

Knowledge partitioning in categorization: Boundary conditions

STEPHAN LEWANDOWSKY and LEO ROBERTS
University of Western Australia, Crawley, Australia

and

LEE-XIENG YANG
National Chung Cheng University, Chia-Yi, Taiwan

Knowledge partitioning refers to the notion that knowledge can be held in independent and non-overlapping parcels. Partitioned knowledge may cause people to make contradictory decisions for identical problems in different circumstances. We report two experiments that explored the boundary conditions of knowledge partitioning in categorization. The studies examined whether or not people would partition their knowledge (1) when categorization rules were or were not verbalizable and (2) when the to-be-categorized stimuli comprised perceptually separable or integral dimensions. When learning difficulty was controlled, partitioning occurred across all combinations of verbalizability and integrality/separability, underscoring the generality of knowledge partitioning. Partitioning was absent only when the task was rapidly learned and people reached a high level of proficiency, suggesting that task difficulty plays a critical role in the emergence of partitioned knowledge.

The knowledge partitioning framework is based on the idea that knowledge, rather than being represented as an integrated whole, may be fractionated into independent parcels that are used selectively without reference to knowledge held in other parcels (see, e.g., Lewandowsky, Kalish, & Ngang, 2002; Lewandowsky & Kirsner, 2000; Yang & Lewandowsky, 2003). As a consequence, people may provide contradictory answers to an identical problem, depending on which knowledge parcel they use to guide their answer.

Knowledge partitioning has been demonstrated with experts (Lewandowsky & Kirsner, 2000), with nonexpert participants in a function learning paradigm (Lewandowsky et al., 2002), and with nonexpert participants in categorization tasks involving numeric (Yang & Lewandowsky, 2003) and perceptual stimuli (Yang & Lewandowsky, 2004). In each instance, a normatively irrelevant context cue (such as a verbal label or stimulus color) served as the basis for the contradictory resolution of an identical problem. This contradiction, elicited by the mere switch of the context cue without alteration of the surface structure of the problem, distinguishes partitioning from other well-known context effects (see Kalish, Lewandowsky, & Kruschke, 2004).

To illustrate, consider the study by Lewandowsky and Kirsner (2000) in which expert firefighters had to predict the spread of simulated wildfires. Firefighters were

provided with the physical predictors of fire spread in two different contexts: *back-burn* or *to-be-controlled*. (A *back-burn* is a fire lit by fire fighters in the path of an advancing to-be-controlled fire to starve it of fuel.) In reality, any fire will preferentially spread with the wind and up slopes. It follows that if the two variables are placed in opposition, the direction of a fire is not immediately obvious: With light downhill winds, slope may be the overriding variable, and the fire may spread uphill. Conversely, if the opposing wind is sufficiently strong, a fire may spread with the wind and hence downhill. Whether the fire is a back-burn or to-be-controlled is physically irrelevant; however, Lewandowsky and Kirsner found that experts expected an identical fire to spread uphill against the wind if presented as a back-burn, and downhill with the wind if presented as a to-be-controlled fire. In contrast to the widely accepted view that expertise is highly integrated (e.g., Bédard & Chi, 1992; Glaser, 1996; see Lewandowsky et al., 2002, for a review of the integration assumption), this finding suggests that experts possess independent parcels of knowledge that are selectively accessed on the basis of context and without regard to knowledge in other parcels. Although the apparent partitioning of knowledge necessarily causes errors in certain situations, the firefighters' context-bound behavior is nonetheless statistically appropriate: Because back-burns are lit only when winds are light, slope tends to be particularly relevant. In contrast, most to-be-controlled fires occur when the wind is strong, making wind particularly salient in that context.

Although the results with experts were suggestive of partitioning, stronger support for the idea that people can acquire multiple and mutually independent parcels of

Preparation of this article was facilitated by a Discovery grant from the Australian Research Council to the first author and Mike Kalish. Please address correspondence to S. Lewandowsky, School of Psychology, University of Western Australia, Crawley, WA 6009, Australia (e-mail, lewan@psy.uwa.edu.au; personal home page, www.psy.uwa.edu.au/user/lewan/).

knowledge for a single task was obtained in the laboratory using nonexperts. Lewandowsky et al. (2002) asked participants to learn a U-shaped (quadratic) relationship between a continuous stimulus dimension and a continuous response variable. The function embodied an abstraction of the relationship just discussed between wind speed and the speed of fire spread on a slope. On each training trial, the stimulus value was accompanied by one of two context labels, each of which was preferentially associated with one half of the function (the ascending or descending component of the quadratic). For the descending part of the function, most stimuli were identified by the context label “back-burn,” reflecting the fact that back-burns involve light winds and spread uphill, and that increasing (downhill) wind speed slows their (uphill) progress. For the ascending part of the function, the likely context label was “to-be-controlled,” reflecting the fact that most to-be-controlled fires occur when winds are stronger and spread increasingly quickly as the wind picks up. The vertex of the function represented the point at which opposing wind and slope were balanced, thus minimizing speed of spread. After learning this function in the two contexts, people on a transfer test not only used the context label to guide their extrapolation, but their performance within each context closely resembled the transfer performance of people in two control conditions who had only learned one half of the function. Specifically, people who had only learned the descending (or ascending) half of the function transferred to the ascending (or descending) half in the same way as people who had learned both halves of the function but in different contexts. This pattern implies that when presented with a stimulus in a new context, participants’ performance was not influenced by their acquired knowledge in the other context, further strengthening the case for the presence of mutually exclusive knowledge parcels.

Yang and Lewandowsky (2003, 2004) extended the knowledge partitioning framework to categorization. In their experiments, participants were provided with a context cue, instantiated either as an arbitrary verbal label or as the color of the perceptual stimulus, and the dimensional values of stimuli were selected from within a two-dimensional category space. In the study of greatest interest (Yang & Lewandowsky, 2004), dimensional values were represented graphically by the height of a rectangle and the horizontal offset of a vertical bar shown within the rectangle. Because the category space used by Yang and Lewandowsky also underlies the present experiments, it deserves to be explained in some detail (see Figure 1).

The space created by the two dimensions (arbitrarily labeled x and y in Figure 1) was divided into two categories by two parallel partial boundaries. Stimuli within the two boundaries (Area 2 in Figure 1) belonged to Category A, whereas stimuli outside the two boundaries (Areas 1 and 3) belonged to Category B. During training, stimuli that clustered around the lower category boundary (termed the *lower cluster* and represented in Figure 1 by the lower dashed oval) were presented in one arbitrary context (in this case, the stimuli were shown in one color), and stimuli




Figure 1. The category structure used by Yang and Lewandowsky (2003, 2004) and in Experiments 1 and 2A. Items within the dashed ovals are those presented during training, whereas items outside the ovals were presented only at transfer. Items in Area 2 belong to Category A, whereas items in Areas 1 and 3 belong to Category B. Two sample stimuli are shown below the stimulus space to illustrate the nature of the task.

clustered around the upper boundary (the *upper cluster*, represented by the upper dashed oval) were presented in another context (i.e., a different color). Because half of the stimuli in each cluster belonged to each category, context by itself did not predict category membership. However, context did predict which of the two partial boundaries correctly classified each stimulus. At transfer, participants were found to use the boundary that was identified by the stimulus context, even if the stimulus was drawn from the far side of the *other* category boundary. For example, the sample stimulus on the right-hand side in Figure 1 would have been classified as belonging to Category B if presented in the color of the training stimuli in the lower cluster, whereas it would have been assigned to Category A if shown in the other color. In general, stimuli at identical locations in the two-dimensional space were classified differently, depending on a normatively irrelevant context cue. These results suggest that people sometimes master a complex categorization task by partitioning it into multiple partial and independent decision rules.¹

