1976) have replicated some of these important findings.

Implications for the Classical View

What exactly are the implications of these results for the classical view? From the results with natural concepts, one could make the following argument.

1. The typicality variations observed when people categorize members of a superordinate concept are highly correlated with variations in family resemblance scores of the members.
2. The variations in family resemblance scores are due to features that are not common to all members. (A common feature would simply add a constant to all family resemblance scores—see Tables 4 and 5—and hence could not influence the correlation with typicality.)
3. Therefore, typicality variations cannot be explained by variations in the defining features of the superordinate concept.
4. Therefore, typicality variations must be accounted for in terms of nondefining features, but the classical view precludes the latter.

We have no quarrel with premises 1 and 2, nor with conclusion

3, but conclusion 4 is fallacious. This is illustrated by Table 6, which is the same as Table 4 save two exceptions: (1) one common feature, F_0 , has been added to all members; this feature is assumed to be defining for the superordinate concept of furniture, and it adds a constant of 6 to the family resemblance scores of all members; and (2) all features listed for a member are assumed to be defining of that subset (for example, chair's defining features are F_0 - F_4). This example is consistent with the first three steps of the argument just presented but is inconsistent with the critical conclusion that typicality variations *must* be accounted for in terms of nondefining features. For Table 6 shows that, at least in principle, typicality variations can be explained in terms of variations in the defining features of the members; that is, the defining features of the individual members vary in frequency, and this variation may be responsible for the concomitant variation in typicality.

We seem to have saved the classical view from family resemblance. There are problems, however, with our rescue mission. One stems from our assumption that the features listed for concept members are defining ones. Inspection of the features actually listed makes this unlikely. For example, many people list as features of chair "made of wood" and "has four legs." Clearly these features are not true of all chairs, and consequently they are not defining of chair. Other problems arise from our assumption that the superordinate concept, furniture, may be represented by a defining feature. If there is such a feature (or features), why did subjects not list it? (No feature was listed for all instances of furniture.) Perhaps the feature was too abstract for naive subjects to verbalize; perhaps it was too obvious for anyone to mention; but then again perhaps it just wasn't there. Furthermore, this assumption about a defining feature is clearly unnecessary to explain the data of interest. Rosch and Mervis found the same relation between typicality and family resemblance with their artificial concepts, and these concepts were constructed in such a way that they had no obvious defining feature.

The problems mentioned above hinge on whether listed features are valid indicators of the true features of concepts. Putting this issue aside, there are further problems when we try to construct a classical-view model for how the representations in Table 6 could be processed so as to yield a correlation between family resemblance scores and categorization time (the latter being known to correlate with typicality). The complexity model won't do. It holds that the critical factor is the number of features in a concept member, whereas the data show that the critical factor is the distribution of the member's features. Let us try to construct another model of the same sort, that is, one in which the features of the probe and

TABLE 6 HOW FAMILY RESEMBLANCE MEASURES CAN BE CONSISTENT WITH DEFINING FEATURES

Article of furniture	Listed features	Family resemblance measure
Chair	F ₀ (6) ^a F ₁ (5) F ₂ (4) F ₃ (3) F ₄ (2)	20
Sofa	F ₀ (6) F ₁ (5) F ₂ (4) F ₃ (3) F ₄ (2)	20
Cushion	F ₀ (6) F ₁ (5) F ₂ (4) F ₆ (1) F ₇ (2)	18
Rug	F ₀ (6) F ₁ (5) F ₃ (3) F ₇ (2) F ₈ (1)	17
Vase	F ₀ (6) F ₁ (5) F ₂ (4) F ₈ (1) F ₁₀ (1)	15
Telephone	F ₀ (6) F ₂ (4) F ₂ (2) F ₁₁ (1) F ₁₂ (1)	14

a. Numbers in parentheses indicate how often each feature occurs in set of instances.

target concepts are accessed and compared, with the comparison process starting as soon as any features are available and continuing until all the features of the target have been matched or at least one has been mismatched.

In this kind of model the distribution of members' features could affect either the access or comparison processes. Both possibilities have their problems. Consider first the possibility that the access process is affected. Since a positive categorization must be based on retrieval of F₀ from this probe (it is the only feature that matches the defining feature of the superordinate concept; see Table 6), we seek a model in which the access time for F₀ can be affected by the access times for the other features in the probe concept. Such a model is embodied in the following assumptions:

1. Assume that the processing capacity for accessing features is limited.
2. Assume further that the amount of capacity needed to access a particular feature decreases with the frequency of that feature in concept members. This means that the amount of capacity needed to access features that define typical members (like F₁ and F₂) will be less than that needed to access features that define atypical members (like F₁₁ and F₁₂).
3. Therefore, more capacity can be devoted to accessing F₀ when it occurs in a typical probe than in an atypical one.
4. Therefore, F₀ should be accessed faster in typical than in atypical members, which in turn implies that typical members should be categorized faster.

