Journal of Experimental Psychology:
Learning, Memory, and Cognition
2003, Vol. 29, No. 4, 663–679

0278-7393/03/$12.00 DOI: 10.1037/0278-7393.29.4.663

Context-Gated Knowledge Partitioning in Categorization

Lee-Xieng Yang and Stephan Lewandowsky
University of Western Australia

According to the knowledge partitioning framework, people sometimes master complex tasks by creating multiple independent parcels of partial knowledge. Research has shown that knowledge parcels may contain mutually contradictory information, and that each parcel may be used without regard to knowledge that is demonstrably present in other parcels. This article reports 4 experiments that investigated knowledge partitioning in categorization. When component boundaries of a complex categorization were identified by a context cue, a significant proportion of participants learned partial and independent categorization strategies that were chosen on the basis of context. For those participants, a strategy used in one context was unaffected by knowledge demonstrably present in other contexts, suggesting that knowledge partitioning in categorization can be complete.

Research on concept learning and categorization is currently guided by three classes of theories. The first class includes exemplar-based models, which posit that people classify objects into categories on the basis of their similarity to previously encountered exemplars (e.g., the generalized context model [GCM], Nosofsky, 1986; attention learning covering map [ALCOVE], Kruschke, 1992). The second class consists of rule-based models, which posit that people use a decision bound (e.g., a line or a curve) to divide a psychological space into categories (e.g., general recognition theory [GRT]; Ashby & Gott, 1988; Ashby & Maddox, 1990, 1992). The rule-based approach reduces to a traditional prototype model under simplifying assumptions about the shape and orientation of the decision bound (Ashby & Maddox, 1992). A third class of models, which has become more prominent during the last few years, combines elements of instance-based and rule-based theories into a hybrid rule-plus-exception approach (e.g., Erickson & Kruschke, 1998; Nosofsky, Palmeri, & McKinley, 1994; Thomas, 1998; Vandierendonck, 1995). According to these models, people seek to identify a simple rule that predicts classification for most items. The rule is augmented by memorization of a small set of exceptional instances or feature combinations.

Common to most theories is the premise that the representation acquired during learning is relatively homogeneous and invariant across different test situations. This premise is most readily understood for the rule-based (e.g., Ashby & Gott, 1988) or prototype-based (e.g., Homa, Sterling, & Trepel, 1981) approaches, which by definition condense disparate experiences into a single integrated representation—a rule or a prototype—that underlies categorization. Although perhaps less obvious, a similar homogeneity of knowledge is also assumed by instance models (e.g., GCM; Nosofsky, 1986), because *all* stored representations are considered for categorization and all are treated qualitatively identically. Specifically, although stimulus dimensions may attract different amounts of attention (e.g., Kruschke, 1992), all memorized instances enter equally into comparisons to test stimuli. In consequence, although individuals may differ in the way in which they approach and represent a classification problem (e.g., Nosofsky et al., 1994), there is much consensus that whatever representation is acquired would be insensitive to, say, the context in which it is tested.

The present article critically examines this consensus and reports four experiments that show that a significant proportion of participants master a complex categorization task by acquiring partial knowledge "parcels." The results also show that once chosen, a knowledge parcel provides the sole basis of performance, to the exclusion of knowledge demonstrably present in other parcels. Our results support and extend related findings in the areas of expertise (Lewandowsky & Kirsner, 2000), function learning (Kalish, Lewandowsky, & Kruschke, 2001; Lewandowsky, Kalish, & Ngang, 2002), and categorization (e.g., Erickson & Kruschke, 2001).

The structure of this article is as follows: We first review evidence that people have a more fluid, flexible, and heterogeneous category representation than is often assumed. We next consider more extreme findings from the related area of function learning that show that people can acquire independent parcels of knowledge that contain mutually contradictory information. These results go beyond the heterogeneity observed in categorization to date and are accommodated by a conceptual framework known as knowledge partitioning. We then present four experiments that show that partitioned knowledge can also arise in category learning.

Lee-Xieng Yang and Stephan Lewandowsky, School of Psychology, University of Western Australia, Crawley, Western Australia.

Preparation of this article was facilitated by a Large Research Grant from the Australian Research Council to Stephan Lewandowsky. We thank Leo Roberts for assistance during manuscript preparation. We gratefully acknowledge contributions by John Kruschke to early discussions of this research and comments by Michael Erickson and Robert M. Nosofsky on an earlier version of this article. We are also indebted to Mike Kalish for his comments and contributions throughout the project.

Additional related work can be found on Stephan Lewandowsky's World Wide Web page at http://www.psy.uwa.edu.au/user/lewan/

Correspondence concerning this article should be addressed to Stephan Lewandowsky, School of Psychology, University of Western Australia, Crawley, Western Australia 6009, Australia. E-mail: lewan@psy.uwa.edu.au

Heterogeneous Category Representations

Turning first to natural categories, we note that people are known to use their pre-experimental knowledge with nearly kaleidoscopic variety. For example, Roth and Shoben (1983) showed that people's typicality judgments for everyday stimuli (e.g., "Which beverage is more typical, tea or milk?") may reverse between contexts (e.g., whether a truck driver or secretary is thought to be consuming the beverage), thus providing evidence against an invariant semantic space and in favor of a more heterogeneous knowledge landscape. Medin and Shoben (1988) reached a similar conclusion using typicality judgments of adjective–noun combinations. For example, when judging the pairwise similarity between black, gray, and white, participants considered black and gray to be more similar than white and gray in the context of clouds, but this judgment was reversed in the context of hair color.

Similarly fluid performance has been observed in category learning with artificially constructed stimuli. For example, Lamberts (1994) presented exemplars from a hypothetical family defined by features such as size of ears or eyes or presence versus absence of hair. Participants then had to determine the family membership of a target person who was presented either as a brother or cousin of one of the training exemplars. Lamberts found that people applied a more rigorous inclusion criterion in the *brother* context (i.e., a brother must be very similar to other members of his family) than in the *cousin* context.

Even stronger evidence for representational heterogeneity was provided by Aha and Goldstone (1992). In their experiment, training stimuli were sampled from two distinct clusters in a two-dimensional category space, each bisected by its own uniquely oriented boundary. Because the design of Aha and Goldstone's study is relevant to our experiments, the stimulus space from their Experiment 1 is presented in Figure 1. Training items are identified by their category membership (either A or B), and transfer items are labeled X, Y, Z, and W. Stimuli varied along two perceptual dimensions—size of a rectangle and position of a vertical line segment within it—with eight possible levels for each dimension. Correct classification of the training instances in the top right cluster can be achieved by locating a vertical linear boundary in between 6 and 7 on the abscissa, whereas the bottom left cluster can be correctly classified by a horizontal boundary in between 6 and 7 on the ordinate. Note that the two boundaries are globally incompatible because when extended further from their cluster, they dictate opposite classifications for the same test items.

	1	2	3	4	5	6	7	8
1	•	•	•	•	•	•	•	•
2	•	•	•	•	•	B	**Z**	•
3	•	•	•	•	•	B	A	•
4	•	•	•	•	•	B	A	•
5	•	•	•	•	•	**Y**	A	•
6	•	A	A	A	**X**	•	•	•
7	•	**W**	B	B	B	•	•	•
8	•	•	•	•	•	•	•	•

Figure 1. Stimulus space used by Aha and Goldstone (1992). The abscissa denotes the size of the rectangle and the ordinate denotes the position of a vertical line segment within it. The training instances are shown as *A* and *B*. The four critical transfer items are marked with one of {*W, X, Y, Z*}.

Aha and Goldstone found that people classified the critical transfer stimuli X and Y predominantly as belonging to Category A and B, respectively, which means that people applied the vertical boundary to Y but the horizontal boundary to X. This shows that people can be sensitive to different patterns of features formed by a subset of stimuli and may apply different strategies to each subset. Aha and Goldstone also showed that the data could be handled only by a model that incorporated unique dimensional weights *for each exemplar*. Specifically, Aha and Goldstone modified Nosofsky's GCM (e.g., 1986) by inserting a weighting function into the computation of similarities between exemplars. The function incorporated unique learned weights for each exemplar, thus creating a heterogeneous internal category representation.

Erickson and Kruschke (2001) recently reported a replication and extension of Aha and Goldstone's (1992) study. In particular, Erickson and Kruschke (2001) showed that the observed heterogeneity was not an artifact of averaging across subgroups of participants who had mastered only one, but not the other, of the training clusters. Erickson and Kruschke also supported Aha and Goldstone's theoretical conclusion, namely that the data required a stimulus-dependent representation that is difficult to account for by current models.

In summary, there is evidence that the representations of natural and artificial categories can be quite heterogeneous. It turns out that research in the related domain of function learning has uncovered an even greater degree of heterogeneity.

Creating Partitioned Knowledge in Function Learning

In function learning experiments, people must learn the relationship between a continuous stimulus and a continuous response variable. For example, people may learn how long to water their lawn as a function of the day's temperature; they may learn the relationship between driving speed and stopping distance or between the number of alcoholic drinks consumed and likely blood alcohol level, and so on. Participants are presented with a sequence of learning trials, each of which involves a unique pairing of stimulus and target response magnitudes. Responses are followed by corrective feedback.

With sufficient practice, people can learn a variety of continuous functions from these discrete learning trials, as revealed by performance on novel transfer stimuli (e.g., Busemeyer, Byun, Delosh, & McDaniel, 1997). Transfer stimuli that fall within the range of training values, thus requiring interpolation, typically elicit highly accurate responses (see Busemeyer et al., 1997, for a review). For stimuli outside the trained range, participants are likewise capable of extrapolation, albeit at a lower level of accuracy (Delosh, Busemeyer, & McDaniel, 1997).

