Gentner D, Brem S, Ferguson R *et al*. (1997) Analogy and creativity in the works of Johannes Kepler. In: Ward TB, Smith SM and Vaid J (eds) *Creative Thought: An Investigation of Conceptual Structures and Processes*, pp. 403–459. Washington, DC: American Psychological Association.
Gopnik A and Meltzoff AN (1997) *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
Guzzetti B and Hynd C (eds) (1998) *Perspectives on Conceptual Change*. Mahwah, NJ: Lawrence Erlbaum.
Keil F (1989) *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press/Bradford Books.
Kuhn T (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
Nersessian N (1992) How do scientists think? Capturing the dynamics of conceptual change in science. In: Giere R (ed.) *Cognitive Models of Science*, vol. 15, pp. 3–44. Minneapolis, MN: University of Minnesota Press.
Thagard P (1992) *Conceptual Revolutions*. Princeton, NJ: Princeton University Press.
Vosniadou S and Brewer WF (1992) Mental models of the earth: a study of conceptual change in childhood. *Cognitive Psychology* **24**: 535–585.

Further Reading

Ball T, Farr J and Hanson RH (eds) (1989) *Political Innovation and Conceptual Change*. Cambridge, UK: Cambridge University Press.
Carey S (2001) *Science education as conceptual change*. [http://www.house.gov/science/carey_03-04.htm]
Dietrich E and Markman AB (eds) (1999) *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*. Mahwah, NJ: Lawrence Erlbaum.
Feyerabend PK (1981) *Realism, Rationalism and Scientific Method*. Philosophical Papers, vol. 1. Cambridge, UK: Cambridge University Press.
Kunda Z, Miller D and Claire T (1990) Combining social concepts: the role of causal reasoning. *Cognitive Science* **14**: 551–577.
Nersessian N (1989) Conceptual change in science and in science education. *Synthese* **80**: 163–183.
Pearce G and Maynard P (eds) (1973) *Conceptual Change*. Dordrecht: Reidel.
Thagard P (1999) *How Scientists Explain Disease*. Princeton, NJ: Princeton University Press.

Conceptual Representations in Psychology

Introductory article

Arthur B Markman, University of Texas, Austin, Texas, USA

CONTENTS

Conceptual representation refers to the way that information about categories is stored and organized.

INTRODUCTION

Concepts are mental representations that are used to divide the world into groups that will be treated as equivalent for some purpose. Concepts may refer to objects, events, or ideas. Concepts may be used for reasoning, prediction, and communication. Some researchers have distinguished between concepts, which are the mental representations of information, and categories, which are sets of objects in the world that are grouped together. Often, however, these terms are used interchangeably.

Psychologists have explored concept representations in detail. This work has examined both within-category representation and between-category structure. Within-category representation refers to the information that describes a particular category such as 'dog'. Between-category structure refers to the relationships among different categories such as that between the categories 'dog', 'cat', and 'animal'.

WITHIN-CATEGORY REPRESENTATION

The central question about within-category representation involves the way people store information about particular concepts that enables them to classify new items (exemplars) as members of a category. Some work has looked at other uses of categories such as making predictive inferences, causal reasoning, and communication, but this discussion will focus on classification. Three broad types of within-category representation are rule-based models, similarity-based models, and theory-based models.

Rule-based Models

The classical approach to concept representation has been to seek a rule that specifies the necessary and sufficient conditions for something to be a member of a category. A property is a necessary condition for being in a category if all members of that category possess the property. A set of necessary conditions is sufficient to specify a category if all exemplars that have that set of properties are members of the category, and no exemplars that have that set of properties are members of any other category. For example, an object is a triangle if it is a three-sided closed figure. This set of features is necessary and sufficient, because all triangles are three-sided and closed. No object that has these properties can be anything but a triangle.

Unfortunately, outside formal domains like geometry, it is difficult (or perhaps impossible) to find a set of necessary and sufficient conditions that specify the members of a category. For example, it might seem at first glance that a bachelor is an unmarried adult man. While this rule correctly classifies most bachelors, there are many dubious cases. For example, Catholic priests and widowers are both unmarried adult men, but one might be hesitant to classify them as bachelors. While it is possible to continue to refine this definition, it is likely that exceptions could be found to any rule that was generated.

One approach that has been tried to save rule-based approaches has been to assume that people generate fairly simple rules that are good for classifying most exemplars, and then store exceptions to the rules separately. In the example above, the rule 'unmarried adult man' would be used to classify most bachelors, but exceptions such as priests and widowers would be considered separately.