However, the generality of the results by Yang and Lewandowsky (2003, 2004) within the realm of categorization remains to be established. Yang and Lewandowsky’s (2004) stimuli were characterized by two properties that are linked to two major current issues in categorization research. First, the stimuli could arguably be classified by a verbalizable rule. For example, the rule “Respond A if the rectangle is taller than the offset of the vertical bar on the left; respond B otherwise” (Yang & Lewandowsky, 2004, p. 1057) could be used to describe the lower bound-

ary of the category space. There is considerable evidence that some, but not all, categorization stimuli are amenable to classification by a verbalizable rule, and that the two classes of stimuli are often processed differently (see, e.g., Ashby, Ell, & Waldron, 2003; Ashby, Queller, & Berretty, 1999; Maddox, Ashby, & Bohil, 2003; Waldron & Ashby, 2001). Second, Yang and Lewandowsky's (2004) stimuli consisted of perceptually separable (rather than integral) dimensions (e.g., Garner, 1974; Garner & Felfoldy, 1970; Goldstone, 1994; Nosofsky, 1986). There is considerable evidence that processing differs greatly between integral stimuli, such as colors varying in saturation and brightness, and separable stimuli, such as circles varying in diameter and angle of an embedded radial line. For example, separable dimensions can be analyzed individually without interference from another separable dimension (e.g., Garner & Felfoldy, 1970), and people can selectively focus their attention on one dimension with greater efficiency when categorization stimuli are composed of separable rather than integral dimensions (e.g., McKinley & Nosofsky, 1996). It follows that knowledge partitioning in categorization may be affected by the two key variables of whether or not a rule can be verbalized and whether stimulus dimensions are separable or integral. We now consider those variables in more detail before reporting two experiments that examined their role with respect to knowledge partitioning.

Verbalizability of Categorization Rules

Turning to the first issue, the verbalizability of a rule, Ashby, Alfonso-Reese, Turken, and Waldron (1998) suggested that different systems underlie the learning of different category structures. Specifically, Ashby et al. (1998) argued for the existence of two separate categorization systems: an explicit rule-based system, which operates when category rules are easily verbalized, and an implicit information integration system, which operates when category rules are difficult or impossible to verbalize.

Ashby et al. (1998) suggested that verbalization of a two-dimensional rule (e.g., a rule based on the diagonal boundary $y = x$) requires both stimulus dimensions to be expressed and measured in the same units, so that an instant dimensional comparison can be made. Thus, whereas comparing the height of a rectangle and the offset of a vertical bar within it is meaningful (because both are measured as distances), comparing (for example) the hue and saturation of a color is not. Therefore, for category spaces with two or more dimensions, the verbalizability of the rule emerges as a property of the stimulus dimensions.

Evidence for the existence of separate verbal and nonverbal systems of categorization involves a variety of dissociations. For example, verbal rules can be learned even when feedback during training is removed (Ashby et al., 1999) or delayed (Maddox et al., 2003), whereas nonverbal information integration tasks are disrupted by those manipulations. Similarly, unlike information integration tasks, rule-based responding is unaffected when response keys or response hands are switched following training

(Ashby et al., 2003). Conversely, the learning of verbal rules, but not of information integration tasks, is impaired by a secondary task (Waldron & Ashby, 2001) as well as by limiting the time to process feedback (Maddox, Ashby, Ing, & Pickering, 2004). These dissociations raise the possibility that knowledge partitioning, which has so far only occurred with verbalizable stimuli, may not be detectable in information integration tasks.

Separability Versus Integrality of Dimensions

The dimensional variable of separability versus integrality, which we will call *dimensional interaction* from here on, affects categorization in several fundamental ways. For example, dimensional interaction affects the observed relationship between identification and categorization performance. Shepard and Chang (1963) discovered that the learning of various category structures could be predicted on the basis of confusions in stimulus identification performance when the to-be-categorized stimuli were integral, but not when they were separable. Shepard, Hovland, and Jenkins (1961) suggested that the lack of generalization from identification to classification for separable dimensions may arise from participants' focusing their attention during categorization on a relevant dimension at the expense of other dimensions. Because integral stimuli, as a result of their unanalyzable nature, are spared from selective dimensional attention during categorization, their classification can be predicted from identification performance.

Since Shepard's seminal work, the concept of selective attention has become increasingly important in models of categorization (see, e.g., Kruschke, 1992; Nosofsky, 1986; Nosofsky & Johansen, 2000), and it has repeatedly been shown to be a powerful explanatory concept (e.g., McKinley & Nosofsky, 1996; for a recent review, see Kruschke, 2005). The way in which dimensional interaction modulates the role of selective attention is particularly relevant here. For example, Nosofsky (1989) showed that when stimulus dimensions are separable, and only one dimension is relevant to the categorization decision, people learn selectively to attend to the relevant dimension at the expense of the irrelevant one. When stimulus dimensions are integral, by contrast, people are considerably less able to devote attention selectively to one dimension (McKinley & Nosofsky, 1996).

One implication of this reduced ability to differentiate between integral dimensions is that in some cases, people may categorize stimuli not on the basis of experimenter-intended stimulus dimensions, but on the basis of some other property that is "emerging" from the integral stimuli. For example, rectangles that differ in height and width—two integral dimensions—may be classified on the basis of their "tallness" or "thinness," expressed as the ratio of the two dimensions (see, e.g., Monahan & Lockhead, 1977). Given that knowledge partitioning relies on participants' recognition of how dimensions interact—specifically, that context gates the nature of the relationship between the remaining two dimensions—it follows that knowledge

partitioning, known to occur with separable stimuli, may not be detectable with integral stimuli.

EXPERIMENT 1

In Experiment 1, we examined the presence of knowledge partitioning in categorization when categorization rules were presumed to be verbal or nonverbal and, orthogonally, when stimulus dimensions were perceptually separable or integral.² All conditions used the category structure of Yang and Lewandowsky (2003, 2004).

Method

Design and Participants

The experiment employed a 2 × 2 between-subjects design with verbalizability of task (verbal rule vs. information integration) and dimensional interaction (integral vs. separable) as the two independent variables.

Eighty-one volunteers participated in exchange for course credit and were randomly assigned to conditions.

Stimuli and Apparatus

All stimuli were generated and displayed on a PC-compatible computer, which also recorded all responses. Figure 2 shows sample stimuli for all four conditions.

Stimuli were assigned to levels of the dimensional interaction variable on the basis of empirical precedent.³ Circles with a radial line, in which the diameter of the circle and the orientation of the line map onto the two dimensions, have been used as examples of separable stimuli for more than 40 years (see Maddox & Ashby, 1998, for a brief overview). Likewise, rectangles with a vertical bar on the lower horizontal border, in which rectangle height and horizontal offset of the bar map onto the two dimensions, were identified as separable by Erickson and Kruschke (1998). Rectangles varying in height and width were confirmed to be integral by Dunn (1983) using a variety of measures and analyses. Finally, the top right-hand cell of Figure 2 shows an example of Fourier descriptors, constructed by varying amplitude and phase of superimposed sine waves. Those stimuli were confirmed to be integral by Op de Beeck, Wagemans, and Vogels (2003) in an elegantly designed visual-search experiment (see also Cortese & Dyre, 1996).

On the basis of the criteria postulated by Ashby et al. (1998), the rectangles varying in height and width and the rectangles with a vertical bar were considered to be verbalizable because the dimensions were commensurate (i.e., both dimensions were expressed by linear

extent). By the same criterion, the circles with a radial line at varying angles were considered to be nonverbalizable because the units of the two dimensions (angle vs. size) remained noninterchangeable. Likewise, the Fourier descriptors were considered nonverbalizable because the two dimensions, amplitude and phase angle, were also expressed in incommensurate units.