Using a U-shaped quadratic function, Lewandowsky et al. (2002) showed that people may acquire partitioned knowledge that is not integrated across training contexts. The manipulation involved the presence of a context label during training that probabilistically identified the local relationship between stimulus and response magnitudes. Specifically, 90% of magnitudes below the vertex of the quadratic function appeared with one context label, whereas 90% of magnitudes above the vertex were presented in the other context. This manipulation created what Lewandowsky et al. called a second-order relationship between context and response magnitudes. That is, owing to the symmetry of the quadratic function, context *did not* predict responses directly (i.e., no first-order relationship), as the mean response value was identical for stimulus magnitudes above and below the vertex.

Lewandowsky et al. (2002) observed context-specific transfer: Instead of extrapolating along the U-shaped quadratic function, participants extrapolated along a different function in each context. Because context had no first-order relationship to responses, this result could not have occurred if context had directly contributed to a single weighted decision rule. Instead, Lewandowsky et al. argued that their results represented an instance of so-called knowledge partitioning, that is, the creation during learning of independent parcels of knowledge that are selectively accessed on the basis of context. On this view, context "gates" access to a parcel of knowledge that is used to perform a task.

Another intriguing aspect of Lewandowsky et al.'s (2002) study was the apparent lack of linkage or integration between parcels. That is, once a parcel of knowledge was chosen and accessed on the basis of context, it governed performance at the exclusion of knowledge demonstrably available in other parcels. This conclusion was based on the fact that context-specific extrapolation was unaffected by whether or not people had experienced the other function component at all. The implications of this finding for categorization are best explored using the earlier study by Aha and Goldstone (1992).

Partitioned Knowledge in Categorization

In Aha and Goldstone's (1992) study, the critical transfer stimuli X and Y were classified according to the boundary within the nearest cluster of training items (see Figure 1). This arguably reflected some partitioning of knowledge because categorization depended on the immediate neighborhood of a transfer item, an assertion supported by the exemplar-specific dimensional weights required to model the data. However, at the same time, participants were sensitive to the global incompatibility of the two local boundaries, because generalizations from the training items were limited to the immediate vicinity of each cluster. For example, Item X was not classified as B despite this being suggested by the nearby vertical boundary. It follows that although people's knowledge was sensitive to the idiosyncrasies of each cluster, and hence heterogeneous, it was also sensitive to the presence of the other cluster.

The function learning results of Lewandowsky et al. (2002), by contrast, revealed that people may selectively access knowledge in one parcel without considering information demonstrably accessible in another context. In terms of Aha and Goldstone's (1992) design, this would be tantamount to showing that people generalize from each cluster of training items on the basis of information contained in that cluster alone, without regard to what they have learned about the other cluster and without awareness of the global incompatibility of the two local decision boundaries.

We now report four experiments that sought evidence for such partitioning of knowledge in category learning. The methodology was inspired by the function learning design of Lewandowsky et al. (2002); hence, a binary context label was used to manipulate and encourage knowledge partitioning. To facilitate presentation, we report and interpret the results from a rule-oriented perspective and identify the rules used by participants in each experiment (cf. Lewandowsky, Kalish, & Griffiths, 2000; Nosofsky et al., 1994). Discussion of the various theoretical alternatives is deferred to the General Discussion section.

Experiment 1

The first experiment used a category structure that was amenable to knowledge partitioning as defined by five criteria: (a) Categories were separated by a single but complex boundary that could be partitioned into simpler components; (b) the simple component boundaries had to be identifiable by a context variable during training; (c) context could not be a valid *direct* predictor (i.e., context had to be a second-order but not a first-order predictor); (d) the complex boundary had to be learnable when context either was absent or had no predictive value (i.e., was not even a second-order predictor); and finally, (e) there had to be a diagnostic set of transfer stimuli that could help determine whether or not knowledge partitioning emerged during training. These five criteria were satisfied as follows.

The principal between-subjects manipulation concerned the role of context. In the randomized-context condition, context labels were randomly assigned to training items and thus had no predictive value. This procedure served to satisfy Criterion D, that people can learn the complex boundary unless partitioning is encouraged. Partitioning was encouraged in the systematic-context condition, in which context was consistently mapped not to an outcome but to subsets of training items, each characterized by a component boundary. This manipulation satisfied Criteria B and C above. (The systematic-context condition introduced a correlation among predictors within a category. Discussion of previous research on the correlated predictors is deferred until after presentation of all results.) Both conditions used the same set of transfer items that were tested in both contexts. This permitted identification of knowledge partitioning (Criterion E) through diagnostic response patterns that are described below.

Method

Participants and Apparatus

Sixty-two volunteers from the University of Western Australia campus community participated voluntarily in exchange for partial course credit. An equal number of participants was randomly assigned to each condition.

The experiment was controlled by a personal computer that presented all stimuli and collected and scored all responses. The same apparatus was used in all experiments.

Stimuli

All stimuli were sampled from the two-dimensional pseudocontinuous category space shown in Figure 2. The bilinear category boundary, represented by the solid line, is described by: $Y = 500 - |X - 400|$. Stimuli below and above the triangular boundary belonged, respectively, to Category A and B. To clarify notation, from here on we use *boundary* to refer to a design feature of the stimulus space and we use *rule* to refer to participants' imputed categorization strategy.

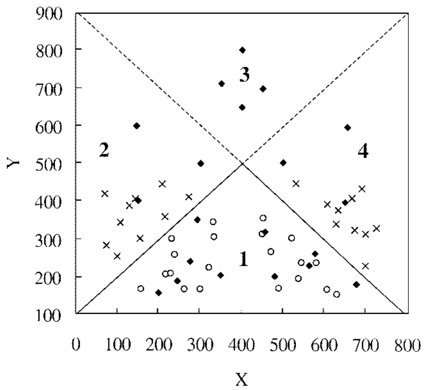

Figure 2. Stimulus space formed by dimensions X and Y used in Experiment 1. The category boundary is indicated by the solid line. All stimuli above the boundary belong to Category B, whereas all stimuli below belong to Category A. The training instances are shown as open circles (Category A) and crosses (Category B). Note that Y < 500 for all training instances. Transfer items are shown as solid diamonds. Numbers 1−4 refer to diagnostic areas. Dashed lines represent the partial boundaries expected to be learned if people partition their knowledge.

The two categories were arbitrarily instantiated as different species of fish, with predictors X and Y labeled as "density" and "depth," respectively. Participants were told that density referred to the amount of food available in the habitat, whereas depth referred to how far below the surface the fish preferred to live.

Training items. All participants were trained on a common set of 40 stimuli, all with Y < 500. Training stimuli were separated into two subsets of 20 with values of X above and below 400. Each subset contained an equal number of instances of each category that clustered around the local diagonal boundary whose orientation was positive for X < 400 and negative for X > 400. Training instances belonging to Category A are represented by open circles in Figure 2 and those belonging to Category B by crosses. The training sequence consisted of 10 blocks, each of which involved presentation of all 40 items in a different random order.

Each stimulus also included a third, binary context feature (C), not shown in the figure. Regardless of condition, both levels of context (C_1 and C_2) were equally represented among training instances and categories. Context therefore was not a direct (i.e., first-order) predictor. In the randomized-context condition, context also had no second-order relationship because it was randomly paired with training instances, subject to the constraint that each context had to occur equally often with each category. In the systematic-context condition, one level of context (e.g., C_1) occurred with stimuli whose value of X was below 400, whereas the other level (e.g., C_2) occurred with values of X > 400. Hence there was a perfect correlation between C and the *relationship* between X and Y that defined the local boundary between Categories A and B.

Context was instantiated as "season" using the two levels "summer" and "winter." Assignment of context labels to magnitudes of X was counterbalanced between participants; for ease of exposition, from here on we refer to context as *left* and *right* to represent whichever label accompanied values of X below and above 400, respectively.

Transfer items. The 20 transfer stimuli are identified by filled diamonds in Figure 2. There were two blocks of transfer trials, each consisting of a different random order of all transfer stimuli. Context alternated between blocks for each transfer stimulus, and each context was represented equally often in each block.

Diagnosticity of transfer items. Transfer items were sampled from the four areas identified by number in Figure 2. The expected pattern of responses differs between and within areas depending on whether or not people partitioned their knowledge.

If partitioning is absent, then the true boundary should apply in both contexts, and all items from Area 1 should be consistently classified as belonging to Category A and items from the remaining areas (2–4) as belonging to Category B. A very different pattern is expected for the systematic-context condition if people partition their knowledge. In that case, one would expect the positive diagonal component of the boundary to be used to classify *all* stimuli in the *left* context. This extended partial rule is indicated by the dotted lines along the positive diagonal in Figure 2. Accordingly, in the left test context, stimuli in Areas 2 and 3 should be predominantly classified as belonging to Category B, whereas those in Areas 1 and 4 are expected to be classified as A. The reverse situation is expected in the *right* context. In that case, the negative partial boundary is expected to be used (also indicated by a dotted line in the figure) and stimuli in Areas 1 and 2 would be classified as A, whereas those in Areas 3 and 4 would be classified as B. Overall, knowledge partitioning would be identified if responses differ between test contexts for Areas 2 and 4 but remain the same in Areas 1 and 3.

