Comparing Exemplar- and Rule-Based Theories of Categorization

Jeffrey N. Rouder[1] and Roger Ratcliff[2]

[1]*University of Missouri-Columbia and* [2]*The Ohio State University*

ABSTRACT—*We address whether human categorization behavior is based on abstracted rules or stored exemplars. Although predictions of both theories often mimic each other in many designs, they can be differentiated. Experimental data reviewed does not support either theory exclusively. We find participants use rules when the stimuli are confusable and exemplars when they are distinct. By drawing on the distinction between simple stimuli (such as lines of various lengths) and complex ones (such as words and objects), we offer a dynamic view of category learning. Initially, categorization is based on rules. During learning, suitable features for discriminating stimuli may be gradually learned. Then, stimuli can be stored as exemplars and used to categorize novel stimuli without recourse to rules.*

KEYWORDS—*human learning; categorization; perception*

We discuss theories of how people cluster items into categories and how to discriminate among these theories. One group of theories is based on *exemplars*. Category membership of a novel item is determined by its similarity to previously encountered stimuli (Medin & Schaffer, 1978). Consider, for example, a radiologist who must classify a suspicious spot on an X-ray either as a tumor or as natural tissue variation. Exemplar-based theories posit that the decision is reached by comparing the current X-ray to exemplars of X-rays in memory. If the X-ray appears more visually similar to X-rays of tumors than to those of normal tissue, the radiologist may classify it as indicating a tumor. A second group of theories is based on *rules* (Trabasso & Bower, 1968). A radiologist using rule-based categorization would observe whether specific properties of the X-ray meet certain criteria; for example, is there an extreme difference in brightness in a suspicious region relative to the other regions? A decision is then based on this property alone. Whereas exemplar-based

approaches rely on large mnemonic stores of previous experience, rule-based approaches[1] rely on simple rules or criteria and do not require extensive memory. Although there are other approaches to categorization problems, the broad dichotomy between stored exemplars and abstracted rules runs deep within the literature. From a practical point of view, understanding how categorization is performed yields insights on how to best teach pattern-recognition activities. For example, is learning the art of evaluating X-rays best accomplished through studying many exemplars or through mastering certain rule-based tips?

DISCRIMINATING AMONG MODELS

The main problem in testing categorization theories is that specific model predictions often mimic each other. Consider the task in which a person decides if a piece of footwear is better classified as a boot or as a shoe. Figure 1 shows how the physical attributes of height and width affect category membership. Boots tend to have greater heights than same-length shoes, but there is some overlap in the categories. A reasonable rule from these category distributions is that a novel piece of footwear should be classified as a boot if its height is greater than its length (this rule is shown as the dotted diagonal). Accordingly, the item denoted with an *X* should be classified as a boot. Yet, this prediction is in accordance with an exemplar-based prediction too. In the figure, there are more boot exemplars than shoe exemplars above the diagonal and fewer below it. Hence, according to exemplar theory, items above the diagonal should be classified as boots, and those below it should be classified as shoes. This experimental-design problem, in which reasonable rules covary with exemplar density, is fairly common. Consequently, it is not surprising that exemplar- and rule-based models have performed about equally well in many previous evaluations.

---

[1]For a rule-based theory to be tenable, some limits must be placed on the types of rules participants may use. Without limits, a theory may be vacuous. We restrict rules to those that describe how physical properties of the to-be-classified stimulus should be combined and evaluated. An example of an invalid rule is "categorize the stimulus as the same as similar neighbors in memory." This rule, which essentially describes an exemplar process, is invalid because it does not specify physical properties as the basis of comparison.

Address correspondence to Jeffrey Rouder, Department of Psychological Sciences, 210 McAlester Hall, Columbia, MO 65211; e-mail: rouderj@missouri.edu.
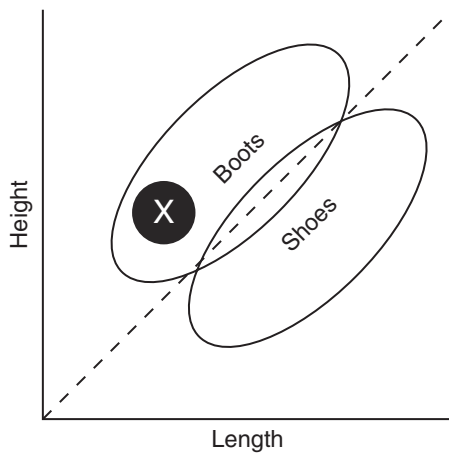
Fig. 1. Overlapping categories in a hypothetical categorization task. Although the categories of boots and shoes overlap, boots tend to have greater heights for comparable lengths than shoes do. An exemplar approach categorizes a novel stimulus (such as the one denoted by the letter *X*) based on its similarity to exemplars from that category; a rule-based approach categorizes based on criteria. The dotted line denotes a reasonable rule: If an item is taller than it is long, it is a boot.