Similarity-based Models

Intuitively, it seems that a new exemplar is classified based on its similarity to the category. For example, an object might be classified as a bird, because it looks like a bird. This intuition has been captured by similarity-based models, which assume that people classify a new exemplar based on its similarity to some stored category representation. Similarity-based models differ from each other primarily in their assumptions about the nature of the stored category representation.

Prototype models assume that people store some average representation of an object. The average need not be identical to any actual exemplar, but rather contains the features most frequently associated with that category. For example, the typical bird might be a small animal that flies, sings, and has feathers. Not all birds have this entire set of properties (e.g. penguins do not fly), but the more of these features an exemplar possesses, the more likely it is to be a bird.

One result often taken as evidence for prototype models is that categories have a graded typicality structure: that is, people have strong intuitions about which members of a category are typical members of that category and which members are atypical. For example, robins and sparrows are generally thought to be typical birds, while chickens and emus are thought to be atypical birds. Generally, the typicality of an exemplar is related to its similarity to the prototype of the category.

A second prominent similarity-based model is the exemplar model, which posits that people store representations of each category member rather than creating a prototype. New exemplars are then classified by comparing them with all of the known exemplars. The more similar a new exemplar is to the known exemplars of a particular category, the more likely it is that the exemplar will be classified as a member of that category. Exemplar models are also able to account for graded typicality structure, because typical exemplars are similar to many members of a category, but atypical exemplars are similar to only a few members of a category.

Theory-based Models

Despite the success of similarity-based models in predicting how people classify new items, there are situations in which people classify items in a manner that violates similarity. For example, at a party, a person might be classified as drunk if he or she dives headfirst into a cake. This person is not being classified on the basis of any similarity to known exemplars of drunk people; rather, common beliefs about drunken behavior are sufficient to classify the person as drunk.

There seems to be a developmental change in people's ability to use theory-based information. When children first learn categories, they often classify on the basis of surface characteristics. For example, if told about a black cat that has a white stripe painted on its back, and a bag of smelly stuff surgically placed inside it, young children will classify it as a skunk. Older children and adults, however, will classify it as a cat, suggesting that they are able to use a theory about biological categories to classify this (rather strange) exemplar.

Which Type of Model is Right?

Of the three types of models, the rule-based models are least often used in conceptual processing. There are some situations in which people must make repeated classifications that involve a rule and exception process. However, empirical studies suggest that even when people are asked to form rules, their ability to apply the rule is influenced by the similarity of a new exemplar to those seen before.

Both similarity-based and theory-based processes are often used in categorization. There are times when people must be able to identify a new item on the basis of the similarity of its properties to those of items seen in the past. In addition, there are cases in which people's theories about a domain influence categorization. Current research is focusing on how to integrate similarity-based and theory-based approaches.

BETWEEN-CATEGORY STRUCTURE

A second important aspect of category representation involves the relationships among categories. In this section, two aspects of between-category structure are examined. First, much research has examined the hierarchical organization of categories. This work is concerned with understanding how people may categorize objects at different levels of abstraction. A second area that has received attention is people's ability to generate categories based on goals. This study of goal-derived categories also provides a window into conceptual processing.

Hierarchical Organization of Categories

If you see a small, curly-haired, four-legged living creature being walked on a leash down the street, you can classify this thing as a poodle, a dog, or an animal. That is, for any given object, there are a variety of categories to which it belongs. Many of these categories differ from each other in their degree of abstraction: 'dog' is a more abstract category than 'poodle', because all poodles are dogs, but not all dogs are poodles. Similarly, 'animal' is a more abstract category than either 'dog' or 'poodle'.

A striking aspect of this category structure is that if you show people a picture of some object, they are most likely to identify it using a category at a middle level of abstraction. For example, shown a picture of the item described in the previous paragraph, people are likely to identify it first as a dog rather than as a poodle or as an animal. This tendency has led psychologists to refer to this middle level of abstraction as the basic level. Categories more abstract than those at the basic level (e.g. animal) are called superordinate categories, and categories more specific than those at the basic level are called subordinate categories (e.g. poodle).