Though this model accounts for the typicality effects associated with Table 6, it has a serious problem. The model implies that the fewer features a concept member contains, the faster it will be categorized (the fewer the features, the more capacity can be devoted to accessing each one); but Malt and Smith (1981a) found that those concept members that had a minimal number of features (as determined by attribute listings) are categorized slowest of all.

A similar situation results if we try to construct a model in which the distribution of members' features affects the comparison process. Now we would assume that the capacity needed to compare a feature decreases with the frequency of the feature in concept members. Consequently, less capacity will be needed for comparing features that define typical members, which means that more comparison-capacity can be devoted to F₀ (the feature that defines the superordinate) when it occurs in typical members, which in turn implies that typical members should be categorized faster than atypical ones. Again, though, this model erroneously predicts that the fewer features a concept member contains, the faster it will be categorized.

What this shows is that some simple classical models cannot explain the relation between certain typicality effects and family resemblance scores. We can construct more complex classical models that will do the job, but the ones we have tried all require some ad hoc assumptions (like the one that the common feature, F₀, though it occurs very frequently, is so complex that it is processed slower than features that occur less frequently). In short, we feel that if the Rosch and Mervis results are to be explained by the classical view, ad hoc processing assumptions are needed, which reflects badly on the view.

To summarize all we have said about the Rosch and Mervis results, the critical result is that typicality variations are due to the distribution of features of concept members. This by itself is not, in principle, inconsistent with the classical view, as is demonstrated by our example in Table 6. But to make the result consistent with the view, we had to make some precarious assumptions about the relations between the listed features and the true features of the concepts. To account for the results in terms of a classical model, we would have to make even more precarious assumptions. Like other findings we have covered and will cover, the present results do not decimate the classical view, but they do provide reasons for lessening our belief in it.

USE OF NONNECESSARY FEATURES

Though Rosch and Mervis's work (1975) centered on an explanation of typicality effects, we noted that one of its main implications

for the classical view involved the possible use of nondefining features. The work we now wish to discuss offers a more direct approach to the issue of nondefining features: it specifically tries to show that people use nonnecessary features in categorization.

A good example of this work is a study carried out by Hampton (1979). One group of subjects listed features that characterized concepts like bird, fruit, tool, and so on. Next, they rated the extent to which subsets of these concepts had the features listed. For example, if subjects listed "flies" for bird, they might specify that robin has this feature while chicken does not. These ratings were then used to predict categorization times for another group of subjects. The more features that were shared by a concept and one of its members, the faster that member could be categorized; that is, the number of shared features between member and concept was a good measure of the typicality of that concept member. There are two critical points:

1. Some features listed for a concept were nonnecessary ones (for example, "flies" for bird).
2. These nonnecessary features were correlated with categorization performance.

It follows by a correlational syllogism that

3. Nonnecessary features are used in categorization.

While this argument is similar to the one ascribed to Rosch and Mervis (1975), it is stronger because it involves a direct assessment of the nonnecessary features of the concept.

The conclusion from the argument given above is difficult to reconcile with categorization models based on the classical view. It is clearly incompatible with the complexity model, which assumes that when one has to decide whether a probe names an instance of a target concept, only necessary features of the target concept are considered. Indeed, any classical-view model that restricts itself to necessary features (an obvious restriction) must be incompatible with the use of nonnecessary features in categorization.

How can the classical view get around this argument? The weak point in Hampton's experiment, of course, is the assumption that the listed features correspond exactly to the true defining features of a concept, a point noted in our earlier discussion of Rosch and Mervis. It seems most unlikely that just anyone off the street can readily list the features of a concept, particularly since such a list may be sensitive to contextual factors, and since sophisticated semanticists have been unable to compose lists of defining features after decades of study. Thus a proponent of the classical view might argue that

the "features" listed for a particular concept are simply an epiphenomenon in the following sense:

1. Categorization is based on defining features that are relatively inaccessible, or at least difficult to report or introspect on.
2. These defining features, however, are correlated with other nonnecessary features that are relatively easy to report.
3. Consequently, the correlation between categorization and nonnecessary features is being mediated by defining features.

Though this rebuttal is legitimate, it is based on the assumption that unspecifiable defining features just happen to be correlated with specifiable undefining ones. The rebuttal thus capitalizes on its own ignorance of defining features—a rather dubious state of affairs, and one that makes us take the findings on nonnecessary features as a serious problem for the classical view.