We provided a new set of paradigms that are more diagnostic than previous ones were (Rouder & Ratcliff, 2004). Stimuli in our experiments varied on a single dimension; in the experiments we review, stimuli were squares varying only in size. Squares were assigned by the experimenter to one of two categories (labeled A and B), but the assignments were not straightforward. Squares of a given size could be from either category, but more were from one category than the other. Figure 2A shows the assignments. Extremely small squares were assigned to category A 60% of the time and to category B 40% of the time. Moderately small squares were always assigned to category A and moderately large squares were always assigned to category B. Extremely large squares were assigned to category A 60% of the time. On an experimental trial, a single square was presented and the participants' task was to decide its category assignment. The optimal strategy was to pick category A for all stimuli other than the moderately large squares. Participants adopting this strategy would always classify the moderate stimuli correctly and classify the extreme stimuli correctly on 60% of the trials.

Rule- and exemplar-based models make different predictions for how participants should classify stimuli. Figure 2B shows rule-based prediction: A reasonable rule is to assign moderately large squares to category B and all other sizes to category A. Corresponding decision bounds are shown as vertical dotted lines; the resulting predictions are shown as a thin, solid line. In rule-based models, it is typically assumed that perception is not quite veridical—there is a modest amount of trial-by-trial variability even for the same participant observing the same stimulus. The effect of this variability is to smooth the predictions (as shown with the thicker solid line). As a result, rule-based theories predict that participants will more likely assign extremely small squares to category A than they will moderately
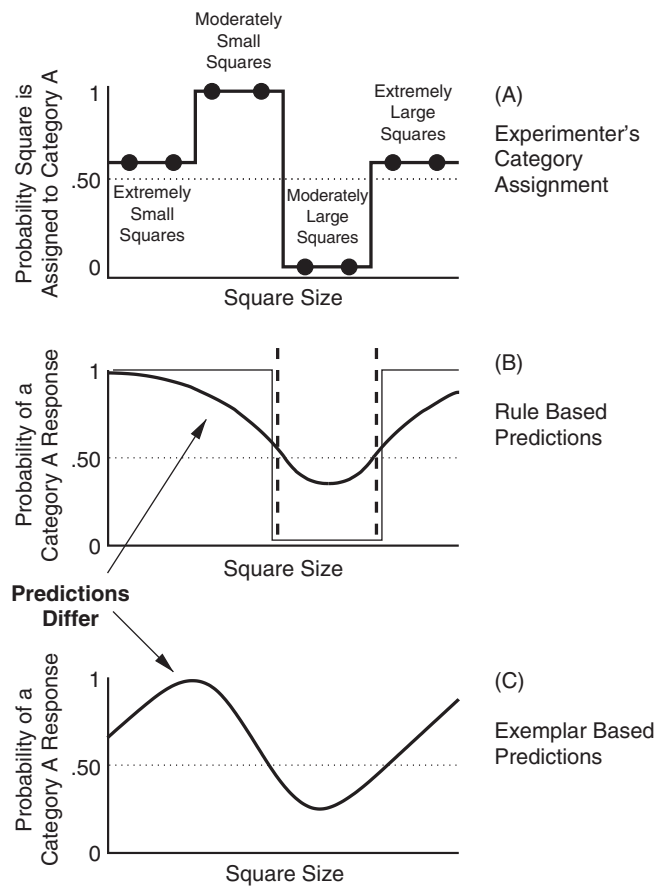

Fig. 2. Category-assignment probabilities as a function of square size (A), as well as rule- and exemplar-based predictions for participants' categorization patterns (B and C, respectively).

small ones. This prediction contrasts with that for exemplar models (Fig. 2C). Here, the proportion of category A exemplars is greater for moderately small squares than it is for extremely small ones; consequently, the category A response probability is greater for moderately small squares than it is for extremely small ones.