Procedure

Participants were tested individually in a quiet booth. On every trial, an unchanging schematic icon of a fish was displayed on the screen. Below the icon, the three labels density, depth, and season were displayed from left to right, with the corresponding feature values displayed immediately underneath. Categorization responses were recorded using the *F* and *J* keys. Assignment of categories to keys alternated across participants. Each trial commenced with display of a fixation signal (a "+") in the center of the screen for 500 ms, followed by presentation of the stimulus, which remained visible until participants responded.

In the training phase, participants were given feedback (the word "correct" or "wrong" shown in the center of the screen) for 1000 ms after each response before the next trial commenced. No feedback was presented during transfer.

Self-paced breaks were inserted after every 40 training trials and after every 20 transfer trials. In addition, there was a break between training and transfer. The experiment lasted just under 1 hr.

Results and Discussion

Training Performance

Training performance was analyzed using the proportion of correct responses in each block as the dependent measure. A 2 (condition) × 10 (training block) between-within analysis of variance (ANOVA) revealed a significant main effect of block, $F(9, 540) = 51.193$, $MSE = .009$, $p < .001$, along with a marginal main effect of condition, $F(1, 60) = 3.889$, $MSE = 0.103$, $p = .053$. The interaction between both variables was not significant, $F(9, 540) < 1$. The main effect for block resulted from an improvement in performance from the first ($M = .552$) to the last block of training ($M = .797$), whereas the main effect of condition reflected the fact that people were better overall in the systematic-context condition ($M = .726$) than in the randomized-context condition ($M = .675$).

The latter effect replicated a related finding by Lewandowsky et al. (2002), who also found that performance during training was better for the systematic-context condition. We suggest that this reflects the enhanced ease of the task when it can be partitioned into independent subcomponents on the basis of context. We next show that this type of partitioning occurred in Experiment 1.

Transfer Performance

The dependent variable for all transfer analyses in this article was the probability of classifying a stimulus as belonging to

Category A. Responses were aggregated across transfer items within each of the four diagnostic areas. Table 1 shows the mean transfer responses for each condition and area.

A 2 (condition) × 2 (test context) × 4 (area) between-within ANOVA revealed a main effect of area, $F(3, 180) = 76.12$, $MSE = 0.12$, $p < .01$, and 2 two-way interactions involving test context and area, $F(3, 180) = 8.47$, $MSE = 0.04$, $p < .01$, and condition and test context, $F(1, 60) = 6.67$, $MSE = 0.04$, $p < .05$. These effects were qualified by an overarching interaction involving all three variables, with $F(3, 180) = 5.17$, $MSE = 0.04$, $p < .01$. No other effects were significant, with the largest $F = 2.19$ ($p > .10$).

The crucial three-way interaction was first explored by computing a separate Test Context × Area effect for each condition, which was found to be significant in the systematic-context condition, $F(3, 90) = 8.95$, $MSE = 0.06$, $p < .01$, but not in the randomized-context condition, $F(3, 90) = 1.22$. Further exploration involved simple comparisons between test contexts for all areas in the systematic-context condition. These tests revealed no significant effect in Area 1, $F(1, 30) = 2.06$, $MSE = 0.02$, $p > .10$, but significant effects in Area 2, $F(1, 30) = 11.14$, $MSE = 0.09$, $p < .01$; Area 3, $F(1, 30) = 10.87$, $MSE = 0.06$, $p < .01$; and Area 4, $F(1, 30) = 3.73$, $MSE = 0.07$, $p = 0.063$.

Knowledge partitioning in categorization. These results suggest a straightforward interpretation: First, people were clearly able to learn the correct bilinear boundary when context was of no predictive value. In the randomized-context condition, people correctly classified transfer items below the bilinear boundary as belonging to Category A, whereas all others—including those in Area 3 furthest from the trained region—were predominantly classified as belonging to Category B. Second, when context predicted the orientation of a local boundary, there was some suggestion that people learned to partition their knowledge. Transfer performance was sensitive to context in Areas 2 and 4 (but not Area 1), as would be expected when partial rules are used.

However, these conclusions are accompanied by two limitations. The first involves the apparent asymmetry of the partial rules. The second concerns the magnitude and consistency of the effect.

Table 1
Mean Probabilities of Category A Responses in Experiments 1 and 2 for Novel Transfer Items in Each Area and Test Context (Left or Right)

Condition	Area 1	Area 2	Area 3	Area 4
Experiment 1				
Randomized context				
Left	.79	.17	.26	.33
Right	.79	.20	.22	.27
Systematic context				
Left	.79	.09	.07	.44
Right	.75	.33	.28	.31
Experiment 2				
Systematic context				
Left	.80	.16	.19	.62
Right	.74	.47	.39	.40

Asymmetry of partitioning. If people in the systematic-context condition consistently applied one or the other partial rule, the pattern of responding in Areas 1 and 3 should remain unaffected by test context. The data partially contradict this expectation: Although there was no effect of context for stimuli in Area 1, there was an unexpectedly large and significant effect in Area 3.

This unexpected result is explored further in Figure 3, which shows the probability of a Category A response, averaged across participants, for each transfer item in the systematic-context condition separated by test contexts. Figure 3A (*left* test context) suggests that a partial rule was applied with a fair amount of consistency to most test items (with the exception of at least one item, in the lower right, for which P(A) was .29 despite it being in the region in which items should be classified as belonging to A). By contrast, Figure 3B (*right* context) does not suggest use of a partial rule. Instead, the data are more consistent with (imperfect) use of the correct bilinear boundary. This asymmetry among test contexts is reminiscent of the known characteristic of function learning that ascending functions are easier to learn than descending functions (e.g., Busemeyer et al., 1997; Lewandowsky et al., 2002).

Magnitude of partitioning: Individual differences. The magnitude of the context effect for Areas 2 and 4 is not indicative of a consistent application of partial rules. Consistent knowledge partitioning by all participants would imply probabilities close to 0 and 1 across the two contexts. This was not observed. For example, in the *left* context, P(A) for Area 4 was .44, which is close to what would be expected by chance or if different individuals applied different strategies.

We explored individual differences by entering participants' response profiles (i.e., the pattern of a participant's A vs. B responses to all transfer stimuli) into a *k*-means cluster analysis that used four predefined clusters (cf. Lewandowsky et al., 2000). Cluster centroids were predefined to represent application of the (a) correct bilinear boundary, (b) the left partial boundary, (c) the right partial boundary, and (d) a horizontal boundary at $Y = 300$. Using a Euclidean distance measure, the analysis assigned each participant's response profile to the closest cluster. Participants who were equidistant from all predefined patterns were considered to be unrecognizable.

Table 2 shows the assignment of participants to clusters for each condition and test context. Confirming the overall transfer analysis, the top panel of the table shows that there was little or no effect of context in the randomized-context condition. The results in the systematic-context condition differed in several important ways: First, fewer participants overall applied the correct bilinear boundary. Second, there was a notable effect of context for the partial rule clusters, with a significant proportion of participants using the partial rule that was appropriate to the context. Third, fewer participants overall applied the descending partial boundary than the ascending one; this again confirms that negative relationships are more difficult to learn than positive ones.

One problem concerns the large percentage of participants who remained unrecognizable. Two related reasons can be cited for this: First, there were few stimuli in the generalization regions (i.e., with $Y > 500$), which made differentiation of some candidate rules more difficult. Second, there were two transfer stimuli (one in

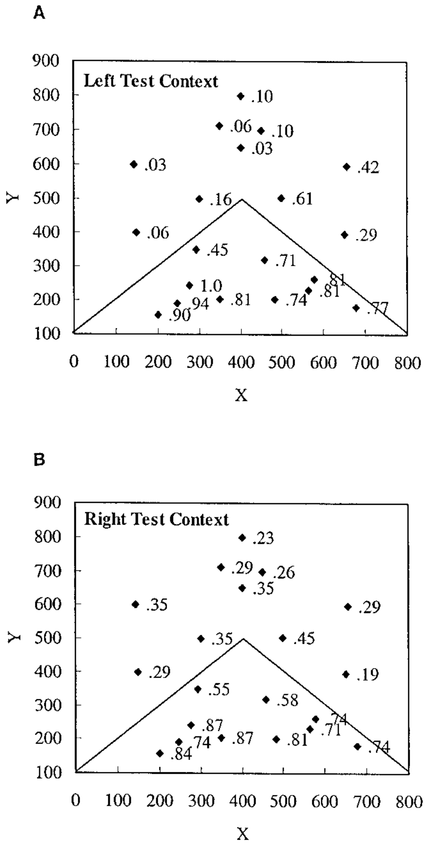

Figure 3. The average probability of Category A responses for each transfer item in the systematic-context condition of Experiment 1. The category boundary is indicated by the solid line. Transfer items are shown as solid diamonds. A: the *left* test context. B: the *right* context.

Area 2 and one in Area 4) that were in close proximity to the centroid of studied items belonging to Category B. The similarity of those test stimuli to studied exemplars may have overridden the application of a partial rule that dictated an opposing classification (e.g., Nosofsky, Clark, & Shin, 1989).

This problem was resolved by Experiment 2, which used the same category space but ensured that transfer items were further away from the trained stimuli. Experiment 2 also included two control conditions to help identify the extent of knowledge partitioning.

Experiment 2

The category structure used in this experiment is shown in Figure 4. It differs from the previous one only with respect to the location of transfer items.

The randomized-context condition was not included in Experiment 2. Instead, Experiment 2 included two single-context learning conditions. In the left-only condition, only the *left* set of instances (i.e., those with X < 400) were presented for training in a single context. In the complementary right-only condition, participants learned the *right* set of instances (X > 400).