Category Space

All stimuli were drawn from the two-dimensional category structure designed by Yang and Lewandowsky (2003) shown in Figure 1. During the training phase, 40 stimuli were sampled from random locations within the training space for each participant, with the constraint that half of the stimuli were located above and half below the $y = x$ diagonal.

The two dimensions of the category space corresponded to the dimensions of the stimuli in each condition (e.g., height and width of a rectangle in the integral–verbal condition, and so on). Correct classification could be achieved on the basis of x and y alone. A third dimension, color of the stimulus, acted as the context variable in all conditions. During training, the 20 stimuli in the upper cluster (upper dashed oval in Figure 1) were consistently presented in one context (e.g., red), whereas the 20 stimuli in the lower cluster (lower dashed oval) were presented in the other context (e.g., green). Assignment of color to cluster was counterbalanced across participants. Because half of the stimuli in each cluster belonged to each category, context signaled cluster membership (i.e., upper or lower) but not category membership.

Training consisted of eight blocks of trials, each involving a random sequence of the 40 training items. Training was followed by a transfer phase that consisted of two blocks of 60 stimuli, presented in a different random sequence in each block. Of the 60 transfer stimuli, 40 were novel (i.e., any point in Figure 1 outside the two ovals), and 20 had been presented during training. Novel transfer items were the same for all participants. Context was switched across the two transfer blocks so that each stimulus appeared once in each context.

Procedure

Participants completed the task individually in a quiet testing booth. They were given the opportunity to take self-paced breaks at the end of each block of trials.

Each trial commenced with a “+” sign that appeared in the center of the screen for 500 msec. The stimulus was then presented centrally, and participants indicated their classification via keypress. The stimulus remained on the screen until a response was made. During training, participants received feedback following category selection via the word *correct* or *wrong* appearing on the screen for 1,000 msec. Feedback was withheld during transfer.

Following completion of the session, participants were asked to describe any rules they had used to classify the stimuli.

Results and Discussion

For all analyses, a participant’s responses were included only if the individual’s performance in the final training block significantly exceeded chance ($\alpha = .05$), which corresponded to a proportion correct of .65 (by a binomial distribution for $n = 40$ items and $p = .5$). By that criterion, the data of 10 participants were excluded from the analysis, so that data from 71 participants were retained (22 in the integral–verbal condition, 16 in each of the nonverbal conditions, and 17 in the separable–verbal condition).

Verbalizability of Rules

To examine the effectiveness of the verbalizability manipulation, the self-reported rules provided by participants




Figure 2. Sample stimuli used in Experiment 1.

upon completion of the experiment were sorted into categories by three independent judges. Judges were instructed to maximize the similarity of rules within each category and to create any number of categories required to do so.

Verbalizability was measured by examining the consistency of reported rules across participants within each condition. If stimuli were classified by a simple verbal rule, then all participants could be expected to report using that rule. For this case of maximal consistency, judges would be expected to cluster all (or nearly all) responses into a single category. Conversely, if verbalization were difficult, consistency would be expected to be minimal, and in the extreme case every participant's rule might be judged to belong to its own unique category. Accordingly, consistency was measured as the proportion of the number of categories created by the judges relative to the number of participants who provided responses ($n = 28$ for the nonverbal conditions and $n = 30$ for the verbal conditions). Small values of this proportion reflect high consistency (hence, high verbalizability), whereas large values of the proportion reflect low consistency (hence, low verbalizability). The consistency proportion was .31 for the verbalizable conditions and .58 for the nonverbalizable conditions, with the difference found to be significant by Fisher's exact test ($p < .05$; the pattern of differences was unchanged if the separable and integral conditions were compared separately). We conclude that the verbalizability manipulation was successful.

Training Performance

Figure 3 shows training performance for all conditions. A 2 (dimensional interaction: separable vs. integral) \times 2 (verbalizability of task: verbal vs. nonverbal) \times 8 (block) ANOVA with between- and within-subjects factors revealed a main effect of block [$F(7,469) = 86.46$, $MS_e = 0.01$, $p < .0001$], a main effect of dimensional interaction [$F(1,67) = 13.02$, $MS_e = 0.05$, $p < .0006$], a marginally

significant main effect of verbalizability [$F(1,67) = 3.39$, $MS_e = 0.05$, $p = .07$], and an interaction between those between-subjects variables [$F(1,67) = 17.29$, $MS_e = 0.05$, $p < .0001$].

Figure 3 suggests that the interaction resulted from the strikingly better training performance in the integral-verbal condition than in the remaining three conditions. Pairwise comparisons confirmed this suggestion, for the integral-verbal condition differed significantly from the separable-verbal [$F(1,37) = 35.39$, $MS_e = 0.05$, $p < .0001$], separable-nonverbal [$F(1,36) = 14.40$, $MS_e = 0.05$, $p < .0005$], and integral-nonverbal [$F(1,36) = 18.59$, $MS_e = 0.01$, $p < .0001$] conditions.

Transfer Performance

All transfer results were analyzed using the proportion of A responses as the dependent variable. Following Yang and Lewandowsky (2003, 2004), the category structure was separated into three diagnostic areas to facilitate identification of knowledge partitioning (see Figure 1). Average classification probabilities were computed separately for items within each area. Separate analyses were conducted for novel items and training items presented at transfer.

Training items were presented twice at transfer: once in a context that was congruent with the one experienced at training, and once in the incongruent context. Because responses to training items might involve memorization, thus reducing diagnosticity, training items were analyzed at a coarse level via a 2 (context: congruent vs. incongruent) \times 3 (area) within-subjects ANOVA that aggregated across all conditions simultaneously, rather than using an omnibus ANOVA consistent with the study's 2 \times 2 design.⁴ This analysis revealed main effects of context [$F(1,70) = 8.92$, $MS_e = 0.08$, $p < .005$] and area [$F(2,140) = 82.51$, $MS_e = 0.12$, $p < .0001$], as well as an interaction of both variables [$F(2,140) = 6.89$, $MS_e = 0.05$, $p < .002$].

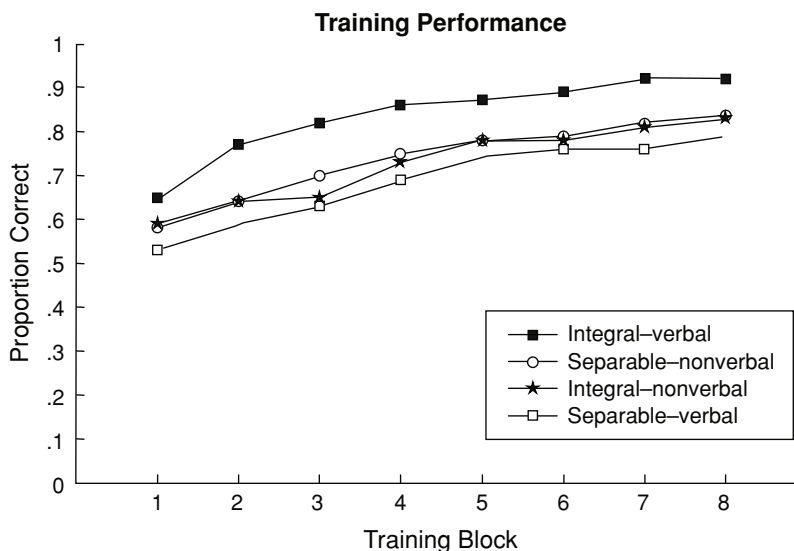


Figure 3. Training performance for all conditions in Experiment 1.

Table 1 shows the pattern of responses for all diagnostic areas. It is evident that participants learned to categorize the stimuli, because items in Area 2 (the part of the space devoted to Category A) were more likely to be classified as members of Category A than were items in Areas 1 and 3. Moreover, the extent to which people relied on the bilinear boundary was attenuated for training items presented in the incongruent context, as reflected by the increased tendency to place such items into Category A when they were from Area 1 or 3. This attenuation is not unexpected, given that training items presented in the incongruent context, strictly speaking, were no longer training items.