There is another source of evidence for the use of nonnecessary features in categorization that does not require subjects to list features. This source is based on multidimensional scaling studies (for example, Rips, Shoben, and Smith, 1973; Caramazza, Hersch, and Torgerson, 1976; Shoben, 1976). Subjects are given pairs of concepts from a particular domain, like robin-sparrow and robin-hawk, where each pair contains two subsets of a generic concept (birds); they are also given pairs that include the concept and one subset (for example, robin-bird, hawk-bird). The task is to rate each pair for its similarity of meaning. These ratings then become input to a scaling program whose output is a geometric space. The points in the space represent the items involved, while the distance between any pair of points reflects the dissimilarity between the two items. Figure 5 illustrates such a space for the concept of bird and 12 of its subsets. The representation seems a reasonable one, since similar birds are close together (for example, hawk and eagle) while dissimilar ones fall far apart (robin and goose).

There are two critical points about this multidimensional space:

1. The horizontal dimension appears to reflect variations in size while the vertical one depicts predatory relations or ferocity, where neither of these dimensions specifies anything necessary about being a bird or being any particular subset of a bird.
2. The distance between bird and any of its subsets correlates highly with how long it takes to categorize that subset as a bird.

Again it follows by a correlational syllogism that




Figure 5 Multidimensional space for bird and 12 of its subsets

3. Nonnecessary properties are being used in categorization. (The properties here take the form of dimensions rather than features, but recall that any dimension can be represented by a set of nested features.)

Though one can raise arguments against taking these results at face value (see, for example, Clark and Clark, 1977, chap. 11; Tversky, 1977), the nature of such arguments differs from those raised against the feature-listing studies. In short, accepting that nonnecessary features are used in categorization is beginning to seem more parsimonious than accepting the arguments needed to salvage the classical view.

NESTED CONCEPTS

A final problem for the classical view stems from its third assumption: If concept X is a subset of concept Y, the defining features of Y are nested in those of X. Figure 6 shows the implica-

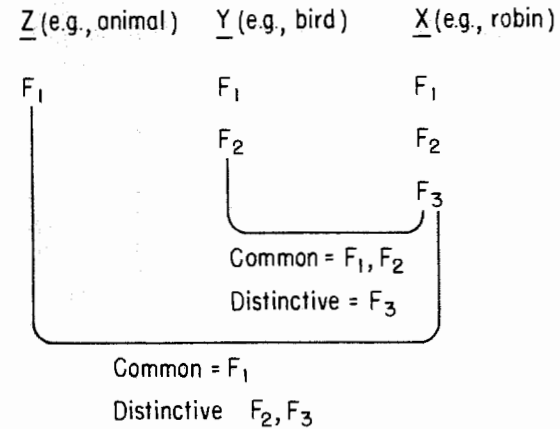


Figure 6 Implications of the nesting assumption

tions of this assumption for a nested triple of concepts X, Y, and Z, where X is a subset of Y and Y a subset of Z. As is clear from the figure, X and Y have more common features and fewer distinctive ones than X and Z, which implies that X should be judged more similar to Y than to Z (Tversky, 1977). Thus the third assumption of the classical view implies that a subset (for example, robin) should always be judged more similar to an immediate superordinate (for example, bird) than to a distant one (for example, animal). That this prediction has sometimes failed is a major problem for the classical view.

To judge the severity of this failing, we need to consider some specific results. One indirect way to measure the similarity between a subset and its superordinates is to give the subset to a group of subjects and ask them to produce its superordinates. The frequency with which a particular superordinate is produced is then a measure of its similarity to the subset. To illustrate, if the subset is rose, and 18 of 20 subjects produce flower as a superordinate, while only 10 produce plant, then rose is more similar to flower than to plant. This technique was used by Loftus and Sheff (1971), and their results showed that a subset was no more likely to produce its immediate superordinate than a distant one. Taken at face value, this finding is at odds with the classical view (see Smith, Shoben, and Rips, 1974).

As usual, there is good reason not to take the results at face

value. The production frequency of a superordinate term probably depends not only on its similarity to the subset but on its general accessibility as well. Thus a better way to measure the similarity of a subset to its superordinates is by direct similarity ratings, that is, by having subjects give numerical ratings of the similarity of a subset to both its immediate and its distant superordinates. Numerous studies have used this technique (Rips, Shoben, and Smith, 1973; Smith, Shoben, and Rips, 1974; McCloskey, 1980; Roth and Shoben, 1980), and it was generally found that the majority of subsets are rated as more similar to their immediate than to their distant superordinates. There are some exceptions to this finding, however; some of them might be due to the use of unfamiliar superordinates such as alloy and mammal (McCloskey, 1980), but others cannot be explained away by familiarity. Thus chicken and duck are consistently rated as more similar to animal than to bird, and these seem to be clear-cut counterexamples of the classical view's prediction that a subset is more similar to its immediate than its distant superordinates. The bottom line is that although the classical view's prediction works in most cases, it does not work in all.