STIMULUS CONFUSION DETERMINES CATEGORIZATION STRATEGY

We found that participants' categorization behavior was not exclusively rule- or exemplar-based. The critical variable that determined behavior was the confusability of the set of stimuli. In one experiment, we simply manipulated the range of square sizes. If the squares covered a narrow range—e.g., they were all fairly similar in size—then a majority of the participants had a pattern of results like that in Figure 3A. This pattern is consistent with the following simple rule: "Place small-sized squares in category A and large ones in category B." When the range of sizes was moderate, more participants adopted a pattern like that in Figure 3B. This pattern is consistent with the two-
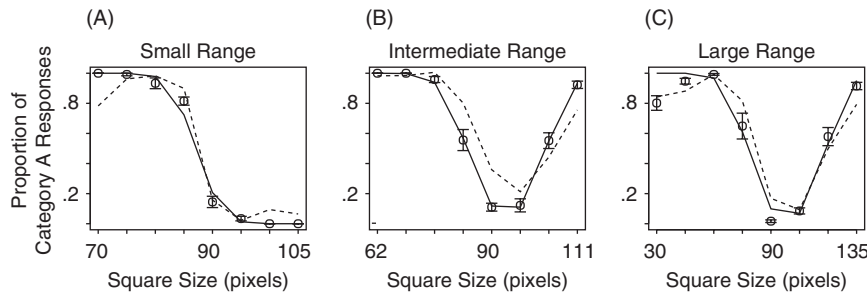
Fig. 3. Selected participants' categorization behavior for three ranges of square sizes. Circles with error bars denote category-A response proportions and standard errors, respectively. Solid and dashed lines denote rule and exemplar predictions, respectively. Panel A shows characteristic results when the square sizes ranged from 70 to 105 pixels. The pattern of responses is consistent with a single decision bound (squares below a criterion are classified in category A; otherwise they are classified in category B). This pattern is not optimal—extremely large squares are systematically misclassified. Panel B shows characteristic results when the square sizes ranged from 62 to 111 pixels. The u-shaped pattern is consistent with two decision bounds (small and extremely large squares all go in category A) and corresponds to better performance than the pattern in panel A. Panel B also characterizes about half the participants' data when the square size ranged from 30 to 135 pixels. A sizable minority of participants in this condition displayed a pattern like that in panel C (reflecting the fact that some extremely small squares are not assigned to category A), consistent with exemplar-based categorization. Results are from Rouder and Ratcliff (2004, Experiment 2).

bound rule of Figure 2B—i.e., "Place all but moderately large squares in category A." When the range was large, about half of the participants had patterns like that in Figure 3B while a sizable minority had patterns like that in Figure 3C. This later pattern is consistent with exemplar theories.

Fig. 4. Category assignments (A) and participants' behavior (B) for ruler experiment. Dip in performance for critical stimuli 2 and 5 when the ruler was presented indicates exemplar categorization. Rouder and Ratcliff (2004, Experiment 4).

In another experiment, we manipulated stimulus confusability by adding a tick-marked ruler to the display with the stimuli. The stimuli were once again assigned to categories; the probabilities are indicated in Figure 4A. The critical stimuli in this design are sizes 2 and 5. The optimal rule is a one-bound solution that separates the smaller squares from the larger ones (rule depicted with a dotted line). With the ruler, the participant could measure the square size, mitigating perceptual uncertainty. Results are shown in Figure 4B. Participants who were not provided the ruler had a decreasing pattern which in itself does not contradict either exemplar- or rule-based processing. Participants who were provided the ruler had a more complex pattern with moderated proportions for the critical stimuli. This pattern is inconsistent with rule-based categorization. The reduction of stimulus confusion with a ruler resulted in a pattern exclusively consistent with exemplar-based processing. More recently, Nosofsky and Stanton (2005) employed a variant of this approach with stimuli that were varied on two dimensions. They found exemplar processing, and their results are consistent with ours.