Turning to the more diagnostic novel transfer items, knowledge partitioning would be present if responses differed according to context for items from Areas 1 and 3, and in opposing directions. Specifically, the probability of classifying an item in Area 1 as belonging to Category A should be greater when the item is presented in the context associated with the lower area than when it is presented in the context associated with the upper area, whereas for an item located in Area 3, the reverse pattern should occur (i.e., the proportion of A responses should be greater when the item is presented in the context for the upper rather than the lower area). Figure 4 shows that this pattern was obtained in three of the four conditions. The one exception was the integral-verbal condition (i.e., rectangles of varying height and width), in which context had no effect on responding in any of the areas. The absence of a context effect, in conjunction with the clear differences in responding between Area 2 and Areas 1 and 3, suggests that people in the integral-verbal condition relied on the bilinear category boundaries.

Statistical confirmation of this pattern was provided by an omnibus 2 (dimensional interaction) \times 2 (verbalizability of task) \times 2 (context: upper vs. lower) \times 3 (area) between-within ANOVA on the novel transfer items. The ANOVA revealed no main effects of verbalizability or dimensional interaction [both $F(1,67) < 1$], but there was an interaction between the two between-subjects variables [$F(1,67) = 4.25, p < .05$]. In addition, the main effects of context [$F(1,67) = 8.00, MS_e = 0.02, p < .01$] and area [$F(2,134) = 172.9, MS_e = 0.04, p < .0001$] were significant, as was the context \times area interaction [$F(2,134) = 22.64, MS_e = 0.04, p < .0001$]. Area also interacted with both between-subjects variables severally and jointly [largest $F(2,134) = 22.10$]; because those effects merely indicate that absolute performance level—expressed as the

difference between Area 2 and the other areas—differed between conditions, they are not explored further. Interactions involving the between-subjects variables and context failed to reach conventional levels of significance [largest $F(1,67) = 3.22, p < .10$].

The main effect of area was due to items in Area 2 being classified as belonging to Category A more often (.68) than those in Areas 1 and 3 combined (.27), which is indicative of people on the whole being sensitive to the trained category structure. This overall sensitivity was qualified by the context \times area interaction, which, as suggested by Figure 4, confirms that knowledge partitioning occurred in the experiment as a whole.

None of the higher order interactions involving context \times area were significant, including the overarching verbalizability \times dimensional interaction \times context \times area interaction [$F(2,134) = 1.93, p > .05$]. Notwithstanding this result, planned comparisons were performed within each condition to further investigate the crucial area \times context interaction. Significant interactions were present in the integral-nonverbal [$F(2,30) = 5.30, MS_e = 0.06, p \approx .01$], separable-verbal [$F(2,32) = 7.32, MS_e = 0.05, p < .005$], and separable-nonverbal [$F(2,30) = 7.75, MS_e = 0.04, p < .002$] conditions, but not in the integral-verbal condition [$F(2,42) < 1$]. These results confirm the pattern apparent from Figure 4, that knowledge partitioning occurred in all conditions except the integral-verbal condition.

To confirm that the significant interactions resulted exclusively from the opposing effects of context in Areas 1 and 3, the analysis was repeated with the items from Area 2 removed. The 2 (context) \times 2 (area) within-subjects ANOVAs yielded significant interactions in the integral-nonverbal [$F(1,15) = 5.48, MS_e = 0.12, p < .05$], separable-nonverbal [$F(1,15) = 9.03, MS_e = 0.08, p < .01$], and separable-verbal [$F(1,16) = 7.16, MS_e = 0.09, p < .02$] conditions, confirming the presence of knowledge partitioning in these conditions. For the integral-verbal condition, the interaction was not significant [$F(1,21) < 1$].

One potential limitation of the observed partitioning is that the magnitude of the context effect in Areas 1 and 3 falls short of what would be expected if all participants consistently partitioned their knowledge, in which case proportions should be close to 0 or 1 across the two contexts. Instead of being near unity, however, the observed differences between contexts hovered around the .2–.3 range, suggesting that the partial boundaries were applied only some of the time or by some of the participants. This issue was resolved by an individual-differences analysis, which we report after the data from all experiments have been presented.

Summary

The results from Experiment 1 are readily summarized: The transfer data indicate that people partitioned their knowledge in three of the four conditions, although the extent of that partitioning across individuals remains to be determined. The one condition in which partitioning was absent involved integral and verbalizable stimuli.

Table 1
Mean Probability of Category A Responses for Training Items Presented at Transfer in Each Context and Each Area in Experiment 1

Context	Area		
	1	2	3
Congruent	.26	.83	.39
Incongruent	.36	.81	.56

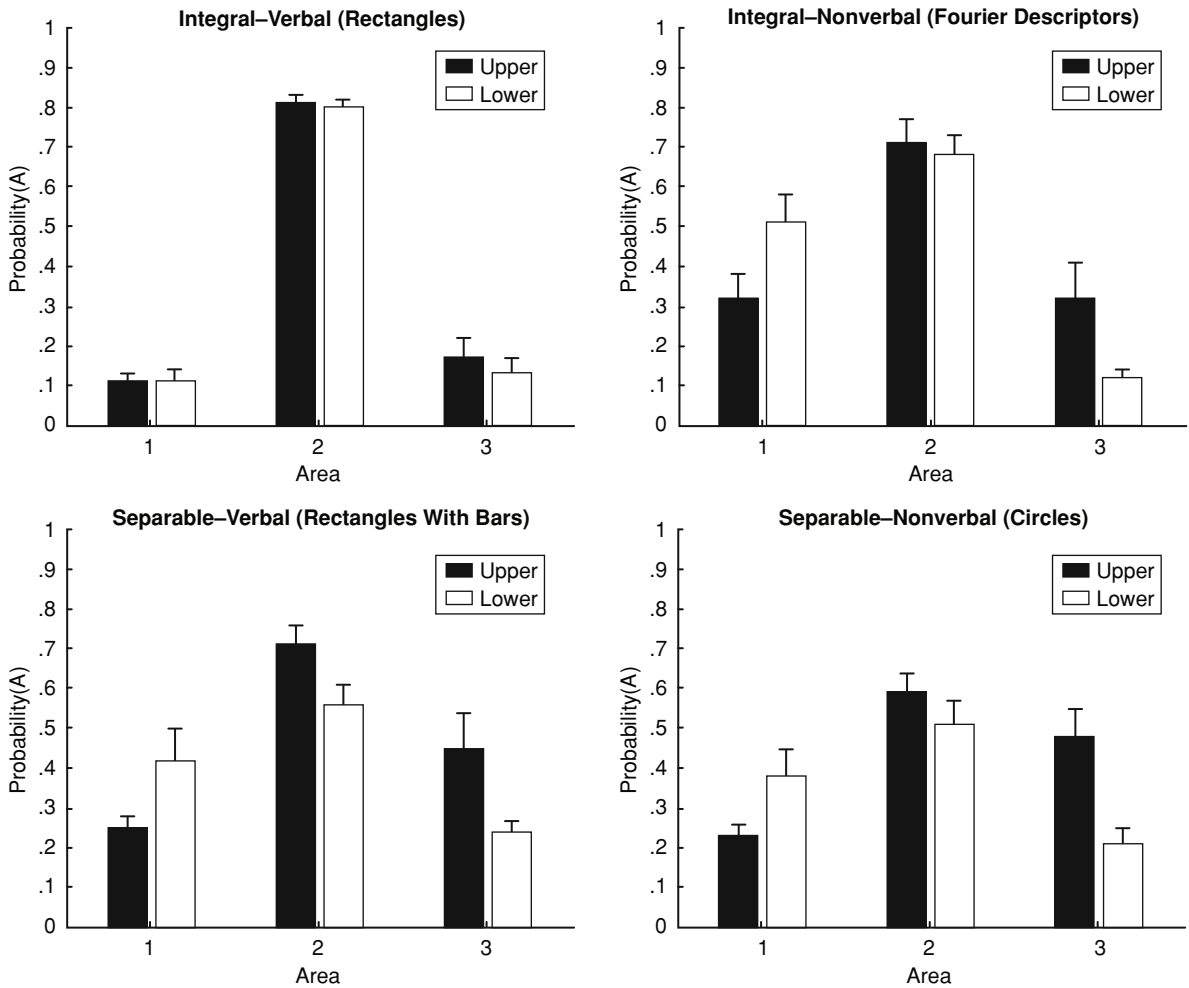


Figure 4. Mean probabilities of Category A responses in Experiment 1 for novel transfer items in each area and test context (upper vs. lower) for all conditions. See Figure 1 for an explanation of the areas.