In addition to the similarity problem, the use of nested triples of concepts in categorization studies had led to another difficulty for the classical view. The problem is that classical-view models, like our complexity one, would predict that a probe concept should be categorized faster when the target concept is a distant superordinate than when it is an immediate one. For example, robin should be categorized faster as an animal than as a bird. But this prediction has often been disconfirmed. The reasoning behind the critical prediction is as follows. For a nested triple, the distant superordinate must contain fewer features than the immediate one—for example, animal has few features than bird—which is just the third assumption of the classical view at work again. And the fewer features there are in the target concept, the fewer must be compared in the comparison stage of our classical-view model, and the less time is needed to decide that the probe concept is indeed a member of the target. This prediction falls out of any classical-view model that assumes categorization is based on a limited-capacity comparison of probe and target features.

We will briefly summarize the experimental literature on this point. Early studies showed that categorizations were faster when the target was an immediate than a distant superordinate (Landauer and Freedman, 1968; Collins and Quillian, 1969; Meyer, 1970). This finding directly contradicts classical-view models of the sort described above. Later studies, however, found few consistent

effects (Smith, Shoben, and Rips, 1974). And the most recent experiments show that in a majority of cases categorizations are faster with immediate than distant superordinates, though there are some exceptions (for example, Roth and Shoben, 1980). The exceptions turn out to be just those cases where the instance was rated as more similar to its distant than its immediate superordinate. For example, it takes longer to categorize chicken as a bird (an immediate superordinate) than as an animal (a distant superordinate).

We can summarize the foregoing discussion by two critical points:

1. With respect to similarity judgments, the classical view predicts an advantage (higher similarity ratings) for immediate over distant superordinates; this prediction works for a majority of cases, but there are definite exceptions.
2. With respect to categorization times, straightforward models based on the classical view predict an advantage (faster times) for distant over immediate superordinates; this prediction fails for a majority of cases, but works in a minority.

We can see no way to overcome point 1. It rests on the notion that similarity increases with common features and decreases with distinctive ones, and to maintain otherwise seems downright implausible. We can try to get around point 2, however, by going to a different kind of categorization model. To illustrate, we might have a model that computes the similarity between the probe and target concepts, and responds affirmatively ("the probe is a member of the target") as soon as the similarity score exceeds some threshold. This model predicts faster times for immediate than distant superordinates, which is consistent with the majority results, but it no longer handles the exceptions. In sum, the results with nested triples are difficult to reconcile with the classical view.

SUMMARY

Table 7 summarizes the four sets of experimental findings just discussed. Again we need to ask, how badly do they damage the classical view? We argued at length that the first set of findings, simple typicality effects, do not really tarnish the classical view because they can be readily explained by the complexity model that is based on this view. The other three sets of results, however, pose serious problems for models based on the classical view.

The second set included Rosch and Mervis's family resemblance results (1975), which showed that typicality variations are cor-

TABLE 7 FOUR EXPERIMENTAL CRITICISMS OF THE CLASSICAL VIEW

1. Simple typicality effects — ratings, categorization times and errors, ease of learning, order of production, and cognitive reference points
2. Determinants of typicality — typicality and the distribution of features across concept members
3. Use of nonnecessary features
4. Nested concepts — similarity and categorization times

related with variations in the distribution of features across concept members. To accommodate these results to the classical view required many ad hoc assumptions; some were needed to explain the relation between listed and true features, while others arose in the effort to specify a classical-view model that could handle the critical results. All told, these results place a heavy burden on the classical view.

The third set of findings consisted of experimental demonstrations that people use nonnecessary features in making semantic categorizations. Taken at face value, these demonstrations constitute strong evidence against the classical view because any model based on the view would presumably be restricted to necessary and sufficient features. One could, however, challenge whether these experiments tapped the real features of concepts, but such challenges seem to invoke more tenuous assumptions.

Finally, we considered findings on nested concepts. The classical view clearly predicts that a subset should be judged more similar to its immediate than it distant superordinate. Although this is generally true, there are some counterexamples. Unless all counterexamples can be explained away by artifacts, they constitute solid evidence against the classical view. We also considered categorization results showing that times are generally faster for immediate than for distant superordinates, but there were some definite exceptions. The general result was inconsistent with most straightforward models based on the classical view. Though we could come up with another classical model that would handle the general result, the model would then be inconsistent with the exceptions.

Each of the last three sets of results offers some evidence against the classical view. In no single case is the evidence unimpeachable, but taken together the three sets of results start to mount a strong case. Furthermore, the three sets of results fit together like a glove. Certain nonnecessary features appear to be used in categorization, and the distribution of these features apparently leads to typicality effects as well as to cases where a subset is judged more similar to a

distant than an immediate superordinate. In contrast, salvaging the classical view from the three critical sets of results involves a great deal of patchwork, where the patching needed in one spot is of no help in another.