The trend in theory building is toward models with both exemplar and rule processes. Erickson and Kruschke (1998) advocate a neural-network model that, under appropriate circumstances, allows behavior to transition from rule-based to exemplar-based processing. Nosofsky, Palmeri, and McKinley (1994) advocate a similar dual-process model: Participants use rules, but may augment them with exemplars if they deem their performance on certain stimuli is too poor. Neither the Erickson and Kruschke nor the Nosofsky et al. model explicitly accounts for stimulus confusion. Yet, because both rely on both exemplar- and rule-based processing, they may be easily adapted to account for our results. In sum, our results provide support for the dual-processing approach as well as give an indication of how these two processes interact.
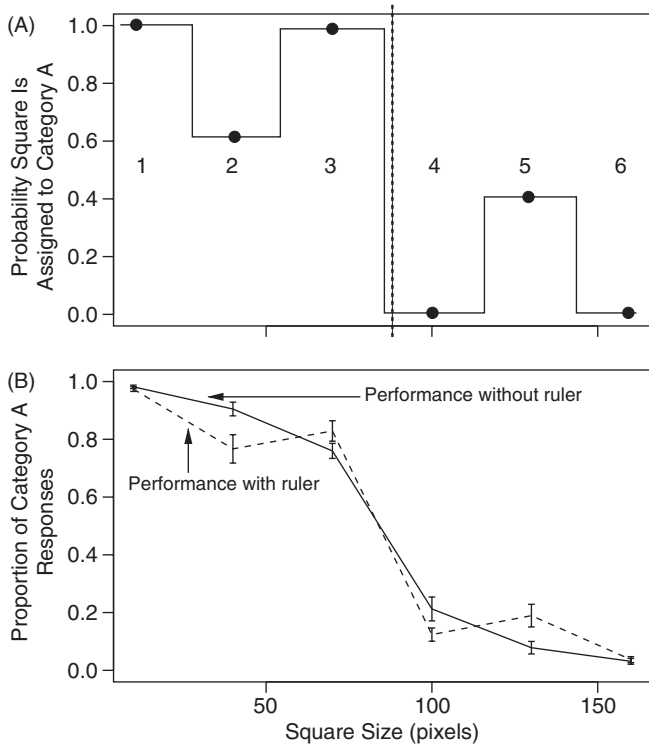
COMPLEX AND SIMPLE STIMULI

Our experiments are based on simple stimuli. One of the oldest problems in the study of human cognition is understanding why complex stimuli like faces, objects, and words are relatively easy to learn while simple stimuli, such as tones and line lengths, are difficult. The magnitude of this disparity is noteworthy. English speakers can learn about 100,000 words (Landauer, 1986); in contrast, even with extensive practice, people can learn no more than about seven well-spaced simple stimuli (Miller, 1956; cf. Rouder, Morey, Cowan, & Pfaltz, 2004). For simple stimuli, there are two factors affecting stimulus confusion: perceptual similarity and mnemonic capacity. If stimuli are too perceptually similar (they vary across too narrow a range), then surely they will be confused. Mnemonic capacity plays a role for simple stimuli too. Even when simple stimuli are perceptually distinct, identification performance is limited. This mnemonic-capacity limit does not hold for complex stimuli. With sufficient practice, a very large number of perceptually distinct complex stimuli may be identified.

Our main claim about stimulus confusion can be speculatively generalized: Categorization is based on exemplars only when exemplars may be reliably stored and maintained. If exemplars cannot be stored and maintained, then categorization is based on rules. For simple stimuli, such as line lengths or square sizes, stimulus confusion reflects both perceptual similarity and the mnemonic limit. Simple stimuli are reliably stored and maintained only when they are physically well separated and relatively few in number. If participants must identify either many simple stimuli or a set of simple stimuli that are perceptually similar, then they necessarily resort to rules. In contrast to simple stimuli, a large number of complex stimuli may be reliably stored and maintained. We claim that as long as complex stimuli are not perceptually confusable, categorization of them is mediated by exemplars. As most everyday cognition involves identifying distinct complex stimuli, most categorization is mediated by exemplars.

Ashby and Ell (2001) first noted that the number of to-be-categorized stimuli may be important determinant of categorization strategy. They posited that when the number of stimuli is small, participants' categorization may be mediated by exemplars. Ashby and Ell posited that when the number of stimuli exceeds capacity, categorization is handled by one of two rule-based modules rather than by exemplars. For simple stimuli, our theory is basically concordant with Ashby and Ell's. We disagree with them for complex stimuli and claim that with sufficient experience with perceptually distinct stimuli, participants' categorization is based on exemplars regardless of the number of stimuli.