Before we seek a possible explanation for why this particular conjunction of stimulus properties prevented partitioning, we focus on the fact that training performance in this condition was better than in the remaining three conditions by a considerable margin (some 10% or more). Taking performance as an index of difficulty, the integral-verbal task may have been so easy that people did not have to partition the category structure. By contrast, the difficulty of the remaining three tasks may have forced people to resort to forming partial context-gated decision boundaries. Experiments 2A and 2B explored this possibility.

EXPERIMENT 2

In Experiment 2, we sought to make the integral-verbal task more difficult without altering the integral nature of the stimulus or the verbalizability of the rule. To this end, the rectangle stimuli were rotated slightly off the vertical without altering any of the remaining task parameters. According to the pervasive oblique effect, performance

on many tasks, including identification of lines (see, e.g., Essock, 1980; Lasaga & Garner, 1983), is inferior when stimuli are presented at an angle (i.e., in an oblique orientation) rather than vertically or horizontally.

It follows that categorization of rectangles, which requires comparison of their height and width, might also be impaired if the stimuli are presented at an oblique angle. For Experiment 2A, the category space of the first study was used with rotated stimuli, whereas in Experiment 2B we additionally rotated the partial boundaries by 90°, thus reversing their slope from positive to negative. Negatively sloped boundaries are known to increase task difficulty in comparison with positive boundaries (see, e.g., Lewandowsky et al., 2002; Yang & Lewandowsky, 2003).

Method

Participants

Seventeen volunteers participated in Experiment 2A, and another 21 volunteers took part in Experiment 2B. These participants either received course credit or were reimbursed \$10 for travel expenses.

Stimuli, Design, and Procedure

The experiment replicated the integral–verbal condition of Experiment 1. In Experiment 2A, the only departure from the method of the preceding experiment was that the rectangles were presented with a counterclockwise rotation of 10° about the bottom left corner. For Experiment 2B, we also changed the category space. Whereas previously the upper and lower category boundaries could be expressed by the function $|y - x - 100| = 200$, in Experiment 2B the boundaries were $|(900 - y) - x - 100| = 200$. Because the nature of the stimuli was unchanged, with the dimensions remaining commensurate, there was no reason to expect verbalizability to be affected by the rotation of boundaries.

Results

Experiment 2A

As in Experiment 1, participants' responses were analyzed only if performance in the final training block surpassed chance. The data of 1 participant were removed on the basis of this criterion, leaving 16 participants for analysis.

Training performance. Training performance is illustrated in Figure 5 and is shown together with results from the integral–verbal condition in Experiment 1 for comparison (the third line in the figure corresponds to Experiment 2B). A one-way within-subjects ANOVA indicated a significant improvement across the eight blocks [$F(7,105) = 14.11$, $MS_e = 0.005$, $p < .0001$]. A planned 2 (condition) \times 8 (block) between–within ANOVA was performed using the data from Experiment 2A and the integral–verbal condition of the first experiment to investigate whether task difficulty in Experiment 2A was enhanced as intended. Neither condition [$F(1,36) = 0.27$] nor the interaction between condition and block [$F(7,252) = 1.41$, $p > .10$] was found to be significant, suggesting that stimulus rotation alone left performance unchanged.

Transfer performance. For training items presented at transfer, a 2 (context: congruent vs. incongruent) \times 3 (area) within-subjects ANOVA revealed a main effect of area [$F(2,30) = 104.27$, $MS_e = 0.04$, $p < .0001$], along with a significant context \times area interaction [$F(2,30) = 3.80$, $MS_e = 0.05$, $p < .05$]. Table 2 shows classification probabilities in each area for training items presented at transfer in the two contexts (i.e., congruent and incongruent). As in Experiment 1, a trend was evident for the difference between responding in Area 2 and Areas 1 and 3 (which reflects sensitivity to the trained category structure) to be attenuated when items were presented in the incongruent context.

For the novel items, transfer responses in each area for each context are shown in Figure 6. Unlike in the integral–verbal condition of Experiment 1, there is a hint of a context effect in Areas 1 and 3 that may be indicative of partitioning, although the overall pattern appears more compatible with nonpartitioned learning of the bilinear boundaries.

A 2 (context: upper vs. lower) \times 3 (area) within-subjects ANOVA revealed a main effect of area [$F(2,30) = 291.04$, $MS_e = 0.01$, $p < .0001$] and a significant effect of context [$F(1,15) = 4.54$, $MS_e = 0.003$, $p = .05$], but no significant interaction [$F(2,30) = 2.13$, $p \approx .14$]. Thus, the slight hint of partitioning in Figure 6 escaped statistical confirmation.

Experiment 2B

As before, participants were excluded from analysis if their performance fell below chance on the last training block. The data of 8 participants were excluded by this criterion, suggesting a high level of task difficulty in this experiment.

Training performance. Figure 5 shows performance during training. The obvious improvement across the eight




Figure 5. Training performance in Experiments 2A and 2B and for the integral–verbal condition in Experiment 1.

Table 2
Mean Probability of Category A Responses for
Training Items Presented at Transfer in Each Context
and Each Area in Experiments 2A and 2B

Context	Area		
	1	2	3
Experiment 2A			
Congruent	.09	.93	.41
Incongruent	.19	.81	.59
Experiment 2B			
Congruent	.48	.83	.51
Incongruent	.65	.87	.78

training blocks was confirmed by a one-way within-subjects ANOVA [$F(7,84) = 18.37$, $MS_e = 0.01$, $p < .0001$].

Figure 5 also confirms that training performance was worse than in Experiment 2A, attesting to the success of the category space manipulation. A planned 2 (experiment) \times 8 (block) between-within ANOVA that compared performance in Experiments 2A and 2B established that the difference between studies was significant [$F(1,27) = 33.69$, $MS_e = 0.06$, $p < .0001$]. In addition, a significant experiment \times block interaction [$F(7,189) = 4.61$, $MS_e = 0.007$, $p < .0001$] reflected the fact that people commenced training with a much lower level of performance in this task than in Experiment 2A but achieved a comparable performance level at the end of training.

In addition, comparison of Figure 5 with the training data of Experiment 1 (Figure 3) suggests that performance in Experiment 2B was roughly equal to that in all conditions in the earlier experiment in which participants were found to partition their knowledge. If task difficulty, as indexed by training performance, is a determinant of whether or not knowledge is partitioned in categorization, then partitioning should be present in this experiment.

Transfer performance. People classified training items in roughly the same ways as in Experiments 1 and 2A (see Table 2). Nonetheless, the 2 (context: congruent vs. incongruent) \times 3 (area) within-subjects ANOVA yielded only a main effect of area [$F(2,24) = 7.46$, $MS_e = 0.08$, $p < .005$], with the previously observed context \times area interaction failing to reach significance here [$F(2,24) = 2.17$, $MS_e = 0.04$, $p > .10$].