When the three critical sets of results are combined with the general criticisms of the classical view discussed earlier, the case against the classical view of concepts starts to look very imposing.

Radical Attempts to Salvage the Classical View

So far we have attempted to defend the classical view by either challenging the basis of the evidence used against it, or by making minimal additions to the view (like the addition of the processing assumptions of the complexity model). As noted, however, these challenges and additions are too piecemeal, and they do not add up to a parsimonious proposal. It is time to consider some more sweeping and radical attempts to salvage the classical view.

In this section we consider three such attempts. The first is based on the assumption that categorization depends on interconcept links, not on the actual features of concepts. The next approach starts by dropping the third assumption of the classical view (the one dealing with the nesting of defining features) and explores the consequences of this move for the findings we have considered. Finally, the third approach takes off from our notion that every concept contains both a core and an identification procedure; while the core may conform to the classical view, the identification procedure need not, and the latter is assumed to be the source of the findings we have considered.

ACCESS LINKS BETWEEN CONCEPTS

Suppose that in addition to acquiring classical-view concepts, people also learn direct links between them that can be used to access one concept from another. This assumption has important implications for performance in a semantic categorization task, the task that has produced most of the experimental evidence against the classical view. Specifically, when asked whether one concept is a subset of another, people check the interconcept links and not the features of concepts. That is, performance in a semantic categorization task does not reflect the contents of concepts at all but rather the ease with which one can move from one concept to another. This idea is illustrated in Figure 7, where the interconcept links are represented by labeled paths. This kind of representation is borrowed from the current network models of knowledge representations (for example, Anderson and Bower, 1973; Collins and Loftus, 1975; Norman and Rumelhart, 1975; Anderson, 1976), but the




Figure 7 Part of a network of interconcept links

present proposal differs from most network models in its claim that the concepts being linked together conform to the classical view. Indeed, the only existing network model that is close to the spirit of Figure 7 is that proposed by Glass and Holyoak (1975), for they also assumed that concepts have necessary and sufficient conditions. Some of the assumptions we make in the following analysis are taken from Glass and Holyoak (1975), while others are borrowed from Collins and Loftus (1975).

There are two critical points about the links or paths in Figure 7. First, the paths between adjacent levels vary in length, with longer paths reflecting longer access times. This assumption is capable of accounting for most simple typicality effects. When asked to judge the typicality of a probe concept to a target one, subjects base their judgments on the length of the path between the two concepts; when categorizing the probe as a member of the target concept, subjects traverse the probe-target path, where these paths are shorter for typical probes (like the robin-bird path) than for atypical probes (like the chicken-bird path); when subset relations between concepts are being learned, relations characterized by shorter paths are acquired earlier.

The second important point about the paths in Figure 7 is that shortcuts are possible; that is, there are direct paths between concepts more than one level apart, and some of these may be shorter than paths between adjacent levels. For example, there is a shortcut between chicken and animal that is shorter and hence more accessible than the path between chicken and bird. This notion helps explain the data on nested concepts. Whenever a subset is judged more similar to its distant than its immediate superordinate, this can be attributed to the existence of a shortcut between the subset and the distant superordinate that is shorter than the path between the subset and the immediate superordinate. And since the shorter path is more accessible, the subset will be categorized faster vis-à-vis its distant than its immediate superordinate. Thus chicken will

be judged more similar to animal than to bird, and will also be categorized faster as animal than as bird. Conversely, whenever a subset is judged more similar to its immediate than its distant superordinate, there are no shortcuts (or if there are, they are relatively inaccessible); hence the subset will be categorized faster at the immediate superordinate level.

Despite its apparent successes, this access-path approach has serious drawbacks. First, although it accounts for some of the experimental findings that embarrassed the classical view, by no means does it account for all. In particular, it offers no explanation of why (1) typicality variations in natural concepts are correlated with variations in the distributions of listed features; (2) typicality variations in artificial concepts can be induced by variations in the distribution of features; and (3) various experiments have revealed evidence for the use of nonnecessary features in categorization. Second, the access-path approach does not even address some of the general empirical arguments raised against the classical view. There is nothing in the approach, for example, that comes to grips with the possible existence of disjunctive concepts, or the failure to specify the defining features of concepts. (Remember, the access-path approach assumes that there are defining features for concepts, even though they are not used in categorization tasks.) Third, in cases where the approach succeeds—as in explaining typicality effects—the access-path approach seems too unconstrained. Though there are some suggested empirical measures of the accessibility of one concept from another (roughly, the frequency with which one concept name leads to the other; see Glass and Holyoak, 1975), the approach still lacks criteria for specifying when paths are formed, what affects their lengths, and so on.