We describe a dynamic trajectory of learning complex stimuli built on long-standing themes in the literature. In the initial stages of learning a domain, participants do not have sufficient knowledge of the relevant features and dimensions of the stimuli to store exemplars that are immune to forgetting. As a consequence, they may abstract simple rules about immediately obvious properties such as size, color, etc. In most cases, these simple rules are inadequate. Fortunately, with experience, participants develop familiarity with the features and dimensions of the domain. With this acquired knowledge, they may store rich exemplars that are immune to interference on a long-term basis. These rich exemplars serve as the basis for category construction and mediate the subsequent categorization of novel stimuli. A good example of a rich exemplar comes from the domain of face recognition. Certain features of faces are invariant to changes in the environment (environmental changes include lighting, angle of presentation, facial expression, hair style, and cosmetics). These features and their relations form a rich exemplar for a face and this exemplar may be associated with other information such as the person's name or gender. The categorization task of determining the gender of a novel person by that person's face can be accomplished by comparing that face to stored male and female faces. In sum, for complex stimuli, rules may be used initially, but with experience, participants use exemplars. The developmental component of this theory—that the features relevant for processing stimuli in a domain must be learned with experience—is characteristic of Gibson and Gibson (1955).

The developmental trajectory from rules to exemplars may be evident in how people learn to read. Although proficient readers recognize words as wholes and often without recourse to phonetic rules, they certainly used the phonetic rules as novices. Evidence for the importance of using phonetic rules in learning to read comes from Allen (1986), who found that the reading proficiency of hearing-impaired graduating high-school students was comparable to that of normal-hearing third graders. It is plausible that the lack of familiarity with phonetic rules prevented hearing-impaired readers from developing sufficiently rich word-level exemplars needed for proficient reading. The implication for development is that while it may be necessary to teach rules to initiate learning, mastery of a domain may have little to do with using these rules.

FUTURE DIRECTIONS

There are a number of interrelated questions that need to be addressed in any theory of category representation and learning. First, although the current research shows how categorization of simple stimuli depends on stimulus confusion, comparable paradigms for complex stimuli have yet to be developed. Appropriate research is needed to test our speculations about the role of exemplars in categorizing complex stimuli such as faces and words. Second, theories of mental representation need to explain the difference in performance with simple stimuli and more complex ones such as words and objects. For example, face-recognition theories need to explain both face recognition and the inability to recognize more than a handful of lines of

different lengths. Few pattern-identification theories also account for poor performance with simple stimuli (c.f., Biederman, 1987; Shiffrin & Nosofsky, 1994). Third, theories of categorization need a developmental mechanism that explains how exemplars or rules come about over the long term (e.g., Erickson & Kruschke, 1998; Seidenberg & McClelland, 1989). Our research holds out the possibility that rules used initially in categorization may be used less with increased experience.

Recommended Reading

Ashby, F.G., & Ell, S.W. (2001). (See References)
Erickson, M.A., & Kruschke, J.K. (1998). (See References)
Rouder, J.N., & Ratcliff, R. (2004). (See References)

REFERENCES

Allen, T.E. (1986). Patterns of academic achievement among hearing impaired students: 1974 and 1983. In A.N. Schildroth & M.E. Karchmer (Eds.), *Deaf children in America* (pp. 161–206). San Diego, CA: College Hill Press.

Ashby, F.G., & Ell, S.W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, *5*, 204–210.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.

Erickson, M.A., & Kruschke, J.K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.

Gibson, J.J., & Gibson, E.J. (1955). Perceptual learning: Differentiation or enrichment? *Psychological Review*, *62*, 32–41.

Landauer, T.K. (1986). How much do people remember? Some estimates of the quantity of learned information in long-term memory. *Cognitive Science*, *10*, 477–493.

Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Miller, G.A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.

Nosofsky, R.M., Palmeri, T.J., & McKinley, S.C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*, 53–79.

Nosofsky, R.M., & Stanton, R.D. (2005). Speeded classification in a probabilistic category structure: Contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 508–629.

Rouder, J.N., Morey, R.D., Cowan, N., & Pfaltz, M. (2004). Learning in a unidimensional absolute identification task. *Psychonomic Bulletin & Review*, *11*, 932–938.

Rouder, J.N., & Ratcliff, R. (2004). Comparing models of categorization. *Journal of Experimental Psychology: General*, *133*, 63–82.

Seidenberg, M.S., & McClelland, J.M. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.

Shiffrin, R.M., & Nosofsky, R.M. (1994). Seven plus or minus two: A commentary on capacity limitations. *Psychological Review*, *101*, 357–361.

Trabasso, T., & Bower, G.H. (1968). *Attention in learning: Theory and research*. New York: Wiley.