Comparison of transfer performance in these two single-context conditions to performance in the matching test context of the systematic-context condition can reveal the extent of partitioning. Specifically, should transfer performance be identical between conditions within each context, then the two partial knowledge components acquired in the systematic-context condition could be considered largely independent. Conversely, should those comparisons reveal an effect of condition within each context, this would identify at least some linkage between knowledge components.

Table 2
Number and Percentage of Participants Identified as Using a Particular Strategy by k-Means Cluster Analysis in Experiment 1

	Test context			
	Left		Right	
Cluster	*n*	%	*n*	%
Randomized-context condition				
True boundary	9	29	7	23
Left boundary	8	26	7	23
Right boundary	3	10	5	16
Y = 300	5	16	5	16
Unrecognized	6	19	7	23
Systematic-context condition				
True boundary	4	13	5	16
Left boundary	13	42	6	20
Right boundary	0	0	5	16
Y = 300	6	16	5	16
Unrecognized	8	26	10	32

Note. N = 31 in each condition.

Method

Participants

Sixty-three undergraduates from the University of Western Australia participated voluntarily in exchange for partial course points. Thirty-one

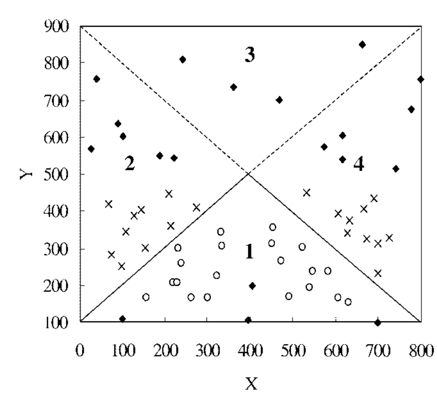

Figure 4. The stimulus space formed by Dimensions X and Y used in Experiment 2. The category boundary is indicated by the solid line. The training instances are shown as open circles (Category A) and crosses (Category B). The transfer items are shown as solid diamonds. Numbers refer to different diagnostic areas. Dashed lines represent the partial boundaries expected to be learned if people partition their knowledge.

participants were randomly assigned to the systematic-context condition and 16 each to the left-only and right-only conditions.

Stimuli and Procedure

The systematic-context condition used the 40 learning instances from Experiment 1. In the left-only condition, participants learned the 20 instances located in the left half of the stimulus space with a single context cue (counterbalanced across participants). Conversely, training in the right-only condition involved the other 20 learning instances. There were 10 blocks of training trials in all conditions. Each block contained a different random sequence of the training items.

The transfer test was identical for all conditions and involved a novel set of 20 stimuli that was further removed from training items (see Figure 4). The remaining details of the transfer test and the procedure were unchanged from Experiment 1. This implied that half the transfer stimuli in the left-only and right-only conditions were presented in a completely new context not encountered during training.

Results and Discussion

Training Performance

The probability of correct responses across blocks increased from .64 (Block 1) to .84 (Block 10). Performance differed considerably between conditions, with the systematic-context condition (M = .70) being worse overall than the left-only (.88) and right-only (.76) conditions. This difference between conditions was not unexpected because the latter condition included twice as many learning instances as each of the two single-context conditions. The single-context advantage may have been further enhanced by the greater complexity of the bilinear category boundary in the systematic-context condition. In addition, the superiority of the left-only over the right-only condition again shows that the ascending partial boundary was easier to learn than its descending counterpart.

These impressions were confirmed by a 3 (condition) × 10 (block) between-within ANOVA that revealed main effects for block, $F(9, 540) = 23.8$, $MSE = 0.01$, $p < .01$, and condition, $F(2, 60) = 10.31$, $MSE = 0.17$, $p < .01$, but no interaction between both variables, $F(18, 540) < 1$. Exploration of the main effect of condition (by Dunnett's test) revealed no differences between the systematic-context condition and the right-only condition but a consistent difference in all blocks between the systematic-context and left-only conditions. The corresponding pairwise differences between the left-only and right-only conditions were significant or marginally significant in Blocks 4 through 8.

Transfer Performance

Knowledge partitioning. We first determined whether or not knowledge partitioning was present by focusing on the systematic-context condition. The average probabilities of Category A responses in the four areas of interest are shown in Table 1.

A 2 (test context) × 4 (area) within-subjects ANOVA revealed significant main effects of context, $F(1, 30) = 5.01$, $MSE = 0.05$, $p < .05$, and area, $F(3, 90) = 25.08$, $MSE = 0.12$, $p < .01$, as well as a significant interaction between both variables, $F(3, 90) = 11.55$, $MSE = 0.08$, $p < .01$. Simple comparisons between the two contexts in all four areas were significant for Area 2, $F(1, 30) = 18.14$, $MSE = 0.08$, $p < .001$; Area 3, $F(1, 30) = 7.87$, $MSE = 0.08$, $p < .01$; and Area 4, $F(1, 30) = 10.44$, $MSE = 0.07$, $p < .01$. The comparison was nonsignificant only for Area 1, with $F(1, 30) = 1.09$.

As in the first experiment, the context effects in Areas 2 and 4 were indicative of the presence of knowledge partitioning, whereas the significant effect for Area 3 was again unexpected. One possible reason underlying the effect in Area 3 was that the descending partial boundary was harder to learn than the ascending one. This possibility can be examined by comparing the right-only and left-only conditions. Figure 5A shows response probabilities for all transfer items presented in the trained context in the left-only condition; Figure 5B shows the same for the right-only condition. For the left-only condition, it is clear that people learned the partial boundary involving the positive diagonal quite well, because all items below that boundary are classified with near uniformity as belonging to Category A, whereas those above it are classified as B. For the right-only condition, by contrast, there is considerably less evidence that people applied the descending partial boundary.

Extent of partitioning. The single-context conditions also afford an assessment of the extent of partitioning by examining

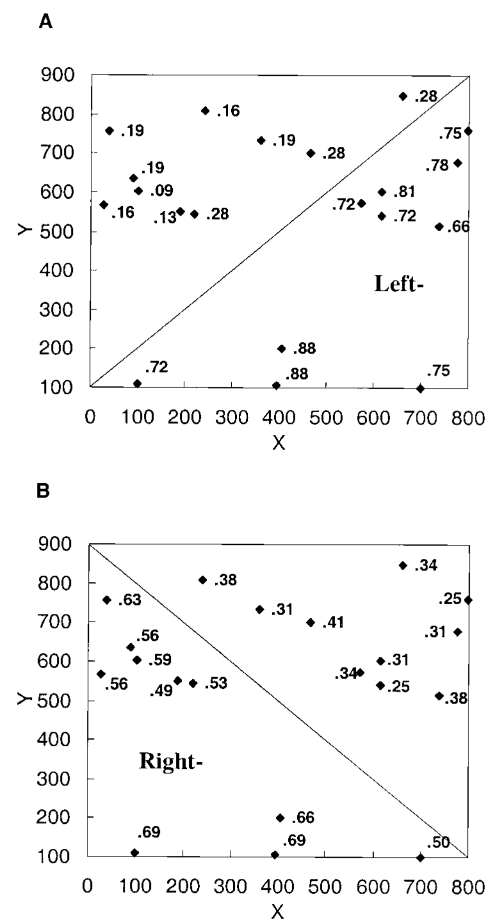

Figure 5. Average probabilities of Category A responses for each transfer item in the two single-context conditions in Experiment 2. A: left-only condition. B: right-only condition. Solid diamonds refer to transfer items. Solid lines refer to the partial linear boundaries that appear to have been learned in the single-context conditions.

whether participants used a partial strategy in its appropriate test context as if the other one had never been learned.

We thus computed correlations between the classification probabilities of each test item in the systematic-context condition and the two single-context conditions. One correlation was computed within the same context (i.e., responses in the *left* context in the systematic-context condition were correlated with responses in the left-only condition and equivalently for the *right* context). The other correlation was computed across the two contexts (i.e., responses in the *left* context in the systematic-context condition were correlated with responses in the right-only condition and equivalently for the *right* context). Note that both correlations were computed between subjects. The correlation was $r = 0.892$ ($N = 20$; regression intercept and slope .04 and .94, respectively, with mean absolute deviation .087) when contexts were the same and $r = 0.19$ (intercept and slope .38 and .20, respectively, and mean absolute deviation .21) when contexts differed.

The magnitudes of these correlations are best interpreted in light of the correlation *between* contexts *within* the systematic-context condition. That correlation, which involved responses by the same participants across contexts, was $r = 0.55$ (intercept and slope = –.08 and 1.05, respectively, mean absolute deviation = .21) and hence considerably smaller than the between-condition correlation within the same context—even though the latter involved responses from different individuals. It follows that in the systematic-context condition, people in each test context classified transfer items in virtually the same way as people in the corresponding single-context condition who had never been exposed to the other context during training.

Individual differences. We conclude by considering the individual differences in the systematic-context condition. The horizontal-boundary cluster used in Experiment 1 was omitted from this analysis because with the revised set of transfer items, this strategy was not distinguishable from use of the true boundary. Hence only three possible strategies were predefined for the *k*-means cluster analysis, corresponding to application of (a) the correct bilinear boundary, (b) the ascending partial boundary, and (c) its descending counterpart. Table 3 shows the resulting classification of participants in both test contexts. Several comments can be made about the classification: First, some 20% of participants seemed to apply the correct bilinear boundary in each context, whereas roughly half were identified as applying the appropriate partial boundary. We therefore conclude that the magnitude of the partitioning effect was greater than that observed in Experiment 1.