These drawbacks are serious, if not overwhelming. The move away from feature-based processes in categorization, and toward path-search processes, seems to have solved few problems. The next two attempts to salvage the classical view maintain a feature-based approach and seem a bit more promising.⁴

TRANSLATIONS BETWEEN FEATURES

Another radical modification of the classical view starts by dropping its third assumption—that the defining features of a concept are nested in those of its subsets. This assumption seems to be the cause of many shortcomings of the classical view: it is solely responsible for the prediction that a specific subset must always be judged more similar to its immediate than its distant superordinate, and it plays a major role in the prediction that there should be no unclear cases, since the nesting of a concept's defining features in

those of its subset is the basis of the classical view's supposed algorithm for determining subset membership. The nesting assumption may even be partly responsible for the failure of semanticists to find defining features for many concepts, because the search may have been overly constrained to features that meet the nesting criterion. Thus, dropping the nesting assumption seems a reasonable starting point for a modification of the classical view.

Under this modification, classical-view concepts are still assumed to be summary representations containing necessary and sufficient features, but now at least some features of a particular concept do not have to be identical to those in the concept's subsets. But in such cases, how can one use features to determine if one concept is a subset of another? The simplest answer is that there are rules or relations for directly translating one feature into another. This idea is illustrated in Figure 8. One defining feature of bird is "animate," and it is listed for robin but not for chicken. However, chicken includes the defining feature "egg-laying," where the latter implies animate. Now one can establish that a target concept includes a probe either by matching the target's features to those of the probe (as in the bird-robin case in Figure 8) or by using the interfeature relations to translate one feature into another (the chicken-bird case).

These ideas can be used to account for most of the experimental findings discussed earlier. To explain simple typicality effects, we assume that the less typical a subset, the more features it contains that require translation when it is compared to its parent concept. In Figure 8, for example, the atypical chicken requires translation of its egg-laying feature, while the more typical robin needs no translation at all. Assuming a translation operation requires more time and is more error-prone than a simple matching operation, we would expect atypical instances to be categorized slower and less accurately than typical ones. To account for Rosch and Mervis's finding (1975) that typical subsets contain features that are fre-




Figure 8 Direct translation between features

quently listed for other subsets of the concept, we might assume that features listed for many subsets are more likely to be part of the concept's representation and hence require no translation. Again, typicality comes down to an inverse measure of the amount of translation needed in relating two concepts.

The data on nested concepts also fit nicely with this translation approach. When the features of a subset do not require any translation vis-à-vis the parent concept, we have a perfect nesting of the concept's features in those of the subset, and we expect the similarity predictions of the classical view to hold; when some features do require translation, we have less than perfect nesting, and exceptions to the similarity prediction are expected. Moreover, since atypical subsets are more likely to require translation, such subsets should constitute the bulk of the exceptions, which seems to be the case.

There is an interesting aspect to the above arguments. In all cases, the typical members of a concept are treated in roughly the way the classical view specifies (few or no translations are needed), whereas atypical members are treated differently (translations are frequently required). The direct-translation approach thus has the character of a rule-plus-exception approach, typical members being handled by the rule (the classical view) and atypical members being handled as exceptions. This approach has the desirable property that much of the machinery of the classical view is salvaged, the view now being limited to the more typical members.

The preceding discussion highlights some strong points of the translation approach. There is, however, one sore point for this approach when it comes to accounting for experimental findings. In explaining Rosch and Mervis's results (1975), we explicitly assumed that features listed of many subsets were often defining of the concept, and implicitly assumed that features listed for any subset were defining of it. Inspection of feature listings like those collected by Rosch and Mervis, however, provides little support for either assumption. Some frequently listed features are clearly non-necessary for parent concept or subset. This relates to a more general problem: the translation approach we have sketched has no natural way of dealing with the use of nonnecessary features in categorization, which is one of the major experimental findings.

Though the direct-translation approach can also be stretched to deal with some of the general empirical arguments raised against the classical view, most of what can be said is quite vague. We have already noted that it may be easier to specify defining features of concepts if we drop the nesting constraint. But this is merely a promissory note, particularly since the direct-translation approach

still assumes a substantial degree of feature nesting between concepts and their typical members. The occurrence of unclear cases may be attributed to the need to translate between features. Perhaps there are subset relations, like "tomato-fruit," in which the needed translation relations are imperfect or unknown. Without some further specification of the actual interfeature relations, though, it is hard to evaluate this possibility. And finally, there is the question of what the direct translation approach has to say about disjunctive concepts. One possible answer might go as follows: Concepts are truly conjunctive, but they may appear disjunctive when we focus on their members and notice that the latter do not share many features; to illustrate with the concept of bird, we might focus on robin and chicken and note that one contains "animate," the other "egg-laying," and mistakenly conclude that bird is disjunctive. Although this answer is possible, we are hard pressed to put much credence in it.