For the novel items, Figure 7 shows that classification performance differed between contexts in Areas 1 and 3, suggesting partitioning. The 2 (context: upper vs. lower) \times 3 (area) within-subjects ANOVA revealed a main effect of area [$F(2,24) = 30.08$, $MS_e = 0.05$, $p < .0001$] and a significant context \times area interaction [$F(2,24) = 4.35$, $MS_e = 0.09$, $p < .05$], but no effect of context [$F(1,12) < 1$]. To assess the effect of greatest interest—the context \times area interaction in Areas 1 and 3—classification probabilities for Area 2 items were excluded from a follow-up analysis. The resulting 2 (context) \times 2 (area) ANOVA yielded a significant interaction [$F(1,12) = 4.66$, $MS_e = 0.17$, $p =$

.052], confirming that knowledge partitioning was present in Experiment 2B.

Individual Differences

Thus far, the data from all experiments have been considered only in the aggregate. Although this suffices for identifying the presence or absence of partitioning, Yang and Lewandowsky (2003, 2004) showed that participants in the conditions in which partitioning was present could be divided into two distinct subgroups: One group of people (roughly a third of participants) was found to partition their knowledge completely, whereas another (slightly larger) group of participants did not partition their knowledge at all. (The strategies of the remaining participants could not be identified.)

Accordingly, we examined individual differences by considering the individual response profiles of all 100 participants retained for the analyses across all experiments. Each participant's responses to the 40 novel transfer stimuli within each context were coded as 1 (Category A) or 0 (Category B). Using a Euclidean distance measure, the resultant response vectors were compared with three predefined cluster centroids that corresponded, respectively, to use of the true boundaries (i.e., the upper and lower partial boundaries, regardless of context), the upper partial boundary, or the lower partial boundary. For example, within the upper-boundary centroid, all items above the upper partial boundary were set to 0 (Category B), whereas all items below the upper partial boundary were set to 1 (Category A). Each participant was assigned to the closest cluster for each context. If a participant's response profile was equidistant to two or more centroids, the person was assigned to a fourth, "unrecognized" group.

The analysis identified 14 participants who were simultaneously classified into the lower-boundary cluster




Figure 6. Mean probabilities of Category A responses in Experiment 2A for novel transfer items in each area and test context (upper or lower). See Figure 1 for an explanation of the areas.




Figure 7. Mean probabilities of Category A responses in Experiment 2B for novel transfer items in each area and test context (upper or lower). Partition boundaries were now shifted to have negative slope, but the areas are functionally equivalent to those used in Experiments 1 and 2A.

in the lower context and the upper-boundary cluster in the upper context. Those participants formed the knowledge-partitioning group, representing a proportion of .14 of all participants. Additionally, 17 participants (.17), showed evidence of partial knowledge partitioning (i.e., they used the upper boundary in the upper context but disregarded the lower boundary in the lower context, or vice versa). Another 58 participants (.58) applied the true boundaries in both contexts. Those participants formed the true-boundary group, whereas the remaining 11 (.11) participants employed unrecognized strategies.

Figure 8 illustrates the differences in strategy between the knowledge-partitioning and the true-boundary groups. The top two panels show the responses of the knowledge-partitioning group in Experiment 2B ($n = 3$), and the middle panels show the responses of the knowledge-partitioning group in Experiments 1 and 2A ($n = 11$). In both sets of panels, the probability of an A response in the upper context progressively increases with an item's distance from the upper partial boundary. Conversely, in the lower context, the probability of an A response progressively increases with an item's distance from the lower partial boundary. The bottom panels show the true-boundary group in Experiments 1 and 2A ($n = 53$; Experiment 2B is not shown because it adds no additional information). It is clear that the probability of an A response is consistently greatest in the central portion of the space and that responding is unaffected by context.

The figure confirms that individuals who were identified as having partitioned their knowledge completely reversed their categorization judgments between contexts in Areas 1 and 3. The proportion of participants who fully partitioned their knowledge was slightly lower than the proportions observed by Yang and Lewandowsky (2003, 2004).

GENERAL DISCUSSION

Summary of Findings

Our objective was to determine whether knowledge partitioning in categorization is affected by the dimensional interaction of stimuli and the verbalizability of the task. Taken together, the two experiments converge on the conclusion that those variables have no effect on the occurrence of knowledge partitioning. Unless learning was extremely rapid, we observed partitioning with all types of stimuli. As in previous studies, partitioning was complete for a significant proportion of participants, and it was incomplete or absent in the remaining participants.

The only instance in which partitioning was not observed involved stimuli that were composed of integral dimensions and that could be readily classified by a verbal rule. However, partitioning was absent only when those stimuli were particularly easy to learn. When task difficulty was increased, and training performance equated with the levels observed with all other stimuli, partitioning emerged. We therefore suggest that knowledge partitioning emerges whenever a task is sufficiently difficult to encourage people to seek simplification by breaking a complex problem into independent components.

Concerns and Limitations

Before we turn to discussing the theoretical implications of our results, we consider three concerns or limitations relating to our experiments: (1) the possible unique status of rectangles, (2) the identifiability of verbal rules, and (3) the possibility that performance might uniformly rely on exemplar memory even when knowledge appeared to be partitioned.

The first issue concerns our use of rectangles varying in height and width as integral-verbal stimuli. There have been repeated suggestions in the literature that people may not consider those stimuli according to their height and width but according to some other variables, such as area and shape (Krantz & Tversky, 1975) or some emergent unidimensional property based on the ratio of their height and width (see, e.g., Monahan & Lockhead, 1977). Our response to this problem is twofold: First, the problem does not appear to be particularly widespread. Dunn (1983) reported that 7 out of 8 observers used height and width when responding to rectangles, with only 1 participant using area and shape instead. Second, and most important, the particular outcome of our study allows us to bypass the problem altogether. Given that the results for the rectangles did not deviate from those for the other stimuli when performance level was controlled, there is little reason to be concerned about the uniqueness of rectangles. The data show that irrespective of how people classify those stimuli, they can give rise to partitioned knowledge in exactly the same manner as the other stimuli used in these studies.

The second issue, concerning the identifiability of a verbalizable rule, has been somewhat contentious. Al-