Table 3
Number and Percentage of Participants Identified as Using a Particular Strategy by k-Means Cluster Analysis in Experiment 2

	Test context			
	Left		Right	
Cluster	*n*	%	*n*	%
Systematic-context condition				
True boundary	6	19	7	23
Left boundary	16	52	7	23
Right boundary	3	10	13	42
Unrecognized	6	19	4	12

Note. $N = 31$.

Statistical support for the effect of test context on classification was provided by computing Cramer's coefficient (ϕ) for the association between test context and cluster membership in each condition. Cramer's coefficient is a transformation of chi-square and is readily interpretable as a measure of association that ranges from 0 to unity (Wickens, 1989). If people change their categorization strategy with test context, there should be a strong association between context and cluster membership. This expectation was confirmed by the magnitude of Cramer's coefficient ($\phi = .407$, $p < .02$).[1]

Summary. This experiment gave rise to some notable results. First, in comparison to Experiment 1, we observed a much clearer indication that knowledge partitioning was present in a significant proportion of participants. Second, by examining the two single-context conditions, we confirmed an asymmetry of the partial boundaries that was independent of the context manipulation and that reflected basic characteristics of category learning. Third, the correlational analyses supported the conclusion that people in our experiment used different context-gated components of knowledge independently. That is, in contrast to the findings by Aha and Goldstone (1992), in which people were shown to be sensitive to the existence of the other cluster when they generalized their responses along one of the boundaries, the present experiment revealed a more extreme case in which people generalized their responses along one of the boundaries as if they had never known the other. This confirms and extends previous results in function learning recently reported by Lewandowsky et al. (2002).

Experiment 3

The results thus far have supported our expectations: Knowledge partitioning arises in category learning in a significant proportion of participants, and when it arises, it can be nearly complete. However, the category structure used thus far entailed two limitations. First, the difficulty of learning differed between the partial boundaries, which may have prevented knowledge partitioning if people were unable to use one of the partial boundaries. Second, the context manipulation was perfectly correlated with one of the predictors: In the systematic-context condition, all values of X below 400 were accompanied with one level of context and those above 400 with the other. It follows that people may have partitioned their knowledge not only on the basis of context, but also on the basis of the value of X—indeed, for this reason the stimulus structure afforded the opportunity for partitioning even in the randomized-context condition. This, in turn, may have reduced the apparent size of the partitioning effect in the first two experiments.

Experiment 3 used a new category structure that circumvented these problems. The category structure, shown in Figure 6, in-

[1] Although Wickens (1989) suggested that repeated observations on the same participants may well be considered independent, a more conservative approach is to divide the underlying value of chi-square by 2, thus accounting for repetition of individuals across contexts (Wickens, 1989). Applying the more conservative $\chi^2/2$ yields a nonsignificant value of Cramer's coefficient ($\phi = .287$, $p = .163$).

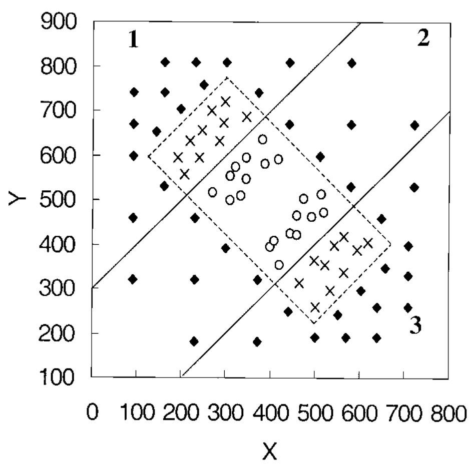

Figure 6. Stimulus space formed by dimensions X and Y used in Experiments 3 and 4. Training instances were randomly chosen for each participant from within the dotted rectangle. Open circles and crosses show a representative set of training items belonging to Category A and B, respectively. The solid diamonds represent common transfer items. The two solid lines represent the true category boundaries. Numbers 1−3 refer to different diagnostic areas.

cluded two parallel ascending boundaries, with one category located in between the boundaries and the other one outside. As before, this implied that the set of learning instances could be divided into two clusters, each separated into two categories by its own local boundary. As in the first two experiments, context in the systematic-context condition identified the two clusters and, by implication, a component boundary of the category space. However, unlike the earlier studies, both component boundaries were ascending and context was not perfectly correlated with either of the two predictors. To verify that participants could learn an integrated representation of this category structure, Experiment 3 again included a randomized-context condition.

Method

Participants

Forty-eight undergraduate volunteers from the University of Western Australia participated voluntarily in exchange for partial course points. An equal number of participants were randomly assigned to the systematic-and the randomized-context condition.

Stimuli

The boundaries in Figure 6 are described by $|Y - X - 100| = 200$. As in the first two experiments, the category task was instantiated as two hypothetical species of tropical fish that had to be classified on the basis of “density” (X), “depth” (Y), and “season” (context).

Training items. Unlike the first two experiments, each participant received a new set of 40 unique training instances that were randomly sampled from within the dotted rectangle in Figure 6 (not including locations in the immediate vicinity of the positive diagonal or the two boundaries). The figure includes a representative sample of items, with open circles and crosses denoting membership of Category A and B, respectively. The set of 40 items was divided into two clusters of 20 stimuli located above (the upper cluster) and below (the lower cluster) the principal diagonal.

In the randomized-context condition, context cues were randomly assigned to training stimuli, subject to the constraint that each context label occurred with an equal number of items from both categories and both clusters. In the systematic-context condition, as in the preceding experiments, context did not predict category membership but predicted the cluster to which a training item belonged. That is, for a given participant, one level of context always occurred with stimuli within the upper cluster and the other with those in the lower cluster. Assignment of context labels to cluster was counterbalanced, and we refer to context generically as *upper* and *lower*.

Regardless of condition, training consisted of eight blocks of trials, each involving a different random sequence of the 40 items.

Transfer items. All participants were tested with the same set of 40 novel transfer stimuli, represented by filled diamonds in Figure 6. As in the first two experiments, there were two blocks of transfer trials, each with a different random order of stimuli. Each context appeared equally often in each block, and context alternated between blocks for each item.

Diagnosticity of transfer items. Transfer items came from three distinct areas in the category space, identified by number in Figure 6. As before, the effect of context was expected to differ between areas depending on whether or not knowledge is partitioned. If partitioning is absent, as would be expected in the randomized-context condition, then context should have no effect in any of the areas, and items in Area 2 should be classified as belonging to Category A, whereas those in Areas 1 and 3 should be classified as B.

The presence of partitioning would be indicated by a context-invariant tendency to classify items in Area 2 as belonging to Category A, accompanied by a strong context effect in Areas 1 and 3. Specifically, in the upper test context, items in Area 1 should be classified as belonging to B and those in Area 3 as belonging to Category A. Conversely, in the lower context, items in Area 1 should be classified as A and those in Area 3 as B.

Procedure

Details of the procedure, including presentation parameters and response keys, were the same as in the first two studies.

Results and Discussion

Training Performance

The training data mirrored the outcome of the first two experiments. Performance improved from .564 (Block 1) to .752 (Block 8) across training, and participants performed better in the systematic-context condition ($M = .717$) than in the randomized-context condition ($M = .652$). Accordingly, a 2 (condition) × 8 (block) between-within ANOVA revealed a main effect of block, $F(7, 322) = 26.981$, $MSE = 0.008$, $p < .01$, a marginal effect of condition, $F(1, 46) = 3.752$, $MSE = 0.105$, $p = 0.059$, and no interaction between these two variables, $F(7, 322) < 1$.

Transfer Performance

Knowledge partitioning. Table 4 suggests that, as expected, test context had no effect on classification probabilities in any area in the randomized-context condition. In the systematic-context condition, there was no effect of context in Area 2 (the middle region within the two component boundaries), whereas there were strong—and opposite—effects of context in Area 1 and Area 3. This pattern conformed to what would be expected under knowledge partitioning.

Table 4
Mean Probabilities of Category A Responses in Experiment 3 for Novel Transfer Items in Each Area and Test Context (Upper or Lower)

	Randomized context			Systematic context		
Group	Area 1	Area 2	Area 3	Area 1	Area 2	Area 3
Overall						
Upper	.10	.71	.20	.12	.65	.50
Lower	.12	.70	.22	.41	.66	.14
True boundary						
Upper				.05	.72	.19
Lower				.05	.58	.06
Knowledge partitioning						
Upper				.14	.74	.88
Lower				.92	.78	.12

The corresponding 2 (condition) × 2 (test context) × 3 (area) between-within ANOVA revealed a highly significant main effect of area, $F(2, 92) = 107.5$, $MSE = 0.06$, $p < .01$. In addition, the main effect of condition approached significance, with $F(1, 46) = 3.37$, $MSE = 0.11$, $p = .073$, whereas there was no discernible effect of test context, $F(1, 46) < 1$. Two of the three two-way interactions were also significant: Condition × Area, $F(2, 92) = 4.3$, $MSE = 0.06$, $p < .05$, and Test Context × Area, $F(2, 92) = 11.97$, $MSE = 0.05$, $p < .01$.

Finally, the overarching three-way interaction involving all experimental variables was significant, $F(2, 92) = 12.16$, $MSE = 0.05$, $p < .001$, and was further explored by separate two-way ANOVAs in each condition. In the randomized-context condition, there was no Test Context × Area interaction, $F(2, 46) < 1$, whereas in the systematic-context condition that interaction was highly significant, $F(2, 46) = 18.27$, $MSE = 0.07$, $p < .001$. As suggested by the table, the Test Context × Area interaction in that condition involved significant effects of context in Areas 1 and 3, $F(1, 23) = 13.62$, $MSE = 0.07$, $p < 0.01$, and $F(1, 23) = 21.55$, $MSE = 0.07$, $p < .001$, respectively, but not in Area 2, $F(1, 23) < 1$.