Some of the problems just mentioned may be alleviated by a more extreme form of the translation approach.⁵ So far we have assumed a direct translation between features, but translation can also operate indirectly in that two defining features (of two different concepts) might be connected by a third feature. To illustrate, a defining feature of fruit might be "seeds," while such a feature of orange might be "acidic." To translate between them, one could use the information that (1) seeded objects are often juicy and (2) acidic objects are often juicy. Hence both seeds and acidic lead to juicy, so juicy serves to translate between the two defining features. Note that this translation process is probabilistic, since the intermediate feature, juicy, is not true of all seeded objects. This kind of indirect translation is consistent with the use of nonnecessary features, for intermediate features may be nonnecessary yet used in the categorization process.

An indirect-translation process can also be elaborated to account for other experimental findings. Two examples should suffice. First, variations in the typicality of concept members might reflect either the necessity of translation or the ease with which the defining features of the members can be translated into the defining features of the concept. Second, the Rosch and Mervis correlations between typicality and the distribution of listed features (1975) might be explained by assuming that the listed features are those used in the translation process; for example, features used to translate many concept members may be more powerful or accessible than features used to translate few members.

The obvious problem with the indirect-translation approach is that it is extremely ad hoc. No constraints of any kind have been

placed on the nature of intermediate features or on their relations to defining ones, and it is unclear what, if any, predictions follow from this approach. Nevertheless, indirect translation at least addresses a wide range of problems and may, if properly developed, have the potential to save the classical view. Such a rescue will come at a high price, however: for all explanations of empirical results will be in terms of intermediate features rather than defining ones. Indeed, the defining features, which are the heart of the classical view, seem to be doing no theoretical work at all, and it becomes unclear why one need posit them to explain categorization. To some extent, the same is true of the direct-translation approach, where most of the explanations hinged on relations between features rather than on the features per se. And as we will see, the identical problem arises in our third way of salvaging the classical view.

GREATER ACCESSIBILITY OF IDENTIFICATION PROCEDURES

Until now we have assumed the following about concept cores and identification procedures:

1. Though many concepts contain an identification procedure as well as a core, the core is more important psychologically because it must always be there and because it determines the contents of the identification procedure.
2. We must thus focus on empirical methods that are likely to involve the core rather than the identification procedure.
3. Therefore, we should concentrate on semantic categorization tasks rather than perceptual ones, since the semantic task requires a consideration of only the cores.

By challenging the third assumption, we generate our third attempt to salvage the classical view. Specifically, we now assume that the identification procedure is more accessible than the core; consequently, semantic categorization is often based on a comparison of the identification features of the target and probe concepts rather than on a comparison of the core's defining features. This new proposal also partly undermines assumption 1 above: if identification procedures are so widely used, it may be misleading to call the core "more important psychologically."

Given that semantic categorizations may be based on a comparison of identification procedures, most of the relevant experimental results fall into place. First, an immediate consequence is that many features used in categorization will be nonnecessary ones, since the features in identification procedures will often be

nonnecessary. Second, simple typicality effects can now be explained in terms of similarity of identification features. The more similar are the identification features of a concept to the identification features of one of its members, the more typical that member is judged to be. Now a categorization model that computes the featural similarity between target and probe concepts will yield faster and more accurate decisions for typical members. Third, the Rosch and Mervis finding (1975)—that typical members contain features common to many other members—can be interpreted solely in terms of identification features. Identification features common to many members are likely to be included in the identification procedure of the concept itself, so again typicality comes down to a matter of similarity of identification features between a concept and its members. Finally, since a concept's identification features need not be perfectly nested in those of its subsets, the similarity prediction of the classical view is no longer expected to hold in all cases.

In addition to accounting for the above results, the assumption that semantic categorizations are based on identification features has another important experimental consequence: semantic categorizations should resemble perceptual ones, since both rely on the same identification features. Hence our earlier stricture against using results from perceptual categorization to evaluate views of concepts must be temporarily suspended, and we need to look briefly at some comparisons of semantic and perceptual categorization.

Though there are few research reports that afford a detailed comparison between semantic and perceptual categorization, what is available shows marked similarities between the two. If one asks subjects to rate a concept's members for typicality, the typicality ordering will be virtually identical for members presented as pictured instances and for members expressed as words (see Smith, Balzano, and Walker, 1978). Moreover, aside from the fact that pictures are responded to slightly faster than words, the effects of typicality are the same in perceptual and semantic categorization (see Guenther and Klatzky, 1977). Also, to the extent that comparable data are available, items that lead to faster categorizations with immediate than with distant superordinates do so regardless of whether the item is presented as a picture (that is, as a specific instance) or as a word (that is, as a subset; compare Smith, Shoben, and Rips, 1974, with Smith, Balzano, and Walker, 1978). And finally, the features listed for items presented as words substantially overlap the features listed for these same items presented as pictures (compare, for example, Hampton, 1979, with Rosch et al., 1976). Although some of these comparisons are tenuous because they in-

volve contrasting results from different studies, they at least suggest that many of the same features are used in semantic and perceptual categorization. This bolsters our new assumption that identification procedures underlie semantic categorization.