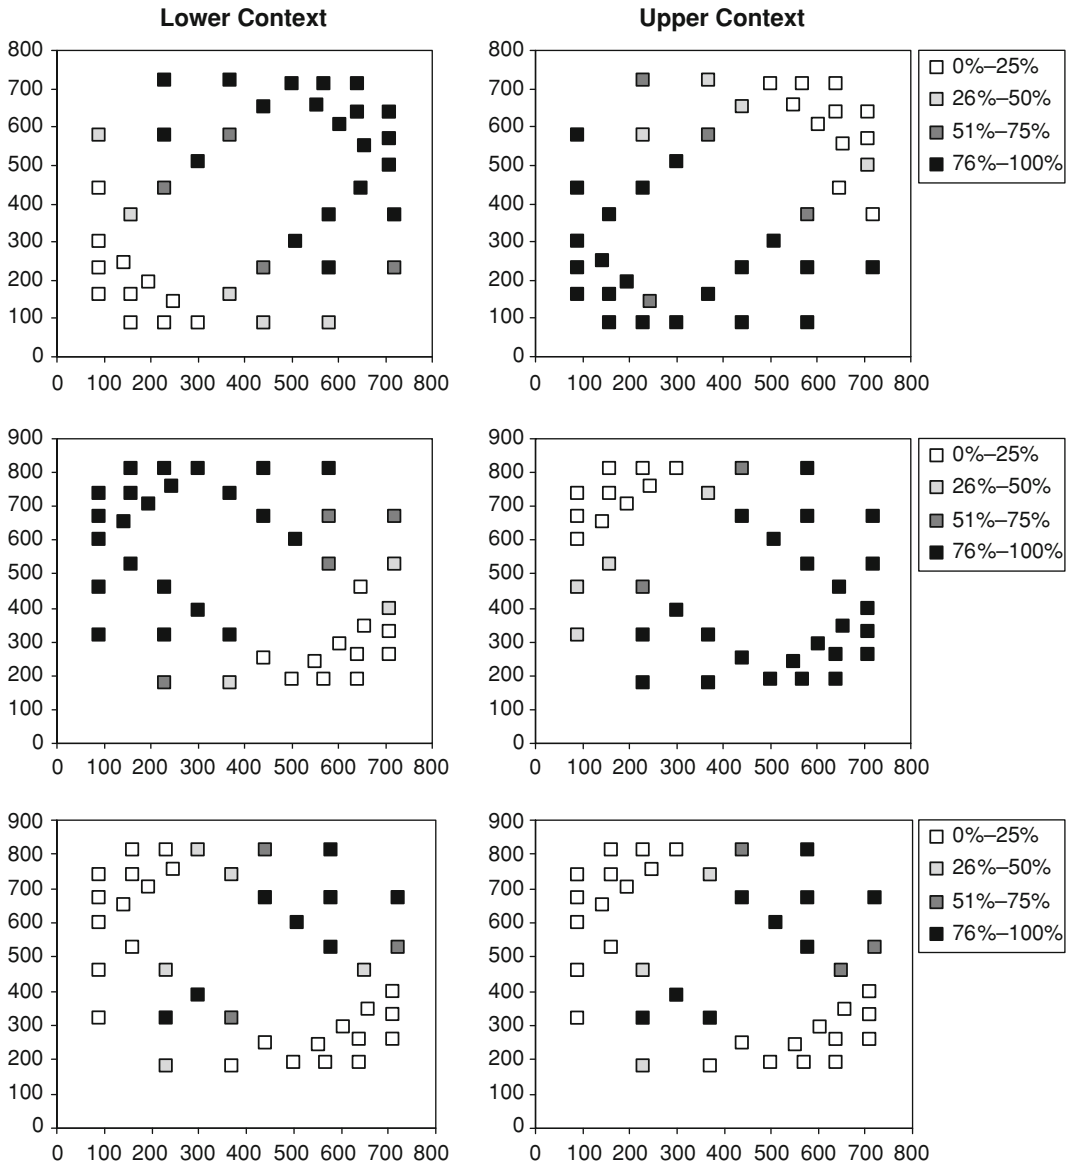


Figure 8. The probability of an A response for each novel transfer item in two contexts (left vs. right panels) for two groups of participants identified by the cluster analysis. The top panels show the responses of the knowledge-partitioning group in Experiment 2B, and the middle panels show the responses of the knowledge-partitioning group in Experiments 1 and 2A. The bottom panels show results for the true-boundary group in Experiments 1 and 2A.

though Ashby et al. (1998) suggested that rule-based responding in two-dimensional categorization tasks occurs only if stimulus dimensions are commensurate, Ashby and colleagues (e.g., Ashby et al., 2003) later acknowledged that conjunction rules (e.g., “respond *A* if the stimulus is small on dimension *x* and small on dimension *y*”; Ashby et al., 2003, p. 1115) may render a multidimensional task verbalizable even when stimulus dimensions are not commensurate. In confirmation, some participants in our separable–nonverbal condition (a circle with radial line) clearly relied on conjunction rules. In particular, people reported various approximations to the rule “Respond *A* if

circles are medium-sized and if the line is vertical, and respond *B* if circles are big or small and if the line is angled to the left or right.”

It follows that our operationalization of the verbal/nonverbal distinction, which was based on Ashby et al.’s (1998) criterion of dimensional commensurability, may have been insufficient to preclude the formulation of rules in the nonverbal conditions. However, Ashby et al. suggested that identification of how people learn categories is “strictly empirical” (2003, p. 1116). Given that people’s self-reported strategies were significantly more uniform when rules were thought to be verbalizable than when they

were not, we suggest that our manipulation was successful to the extent possible given currently available theorizing and empirical precedent.

The third problem concerns the category representation used in the present experiments. In particular, it is conceivable that people responded entirely on the basis of the similarity between test items and instances stored in memory, as would be expected from exemplar-based theories such as the generalized context model (Nosofsky, 1986) or ALCOVE (Kruschke, 1992). If that were the case, the observed partitioning would be more apparent than real, because it would derive from an integrated (albeit exemplar-based) internal representation. This possibility was ruled out by Yang and Lewandowsky (2004), who employed the same category structure and the same separable-verbalizable stimuli used in one of the present conditions, and who showed that pure exemplar-based models such as ALCOVE were unable to fit the data of people who partitioned their knowledge. By contrast, Yang and Lewandowsky (2004) also showed that a hybrid categorization model, ATRIUM (Erickson & Kruschke, 1998), which combines exemplar- and rule-based representations, was able to accommodate their knowledge partitioning data. Specifically, Yang and Lewandowsky showed that partitioning arose in ATRIUM because the model learned to associate different rule modules with each context.

Theoretical Implications

Our results suggest that knowledge partitioning is a potentially general aspect of human categorization performance. First, partitioning transcended the distinction between separate cognitive systems advocated by Ashby and colleagues (e.g., Ashby et al., 1998). Unlike several previous studies that produced performance dissociations depending on whether or not category rules were verbalizable (e.g., Ashby et al., 2003; Ashby et al., 1999; Maddox et al., 2003), no such distinction emerged in the present studies. Our data therefore sit very comfortably with theories that posit that categorization is handled by a single unitary memory system (e.g., Nosofsky & Johansen, 2000). Nonetheless, multiple-systems theories can also accommodate our results by recognizing knowledge partitioning as a common property of both categorization systems proposed by Ashby et al. (1998).

Second, knowledge partitioning occurred irrespective of whether stimuli were integral or separable. This result is noteworthy because people's ability to selectively attend to one dimension or another is significantly curtailed with integral relative to separable stimuli (see, e.g., McKinley & Nosofsky, 1996). It follows that partitioning does not require a selective shift of attention between the x - y dimensions of the stimulus.

The apparent generality of partitioning provides additional empirical support for the theoretical assertion by Yang and Lewandowsky (2004) that partitioning is a pervasive property of human categorization. In their exploration of ATRIUM, Yang and Lewandowsky (2004) showed that the model "partitioned" its knowledge even when simulating a condition in which stimuli were randomly

assigned to context during training—except that in that case partitioning was performed on the basis of the area (in x - y space) from which stimuli were sampled rather than on the basis of context. Kalish et al. (2004) similarly showed by experimentation as well as by simulation that partitioning can be considered an omnipresent aspect of the learning of continuous functions. We therefore suggest that the partitioning of a complex task into several simpler ones that are learned independently of each other is a pervasive attribute of human concept learning.

The Role of Complexity

We observed one exception to the generality of partitioning: It was absent when the task was simple and rapidly learned (i.e., in the integral-verbal condition in Experiments 1 and 2A). The absence of partitioning under those circumstances is consonant with earlier research by Lewandowsky et al. (2002), who, in a function-learning task, likewise observed partitioning to be absent when a positive linear function was learned particularly rapidly. Lewandowsky et al. observed partitioning in all other circumstances, including the learning of a linear function with *negative* slope. The emergence of partitioning with a negative function exactly parallels the present effects of rotating the category boundaries.

Conclusions

The findings presented here reinforce the conclusion that knowledge partitioning is a general phenomenon of perceptual categorization. They demonstrate that partitioning transcends two important distinctions in categorization—namely, whether or not rules can be verbalized and whether stimulus dimensions are integral or separable.