In extension of the first two experiments, this finding suggests that participants applied different partial knowledge in different contexts. One critical difference from the earlier studies was the apparent symmetry of the effect: The size of the context effect was virtually identical for Areas 1 and 3, presumably because both partial boundaries were ascending and did not differ in difficulty.

However, as in the earlier studies, the absolute magnitude of the effect warranted further exploration. Specifically, if people used context-appropriate partial boundaries, then items in Areas 1 and 3 should be consistently classified as belonging to Category A in the lower and upper contexts, respectively. Indeed, given their greater distance from the boundary, items in those areas should be even more uniformly classified as A than items in the central area, some of which may on occasion be considered to lie on the other side of a partial boundary. In actual fact, however, the probability of classification was near 0.5 in both areas. Similar to the earlier experiments, this could have either reflected chance performance or the fact that some individuals partitioned their knowledge whereas others did not. These alternatives were teased apart by the individual differences analysis.

Individual differences. Using three predefined clusters consisting of (a) the correct parallel boundaries, (b) the upper partial boundary, and (c) the lower partial boundary, individual response profiles were entered into a *k*-means cluster analysis for each condition and test context separately. Several comments apply to the results, which are shown in Table 5.

First, in contrast to Experiments 1 and 2 (see Tables 2 and 3) there were very few cases (only 3 people) that escaped identification as using one of the three strategies. This further confirmed that the current category structure eliminated some of the problems associated with the earlier stimuli. Second, in the randomized-context condition, almost nobody used partial boundaries, whereas a considerable proportion of participants uniformly applied the true boundary across both contexts. The negligible value of Cramer's coefficient ($\phi = .059$) supports the conclusion that in this condition people's choice of strategy was unaffected by test context. In the systematic-context condition, by contrast, between a third (lower context) and half (upper) of all participants applied the context-appropriate partial boundary. The crucial role of context was affirmed by a highly significant Cramer's coefficient ($\phi = .677$, $p < .001$). Even with the conservative correction ($\chi^2/2$), the

Table 5
Number and Percentage of Participants Identified as Using a Particular Strategy by k-Means Cluster Analysis in Experiment 3

	Test context			
	Upper		Lower	
Cluster	*n*	%	*n*	%
Randomized-context condition				
True boundary	20	83	19	79
Upper boundary	3	13	4	17
Lower boundary	1	4	1	4
Unrecognized	0	0	0	0
Systematic-context condition				
True boundary	13	54	13	54
Upper boundary	11	46	0	0
Lower boundary	0	0	8	33
Unrecognized	0	0	3	13

Note. $N = 24$ in each condition.

association between context and cluster remained significant (ϕ = .479, $p < .012$).

To further illustrate the response pattern associated with use of the true and partial boundaries, Table 4 also shows the data for two subgroups of participants in the systematic-context condition. The panel labeled true boundary shows the classification probabilities from 11 participants who uniformly applied the true boundary in both contexts, whereas the knowledge partitioning panel shows the data of 8 participants who showed complete partitioning by using the appropriate partial boundary in both of the two contexts. This finding confirms that individuals who were identified as having partitioned their knowledge by a fairly loose criterion—exhibiting a generalization profile that was more similar to that expected under partitioning than to the true boundary—completely reversed their categorization judgment between contexts in Areas 1 and 3. Conversely, there was no effect of context in Areas 1 and 3 for those people who were identified as using the true boundary by an equally loose criterion. It follows that the overall classification probabilities arose from a probability mix of two distinct groups of participants, rather than uniform chance performance.

Experiment 4

In all experiments thus far, context did not directly predict classification. Nonetheless, context was uniformly found to gate performance, at least in a significant proportion of participants, when it predicted which partial boundary to apply to the task. The fourth and final experiment examined whether people were aware of the fact that context did not directly predict classification by adding a contingency rating task (e.g., Wasserman & Berglan, 1998; Williams, Sagness, & McPhee, 1994). The contingency rating task, presented after the conventional transfer test, asked people to provide a numeric estimate of the contingency between values of each predictor by itself and the outcome.

Experiment 4 included only a systematic-context condition, which was nearly identical to that of Experiment 3, except that the transfer test included several training items in addition to the usual novel stimuli.

Method

Participants

Twenty-four undergraduate volunteers from the University of Western Australia participated in exchange for partial course credit.

Stimuli

The category structure and stimuli were the same as in Experiment 3. The transfer test differed by additionally including 20 randomly chosen training instances (5 items from each category in each cluster).

In the final contingency rating task, one predictor was shown at a time, and ratings were obtained for four possible values of X (density: 200, 400, 600, and 800), four possible values of Y (depth: 100, 300, 500, and 700), and the two possible levels of context (summer and winter). Owing to a software error, some predictor values were repeated, and for a small number of participants some values of Y were omitted.

Procedure

The procedure was the same as in Experiment 3, except that the contingency test followed the transfer test. On each contingency rating trial, participants used the right- and left-arrow keys to move a vertical bar on the screen along a horizontal scale to indicate their rating.

The scale was marked by a category label at each end, with the assignment of labels counterbalanced across participants. Each end of the scale represented 100% certainty that an item with the given predictor value would belong to that category.

Results and Discussion

Training Performance

Performance improved across the eight training blocks from .547 (Block 1) to .8 (Block 8), which was confirmed to be significant by a one-way within-subjects ANOVA, $F(7, 161) = 23.45$, $MSE = 0.008$, $p < .01$. In absolute terms, performance was nearly indistinguishable from the systematic-context condition in Experiment 3.

Transfer Performance

Transfer responses to training items. The training items were presented in both contexts at test. One of those (called "congruent") occurred at training, whereas the other one was novel for the particular conjunction of X and Y values ("incongruent"). Strictly speaking, incongruent items were thus no longer old exemplars.

Classification probabilities for training items are shown by context and area in Table 6. In the congruent context, responses reflected knowledge of the true boundary, with items in the central Area 2 being predominantly classified as A and those in the exterior Areas 1 and 3 as B. In the incongruent context, by contrast, there was no evidence that people used the correct boundary. Categorization probabilities reversed for the two exterior areas, giving rise to an effect of context in Area 1, $F(1, 23) = 15.4$, $MSE = 0.08$, $p < .01$, and in Area 3, $F(1, 23) = 26.74$, $MSE = 0.09$, $p < .01$.

Transfer responses to novel test items. Table 7 shows that responses to the novel transfer items mirrored the results of the previous experiment. The relevant 2 (test context) × 3 (area) within-subjects ANOVA revealed a significant main effect of Area, $F(2, 46) = 33.05$, $MSE = 0.08$, $p < .001$, and a significant interaction between both variables, $F(2, 46) = 27.34$, $MSE = 0.05$, $p < .001$. The follow-up simple comparisons between test contexts were significant in Area 1, $F(1, 23) = 12.53$, $MSE = 0.05$, $p < .01$, and in Area 3, $F(1, 23) = 28.26$, $MSE = 0.08$, $p < .001$, but not in Area 2, $F(1, 23) < 1$.

The *k*-means cluster analysis showed that about half of the participants ($N = 10$) applied the true boundary in both contexts, whereas about a third of the participants ($N = 7$) showed complete and symmetric partitioning by using the upper boundary in the

Table 6
Mean Probabilities of Category A Responses in Experiment 4 for Old Transfer Items in Each Area and Test Context

Test context	Area 1	Area 2	Area 3
Congruent	.08	.75	.24
Incongruent	.39	.76	.69

Table 7
Mean Probabilities of Category A Responses in Experiment 4 for Novel Transfer Items in Each Area and Test Context (Upper or Lower)

Group	Area 1	Area 2	Area 3
Overall			
Upper	.11	.70	.56
Lower	.35	.66	.13
True boundary			
Upper	.07	.80	.19
Lower	.13	.72	.09
Knowledge partitioning			
Upper	.18	.70	.92
Lower	.82	.76	.19

upper context and the lower boundary in the lower context. Cramer's coefficient (ϕ = .676, p < .001) confirmed that context affected strategy choice. The effect persisted even with the conservative ($\chi^2/2$) adjustment: ϕ = .478, p < .012. The transfer responses to novel items for those two subgroups of participants are shown in the corresponding panels of Table 7.

Comparison of novel and training items. The two classes of items were compared by computing the differences between upper and lower test contexts (*not* between congruent and incongruent)[2] for each item class and each area. Those differences were entered into a 2 (item type: training vs. novel) × 3 (area) within-subjects ANOVA. The only significant effect was a main effect of area, with $F(2, 46) = 32.14$, $MSE = 0.19$, $p < .01$. Neither the effect of item type nor the interaction between both variables was significant (both Fs < 1). This finding implies that people considered training items presented in a new context at test as being completely novel and indistinguishable from a new pairing of values for X and Y.

An effect of memory for training instances was observed only when training items were presented in the same context at test: As shown in the earlier Table 6, it was only under those narrowly defined circumstances that classifications were compatible with the true bilinear category boundary.

Contingency Ratings

Contingency ratings were analyzed by expressing responses as the judged probability of an item belonging to Category A. Table 8 shows the average ratings for all tested levels of the three predictors (X, Y, and context) together with the normative values (computed for the overall set of training items). Responses are shown averaged across all participants and for the two major subgroups identified by the *k*-means analysis separately.