Basic level categories have been shown to have a number of characteristics. First, they tend to be the most abstract categories whose members have a common shape, and whose shape differs from other contrasting basic level categories. For example, dogs tend to be shaped similarly to each other and differently from other animals; in contrast, animals come in many different shapes. Second, basic level categories tend to have shorter labels than either subordinate or superordinate categories; for example, 'car' is the label for a basic level category, and the labels 'vehicle' (for the superordinate) and 'sports car' (for a typical subordinate category) are both longer than the basic level label. Basic level categories are also the most abstract level for which the category members share the same set of parts: cars all have wheels, brakes, and engines, whereas there are many other vehicles (e.g. helicopters and boats) that do not share these parts. Finally, children tend to learn basic level labels for objects before learning the labels for categories at other levels of abstraction. The factors that characterize basic level categories can be summarized as follows:

- Objects are identified first at the basic level.
- Objects are classified fastest at the basic level.
- The basic level is the most abstract level at which the members tend to have the same shape.
- The basic level is the most abstract level at which the members tend to share parts.
- The basic level is the most abstract level at which people interact with the members using similar motor movements.
- The basic level is the most abstract level at which the category members tend to be similar.
- The basic level is the most abstract level at which category members tend to be dissimilar from members of contrasting categories.
- Children often learn basic level labels before labels at other levels of abstraction.
- Basic level labels are shorter than labels for categories at other levels of abstraction.

The hierarchical organization of categories seems to be strongest for object categories. Some research has been done on the between-category structure of abstract concepts such as events and ideas. People also have categories at different levels of abstraction for these concepts, but the basic level does not have as much of an advantage relative to subordinate and superordinate categories.

Goal-derived Categories

The previous section suggested that an object might belong to many different categories, and that these categories generally differ in their level of abstraction. There are some categories, however, that are organized around people's goals rather than around the overall shape and parts of objects. Some of these categories are ones that we use all the time, and their labels have become words. For example, a 'pet' is a domesticated animal that is kept as a companion. Thus, membership in this category is determined by whether an object serves a particular goal.

An important observation is that people can also generate goal-derived categories as they are needed. For example, you may never have considered the category 'Things to take out of a house in the event of a fire'. Now that this category has been suggested, however, it is easy to generate members of the category (e.g. children, jewellery, photographs).

Novel goal-derived categories are called *ad hoc* categories. A striking finding is that, although they are being generated on the fly, they share many characteristics with categories that were previously learned. For example, like regular categories, *ad hoc* categories exhibit a graded typicality structure: that is, people find it easy to determine which members of an *ad hoc* category are typical or atypical. For example, people might agree that old photographs of family members are good examples of things to take out of the house in the event of a fire, but that an old sofa is a poor example.

One key difference between regular categories and goal-derived categories is in the way that typicality is assessed. For regular categories, an object is typical to the extent that it is similar to the average member (or prototype) of the category. In contrast, for goal-derived categories there is often an ideal member, and items are more typical to the extent they are similar to the ideal. For example, someone might create the goal-derived category 'diet foods'. The ideal member of this category tastes great and has no energy content. A new object will be typical of a category to the extent that it is similar to its ideal.

CONCLUSION

Psychologists have explored the internal (within-category) representation and external (between-category) structure of category representations. Research on within-category representation has focused on the role of rules, similarity, and theory in determining category representation. Research on between-category structure has focused both on relationships among categories at different levels of abstraction and on goal-derived categories.

Further Reading

Barsalou LW (1983) Ad hoc categories. *Memory and Cognition* **11**: 211–227.

Keil FC (1989) *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press.

Medin DL, Lynch EB and Solomon KO (2000) Are there kinds of concepts? *Annual Review of Psychology* **51**: 121–147.

Morris MW and Murphy GL (1990) Converging operations on a basic level in event taxonomies. *Memory and Cognition* **18**(4): 407–418.

Murphy GL and Medin DL (1985) The role of theories in conceptual coherence. *Psychological Review* **92**(3): 289–315.

Nosofsky RM (1986) Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* **115**(1): 39–57.

Nosofsky RM, Palmeri TJ and McKinley SC (1994) Rule-plus-exception model of classification learning. *Psychological Review* **101**(1): 53–97.

Rosch E and Mervis CB (1975) Family resemblances: studies in the internal structure of categories. *Cognitive Psychology* **7**: 573–605.

Rosch E, Mervis CB, Gray WD, Johnson DM and Boyes-Braem P (1976) Basic objects in natural categories. *Cognitive Psychology* **8**: 382–439.

Smith EE and Medin DL (1981) *Categories and Concepts*. Cambridge, MA: Harvard University Press.