Moreover, we identified task difficulty as one boundary condition of partitioning in categorization. Unless tasks are very simple and rapidly learned, a significant proportion of people employ partitioning to facilitate learning.

REFERENCES

- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN, A. U., & WALDRON, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442-481.
- ASHBY, F. G., ELL, S. W., & WALDRON, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, *31*, 1114-1125.
- ASHBY, F. G., QUELLER, S., & BERRETTY, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, *61*, 1178-1199.
- BÉDARD, J., & CHI, M. T. (1992). Expertise. *Current Directions in Psychological Science*, *1*, 135-139.
- CORTESE, J. M., & DYRE, B. P. (1996). Perceptual similarity of shapes generated from Fourier descriptors. *Journal of Experimental Psychology: Human Perception & Performance*, *22*, 133-143.
- DUNN, J. C. (1983). Spatial metrics of integral and separable dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, *9*, 242-257.
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107-140.
- ESSOCK, E. A. (1980). The oblique effect of stimulus identification considered with respect to two classes of oblique effects. *Perception*, *9*, 37-46.

- GARNER, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- GARNER, W. R., & FELFOLDY, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, **1**, 225-241.
- GLASER, R. (1996). Changing the agency for learning: Acquiring expert performance. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports, and games* (pp. 303-311). Mahwah, NJ: Erlbaum.
- GOLDSTONE, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, **123**, 178-200.
- KALISH, M. L., LEWANDOWSKY, S., & KRUSCHKE, J. K. (2004). Population of linear experts: Knowledge partitioning and function learning. *Psychological Review*, **111**, 1072-1099.
- KRANTZ, D. H., & TVERSKY, A. (1975). Similarity in rectangles: An analysis of subjective dimensions. *Journal of Mathematical Psychology*, **12**, 4-34.
- KRUSCHKE, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22-44.
- KRUSCHKE, J. K. (2005). Learning involves attention. In G. Houghton (Ed.), *Connectionist models in cognitive psychology* (pp. 113-140). Hove, U.K.: Psychology Press.
- KRUSKAL, J. B. (1964). Nonmetric multidimensional scaling: A numerical method. *Psychometrika*, **29**, 115-129.
- LASAGA, M. I., & GARNER, W. R. (1983). Effect of line orientation on various information-processing tasks. *Journal of Experimental Psychology: Human Perception & Performance*, **9**, 215-225.
- LEWANDOWSKY, S., KALISH, M., & NGANG, S. K. (2002). Simplified learning in complex situations: Knowledge partitioning in function learning. *Journal of Experimental Psychology: General*, **131**, 163-193.
- LEWANDOWSKY, S., & KIRSNER, K. (2000). Knowledge partitioning: Context-dependent use of expertise. *Memory & Cognition*, **28**, 295-305.
- MADDOX, W. T. (1992). Perceptual and decisional separability. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 147-180). Hillsdale, NJ: Erlbaum.
- MADDOX, W. T., & ASHBY, F. G. (1998). Selective attention and the formation of linear decision boundaries: Comment on McKinley and Nosofsky (1996). *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 301-321.
- MADDOX, W. T., ASHBY, F. G., & BOHIL, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 650-662.
- MADDOX, W. T., ASHBY, F. G., ING, A. D., & PICKERING, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, **32**, 582-591.
- McKINLEY, S. C., & NOSOFSKY, R. M. (1996). Selective attention and the formation of linear decision boundaries. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 294-317.
- MONAHAN, J. S., & LOCKHEAD, G. R. (1977). Identification of integral stimuli. *Journal of Experimental Psychology: General*, **106**, 94-110.
- NOSOFSKY, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**, 39-57.
- NOSOFSKY, R. M. (1989). Further tests of an exemplar-similarity approach to relating identification and categorization. *Perception & Psychophysics*, **45**, 279-290.
- NOSOFSKY, R. M., & JOHANSEN, M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review*, **7**, 375-402.
- OP DE BEECK, H., WAGEMANS, J., & VOGELS, R. (2003). The effect of category learning on the representation of shape: Dimensions can be biased but not differentiated. *Journal of Experimental Psychology: General*, **132**, 491-511.
- SHEPARD, R. N. (1991). Integrality versus separability of stimulus dimensions: From an early convergence of evidence to a proposed theoretical basis. In G. R. Lockhead & J. R. Pomerantz (Eds.), *Perception of structure: Essays in honor of Wendell R. Garner* (pp. 53-71). Washington, DC: American Psychological Association.
- SHEPARD, R. N., & CHANG, J. J. (1963). Stimulus generalization in the learning of classifications. *Journal of Experimental Psychology*, **65**, 94-102.
- SHEPARD, R. N., HOVLAND, C. I., & JENKINS, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, **75**(13, Whole No. 517), 1-42.
- WALDRON, E. M., & ASHBY, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, **8**, 168-176.
- YANG, L.-X., & LEWANDOWSKY, S. (2003). Context-gated knowledge partitioning in categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 663-679.
- YANG, L.-X., & LEWANDOWSKY, S. (2004). Knowledge partitioning in categorization: Constraints on exemplar models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 1045-1064.

NOTES

1. Throughout this article, we couch the discussion in terms of a rule-based framework, rather than adopting a theoretically more neutral language that would be compatible with alternative approaches such as exemplar-based responding. This choice reflects the fact that Yang and Lewandowsky (2004) presented evidence that an exemplar-based account cannot accommodate knowledge partitioning in categorization. We will discuss this evidence more fully after all data have been presented.

To clarify our notation, from here on we will use *boundary* to refer to a design feature of the stimulus space and *rule* to refer to participants' imputed categorization strategy.

2. Some authors have suggested that separability and integrality are not unitary concepts but ought to be further differentiated into two empirically correlated but conceptually independent components—namely, *perceptual* and *decisional* separability or integrality (see, e.g., Maddox, 1992; Maddox & Ashby, 1998). According to this view, "perceptual separability and integrality are properties of particular pairs of stimulus dimensions and their interaction within the perceptual system . . ." (Maddox, 1992, p. 178), whereas decisional separability, though under the volitional control of the participant, is thought to be affected by variables such as the nature of the category space, instructions, and motivational factors. Because those variables were kept constant across all conditions in all experiments reported here, the effects of decisional separability and integrality were not manipulated in the present study, and we do not consider these concepts any further. Instead, our experimental manipulations only concerned the perceptual separability and integrality of the stimuli.

3. We also conducted a confirmatory similarity-rating experiment (not reported in detail here) to ascertain the dimensional interaction of our stimuli on the basis of an MDS analysis and identification of the preferred metric (Euclidean for integral stimuli and city-block for separable stimuli; see Shepard, 1991, for a review). As expected from the literature, the similarity structure among rectangles was better described by a Euclidean metric (badness-of-fit value of .139, using Kruskal's, 1964, stress measure) than the city-block metric (.149). The same pattern was observed for Fourier descriptors (.093 vs. .100). For circles, the reverse was observed, with the Euclidean metric giving rise to greater stress (.241) than the city-block metric (.150). The rectangles with a vertical bar were not included in the similarity-rating experiment because Erickson and Kruschke (1998) reported a scaling solution for those stimuli and confirmed their separability.

4. An omnibus 2 (dimensional interaction) \times 2 (verbalizability of task) \times 2 (context) \times 3 (area) between-within ANOVA revealed no additional effects that would have led us to abandon the simpler combined-condition analysis. Specifically, the only additional effects involving area or context were the interaction between area and verbalizability [$F(2,134) = 3.95$, $MS_e = 0.09$, $p < .05$] and a qualifying three-way interaction involving area, verbalizability, and dimensional interaction [$F(2,134) = 12.67$, $MS_e = 0.09$, $p < .0001$]. Because this interaction was merely indicative of performance differences between the four conditions, it is not pursued further.