In order to determine whether judged contingencies departed from chance, the mean response at each level of each predictor was compared to .5 by a single-sample *t* test. Using a significance level of .05 for each test, we found that three predictor levels departed from chance; those are identified by asterisks in the table. With a more stringent Bonferroni adjustment to significance, none of the 10 responses differed from chance.

Thus, there was some suggestion that people were sensitive to the role of predictors X and Y, but their judged contingencies clustered relatively close to chance and were not nearly as extreme as the normative values. There was little evidence that the two groups of participants differed from each other. For the context variable, there was no evidence that people erroneously judged it to be predictive, as their responses were close to the normatively correct chance value (largest $t \approx 1.00$) and also did not differ between the two contexts: $t(23) < 1$, using all participants. The judged contingency of context also did not differ between groups of participants.

General Discussion

Summary of Results

All four experiments showed that people can learn to use context to gate use of partial categorization knowledge. In Experiments 3 and 4, the usage of partial knowledge was completely symmetrical. Experiment 2 additionally showed that when people partitioned their knowledge, they used each partial knowledge component as though the other one had not been acquired. That is, performance in each context was very similar to the performance of people in comparison groups who had been trained only on half of the stimulus space in that one context, suggesting that the partitioned components of knowledge were largely independent of each other. Finally, Experiment 4 confirmed that people were at least somewhat sensitive to the normative role of each predictor in isolation, including the fact that context by itself did not predict category membership.

Another consistent outcome was that training performance was better whenever knowledge partitioning was observed at transfer. Thus, in Experiments 1 and 3 performance in the systematic-context condition was significantly better than in the randomized-context condition.

The pattern of training and transfer performance parallels the related findings reported by Lewandowsky et al. (2002) in a function learning paradigm. This similarity of findings reinforces the close empirical and conceptual connection between those two modes of concept learning and extends the generality of the knowledge partitioning framework. However, before considering the wider implications of our results, several limitations must be acknowledged.

Limitations of the Present Studies

The first limitation concerns the magnitude of the knowledge partitioning effect. Across experiments, the proportion of participants who partitioned their knowledge when given the opportunity to do so ranged from 20% to at most 50%. A significant proportion of the remaining participants applied the correct complex boundary, whereas still others were not clearly identifiable as using one or the other strategy. This finding suggests that knowledge parti-

[2] For training items presented at transfer, the incongruent context differed between areas. In Area 3, it was the upper context that was incongruent at test, because no item below the principal diagonal appeared in that context during training. Conversely, in Area 1, it was the lower test context that was incongruent. To simplify comparison with novel transfer stimuli, this analysis distinguished between upper and lower contexts irrespective of congruency.

Table 8
Contingency Rating Results in Experiment 4 and Significance Tests for Deviation From Chance (.50)

	X				Y				Context	
Group	100	300	500	700	200	400	600	800	Upper	Lower
Normative	.00	.52	.53	.00	.00	.52	.52	.00	.50	.50
Overall	.47	.61*	.48	.47	.59	.61*	.37	.28**	.45	.54
KP	.47	.55	.48	.54	.64	.68	.36	.20*	.40	.56
TB	.54	.66	.54	.54	.63	.59	.36	.35	.44	.57

Note. X and Y refer to the stimulus dimensions. KP = participants who showed complete partitioning by using the appropriate partial boundary in each of the two contexts; TB = participants who uniformly applied the true boundary in both contexts.
* $p < .05$. ** $p < .01$.

tioning is but one of several ways in which people may choose to master a complex categorization task.

A second potential limitation concerns the use of pseudocontinuous predictors involving numeric labels. Although this is not without precedent in categorization research (e.g., Erickson & Kruschke, 1998; Lewandowsky et al., 2000), the practice has been criticized by Nosofsky and Johansen (2000). It is therefore important to note that Yang and Lewandowsky (2002) obtained knowledge partitioning in two studies with perceptual stimuli similar to those used by Nosofsky and Johansen (2000).

Relationship to Other Findings

Complete Partitioning Versus Limited Heterogeneity

The present results go beyond the limited heterogeneity reported by Aha and Goldstone (1992) because those people who partitioned their knowledge were found to use each partial knowledge component as though the other had not existed. This situation resulted in contradictions and inconsistencies—that is, opposing classifications of the same item in different contexts—that were absent in the study by Aha and Goldstone (see also Erickson & Kruschke, 2001). We are not aware of any further related precedents in categorization; however, Kalish et al. (2001) reported similar contradictions when people learned functions that contained exceptions. Specifically, one of the experiments used a linearly increasing function that contained three gaps, each of which contained an outlying observation. The three outliers, in turn, defined a (sparse) linearly *decreasing* function. People were found to extrapolate either on the basis of the increasing or the decreasing function. On any given transfer trial, one or the other learned function was chosen to govern the response, without any apparent attempt at integration between the two conflicting response alternatives. These trial-by-trial contradictions resemble the context-specific opposing classifications in the present experiments.

Correlated Features in Categorization

The context effects in our studies occurred when the experimental manipulation introduced a correlation with the other predictor(s). There has been much research involving correlated predictors, and it is known that when people learn about a category, they not only learn about each feature independently, but they also learn about the relations between them (Medin, Altom, Edelson, & Freko, 1982; Thomas, 1998). This can express itself in people's ability to predict the absent feature values of novel stimuli (e.g., Anderson & Fincham, 1996) or in the ability to predict category membership from correlations among predictors (Medin et al., 1982).

Most relevant here is another consequence of correlations among predictors, namely that people sometimes learn to rely on a normatively irrelevant predictor. For example, Gluck and Bower (1988) trained people to diagnose fictitious patients as having one of two diseases (i.e., category) on the basis of four symptoms, each of which could take on two values (e.g., fever vs. no fever). On a final test, participants had to judge what proportion of patients with a particular pattern of symptoms would suffer a given disease. People's predictions were affected by the relationship among predictors in addition to the relationship between each predictor and the categories. Specifically, a critical predictor (P_c) that was normatively irrelevant (because each disease occurred equally often in its presence) was erroneously judged to be predictive of one of the diseases (call that *R*). This finding can be taken to result from a correlation among predictors because when P_c occurred with disease *R* during training, it tended to occur on its own or was accompanied by predictors whose associations to the same category were relatively weak.

Our results extend the findings of Gluck and Bower (1988) in several ways: First, in our studies the conditional probability with which a context occurred given a category was equal across categories—that is, $P(\text{Context}_x|\text{Category}_1) = P(\text{Context}_x|\text{Category}_2)$. This stands in contrast to Gluck and Bower's (1988) study, where the conditional probability with which P_c occurred was greater for category *R* than the other (P_c was normatively irrelevant only because the base rates of the two categories differed considerably; hence *R* stands for "rare" category.) Second, in the study by Gluck and Bower people were shown to rely on the irrelevant cue at test when it was presented in isolation. In the present experiments, by contrast, people recognized the irrelevance of context on its own (contingency ratings in Experiment 4) but nonetheless used it when it co-occurred with other predictors at test.

Theoretical Implications

Thus far, we have adopted a rule-based perspective for presentation and discussion of the results. We now examine potential alternative accounts before we develop the rule-based framework further.

Nonlinear Separability

Feature-based theories [also known as component cue models; e.g., Gluck, 1992; Gluck & Bower, 1988; attention to distinctive input (ADIT), Kruschke, 1996; attentional connectionist model (ACM), Shanks, 1992] explain categorization by relying on direct connections between each feature and the possible responses. In consequence, all component cue models necessarily rely on linear (or quasilinear) decision planes that are orthogonal to the vector of weights between the features and each output unit.

It turns out that linear decision planes are incompatible with knowledge partitioning. Consider the stimulus space from the first two studies (e.g., Figure 2). The observed opposing classifications to transfer items in Areas 2 and 4 in different contexts imply that people's representation of the stimulus space was three-dimensional, with the third dimension representing context. Crucially, there is no single linear decision plane within that three-dimensional space that can predict the outcome. Instead, the opposing classifications in Areas 2 and 4 mandate a nonlinear decision plane that is a projection onto the positive X–Y diagonal in one context and a projection onto the negative diagonal in the other. This requires that the plane be twisted by 90° along its trajectory through the context dimension.[3]

Instance Models and Attention to Context

Unlike feature-based theories, instance models such as GCM (e.g., Nosofsky, 1986) or ALCOVE are not restricted to linear decision boundaries. One characteristic of instance models is that the dimensions used for categorization may attract different amounts of *attention*. Dimensions that are particularly important for classification attract much attention and are thus stretched, whereas irrelevant dimensions shrink relative to the others (e.g., Kruschke, 1992). Accordingly, there is every possibility that instance models could accommodate our results if the context dimension was stretched relative to the other two. However, what remains unclear is *why* an instance model would pay attention to context in the first place.

In the GCM, attention weights are sometimes assumed to be optimally distributed across dimensions (Kruschke & Johansen, 1999), but they are also often estimated as free parameters (e.g., Nosofsky & Johansen, 2000). It follows that the GCM may handle the present results if the attention to context is a free parameter and estimated to be suitably large.

ALCOVE (Kruschke, 1992), by contrast, uses error-driven learning to shift attention to relevant dimensions during training. That is, instead of estimating parameters or assuming optimality, the distribution of attention is generated by the model itself through iterative error reduction. Kruschke (1992) showed that ALCOVE can learn to attend to correlated predictors, even if they are individually nonpredictive, by modeling the results of a study by Medin et al. (1982). At first glance, this suggests that ALCOVE may also be capable of handling the present results.