The idea of accessible identification procedures can also offset one of the general empirical arguments raised against the classical view, namely, that some concepts are disjunctive. As mentioned earlier in this chapter, some concepts may appear disjunctive when we mistakenly focus on their identification features rather than on their core features. However, a reliance on identification features is of less help in warding off other general empirical arguments. Positing accessible identification features does not explain the prolonged failure to specify the defining features of the core; nor does it get rid of the problem of unclear cases: for example, if asked, "Is a tomato a fruit?" and given sufficient time to mull it over, one should be able to use the core features to resolve the matter.

Despite the two deficiencies just noted, the present approach gets rid of many of the classical view's empirical problems and correctly predicts a similarity between semantic and perceptual categorization. And it does not seem as unconstrained as the other attempts to salvage the classical view, since identification features would at least be restricted to properties that people actually use in deciding that a physical object is an instance of a concept. However, as was the case with the translation approach considered earlier, the present approach seems to save the classical view by shifting all the theoretical action away from the defining features of the core. Again we may raise the question, why bother to posit classical-view concept cores at all?

One possible answer that seems reasonable to us is that while the identification procedure may form the front line of categorization, the core is used as a backup procedure (this is similar to an argument made by Katz, 1977). That is, difficult categorizations, ones that cannot be done by the identification procedure, are eventually tackled by the classically defined core. This claim is seriously challenged, however, by the presence of unclear cases. Another possible answer is that while an identification procedure generally takes care of categorization, the core plays a major role when we do things with concepts other than categorizing. We may, for example, work mainly with classically defined cores when we combine simple concepts into complex ones (the conceptual-combination function of concepts), or when we draw inferences from existent propositional representations. We will not take up the pros and cons of this proposal since it would take us too far from the mainstream of this book. All we can conclude is that,

when restricted to categorization phenomena, the proposal of accessible-identification procedures essentially salvages the classical view by ignoring it.

SUMMARY

All three attempts to salvage the classical view have their problems. The first, or access-path approach, seems the most problematic. It fails to deal with most criticisms of the classical view—namely, the correlation between feature distributions and typicality, the use of nonnecessary features, the apparent presence of disjunctive concepts, and the inability to specify defining features. Moreover, the access-path approach runs up this rather impressive list of failings while imposing overly powerful and unconstrained assumptions.

The translation approach fares better. Direct translation seems to handle readily several experimental findings in a parsimonious fashion. However, it lacks convincing accounts of the general empirical problems of the classical view (for example, failure to specify defining features), and it cannot explain the use of nonnecessary features in categorization. The latter problem could be solved by positing indirect translation, with its notion of intermediate features, but at the cost of introducing many unconstrained assumptions.

The final approach, accessible identification procedures that determine semantic as well as perceptual categorizations, appears to have the most potential. However—and this is the critical point—all of its potential in handling categorization phenomena seems to be dependent on the nonnecessary (identification) features, and such nonnecessary features are the backbone of the probabilistic view. In short, the most promising attempt to salvage the classical view is promising because it moves toward the probabilistic view.

4 | The Probabilistic View: Featural Approach

IN DEALING WITH the classical view we restricted ourselves to featural descriptions of concepts, which helped us to state the view's assumptions with a bit of precision. Unfortunately, things are different with the probabilistic view. Some proponents of this view have used features to describe concepts, others have employed dimensions, and still others have rejected property descriptions altogether and favored holistic patterns instead (see Chapter 2 for a discussion of these issues). As a result, it is difficult to say much that is substantial about the probabilistic view in general. We can only come up with the following two assumptions: (1) the representation of a concept is a summary description of an entire class; and (2) the representation of a concept can *not* be restricted to a set of necessary and sufficient conditions; rather, it is some sort of measure of central tendency of the instances' properties or patterns. The first assumption is identical to the initial one of the classical view, so it cannot tell us anything about the differences between the views. The second assumption, which ideally should contain the essence of the view, is partly negative, partly vague. It can only be made more positive and more precise when one considers separately the featural, dimensional, and holistic approaches to the probabilistic view.

This is the tack we will take in studying the probabilistic view. In this chapter we take up the featural approach; we will consider how its representational assumptions avoid the difficulties faced by the classical view, how the process models it has generated can account for experimental results that plagued its classical predecessor, and what prices are being paid for these gains. The succeeding two chapters will provide the same kind of treatment for the dimensional approach and the holistic approach to the probabilistic view.