However, there is an important difference between the experiments of Medin et al. (1982) and those reported here. In the case of Medin et al., the two correlated predictors jointly permitted perfect classification of the training items, whereas an alternative strategy involving two other uncorrelated predictors permitted only 75% accuracy. Hence, the error-driven learning in ALCOVE eventually shifted attention away from the uncorrelated predictors to the correlated pair. In our experiments, by contrast, attention to dimensions X and Y alone could completely eliminate error. Conversely, context alone could not afford correct classification in our experiments, which implies that the error-driven learning in ALCOVE should eventually learn to pay no attention to context. In confirmation, Yang and Lewandowsky (2002) showed that ALCOVE cannot predict knowledge partitioning when attention is distributed equally across all dimensions at the outset of learning. ALCOVE could predict partitioning only when learning commenced with context already attracting considerably more attention than the other two dimensions. It follows that ALCOVE, similar to the parameter estimation likely required for GCM, must rely on processes outside the model's explanatory scope to handle knowledge partitioning.

Configural Processing in Categorization

In Experiment 4, training items presented at transfer were classified according to the true boundary only when contexts at training and test were congruent. When test context was incongruent, old conjunctions of X and Y were classified exactly like novel transfer stimuli. One interpretation of this sensitivity to partial change is that people considered the stimuli in a holistic or "configural" manner (e.g., Shanks, Charles, Darby, & Azmi, 1998; Shanks, Darby, & Charles, 1998). Configural processing refers to the coding of "stimuli as indivisible elements rather than as combinations of features or elements" (Shanks, Darby, & Charles, 1998, p. 138). It stands in contrast to elemental processing, which refers to the decomposition of a complex stimulus into its constituent elements or features (Williams et al., 1994).

Evidence for configural processing comes from a variety of sources. For example, Williams et al. (1994) showed that the classic blocking effect (e.g., Kamin, 1969) can be eliminated by configural processing. Similarly, Shanks, Darby, and Charles (1998) showed that the effect of interference also depends on the extent to which processing is configural. In their study, the initially learned associations between two predictors and their respective outcomes (e.g., P → A and Q → B) were found to be intact after interpolated learning during which an opposing outcome (A–) was associated with a compound stimulus involving the initial predictors (e.g., {P, Q} → A–). This contrasts to the known fact that interference would result if a single predictor were reassigned during interpolated learning (e.g., P → A replaced by P → A–). Shanks et al. interpreted their results as evidence of configural processing, such that the compound stimulus {P, Q} is processed independently of its constituent parts (P and Q).

[3] The asymmetry with which people partitioned their knowledge in Experiments 1 and 2 may mandate further contortions of this surface without however eliminating its nonlinearity.

At a surface level, those precedents relate to our Experiment 4, which found that disruption of the configurality of a training instance by presenting it in a new context rendered it indistinguishable from completely novel transfer items. However, further examination of the implications of configural processing reveals that this approach, too, is incompatible with our results.

If people engaged in configural processing, they could have considered the stimuli as one of two possible compounds. First, each item may have formed an integrated triplet consisting of a unique conjunction of X, Y, and context. Because configural processing considers every compound as an indivisible unit, a training instance would then become a "new" item when tested in an incongruent context. This was seemingly observed in Experiment 4. However, other aspects of the data from that experiment rule out this representation. Specifically, consider two training triplets presented in context C_1, one located in Area 2 (call that $\{X_1, Y_1, C_1\}$) and another one in either Area 1 or 3 (call that $\{X_2, Y_2, C_1\}$). If those items were processed as indivisible configural stimuli, what would be the effect of testing them in the incongruent context (i.e., as $\{X_1, Y_1, C_2\}$ and $\{X_2, Y_2, C_2\}$)? Configural theories should expect the extent of disruption (or, equivalently, remaining generalization; Gluck, 1991; Shanks & Gluck, 1994) to be equal for both items. However, in Experiment 4 a change in context did not affect responses to training items in Area 2 (i.e., $\{X_1, Y_1, C_1\} = \{X_1, Y_1, C_2\}$), whereas it had a large effect in Areas 1 and 3 (hence, $\{X_2, Y_2, C_1\} \neq \{X_2, Y_2, C_2\}$). This appears incompatible with a configural representation involving unified triplets.

A second possibility is that, instead of forming triplets, people treat X and Y as a configural stimulus (call that P) that is accompanied by a (componential) context. A training item in context C_1 can thus be represented as $P + C_1$. Supposing that this item belongs to Category A, the learned association can be represented as $P + C_1 \rightarrow A$. Suppose furthermore that there is another compound of X and Y (call that Q) that is associated with the other category in the other context (i.e., $Q + C_2 \rightarrow B$). It follows that when old items are presented in an incongruent test context (i.e., $P + C_2$ or $Q + C_1$), they should be classified into either category with equal probability. This is because the association of P to Category A is equivalent to the association of C_2 to B, thus predicting either category with equal probability for $P + C_2$ (and $Q + C_1$). Our data disconfirm this prediction because people reversed their classification judgment for training items when the context became incongruent at test. It follows that people did not form a partially configural representation.[4]

Gated Modules

We now consider our preferred explanation within the knowledge partitioning framework, namely that people learned to associate partial rules with each level of context and applied each context-specific rule without further consideration of other available knowledge.

Our approach derives from the attention to rules and instances in a unified model (ATRIUM) theory of Erickson and Kruschke (1998, 2001), which uses a so-called mixture-of-experts approach to combine instance memory with rules. ATRIUM contains one module for instances (implemented by a version of ALCOVE) and one module for each linear one-dimensional rule. Learning adjusts the relative importance of these modules (instances vs. rules) for each exemplar, the location of the decision rule on each of the categorization dimensions, the amount of attention devoted to each dimension, and the associations between each exemplar and possible responses. Support for ATRIUM includes the finding that generalization to novel stimuli is governed by one or the other module, depending on the relative proximity of a novel item to exceptional stimuli (Erickson & Kruschke, 1998). An extended version of ATRIUM has been shown to handle Aha and Goldstone's (1992) results (Erickson & Kruschke, 2001).

Nonetheless, in its current state of development, ATRIUM cannot be immediately applied to our data because its rule module(s) are limited to single dimensions and thus cannot handle the diagonal boundaries used here. One possible solution to this problem emphasizes that the dimensions in ATRIUM are *psychological* rather than physical: In consequence, participants in our experiments may have effectively reduced the dimensionality of the task by noting the additive or subtractive relationship of X and Y (M. A. Erickson, personal communication, December 4, 2001). For example, in Experiments 1 and 2, the ascending partial boundary could be reexpressed by a single variable $Z = Y - X$ with a boundary at $Z = 100$. However, even with this auxiliary assumption, it is still unclear how ATRIUM would handle the simultaneous coexistence of partial rules involving the same predictors. That is, in the present case, predictors X and Y would need to be simultaneously reexpressed in two different ways (Experiments 1 and 2) or twice in the same subtractive relationship but with different boundary placement (Experiments 3 and 4). The latter would present a particular difficulty to ATRIUM because it has never posited more than one rule module per dimension (M. A. Erickson, personal communication, December 8, 2001). We therefore conclude that ATRIUM, despite being a closely allied precedent, cannot be applied to our results without modification.

Our preferred approach, then, retains ATRIUM's architecture of separate modules that compete for a categorization response. Rule modules can be one-dimensional or two-dimensional, and a number of such candidate modules may be available during learning. Each rule module would be associated with a set of trained instances, which in the systematic-context condition, might all share a common context. At test, the rule module that was associated with the most similar training items is chosen to drive the response. The most critical difference between our proposed approach and ATRIUM concerns the way in which information from different modules is considered. In ATRIUM, modules cooperate by passing "a proportion of the activation from both sets of . . . nodes to a final output set of category nodes" (Erickson & Kruschke, 1998, p. 118). In our approach, modules would instead be chosen probabilistically and in a discrete manner. Once chosen, a module would exclusively drive the response. This is to accommodate the fact that people seemingly ignore knowledge encapsulated in other modules (Experiment 2).

[4] An objection to this argument might assume that C is somehow more salient than Q or P, in which case categorization may indeed reverse with context. This objection runs counter to the fact that C on its own, unlike P or Q, is never a valid predictor during training, which renders it unlikely to be more salient than the accompanying valid compound predictor.

The approach just described appears to be a promising candidate for formal implementation. It is noteworthy that Kalish et al. (2001) presented a computational model of function learning that is similar in many respects. Their model, known as population of linear experts (POLE), can handle most existing data on function learning, including the partitioning effects reported by Lewandowsky et al. (2002), by assuming that multiple candidate modules compete for a response. Once a candidate is chosen, knowledge in other modules is ignored—this exclusivity of access may be an essential feature of any theory capable of handling the growing body of evidence for knowledge partitioning.

Knowledge Partitioning: Current Status

The partitioning of knowledge has, by now, been observed in several different settings. In the domain of expertise, Lewandowsky and Kirsner (2000) asked experienced fire commanders to predict the spread of simulated wild fires from a set of physical predictors. When the two primary predictors (wind and slope of terrain) were in opposition—thus creating uncertainty about the direction of spread—the experts' predictions were found to depend on a physically irrelevant context. When a fire was presented as one that had to be brought under control, it was expected to spread with the wind. The identical fire presented as a back burn led to the opposite prediction, that it would spread with the slope and into wind. (Back burns are lit by fire fighters in the path of the main fire to starve it of fuel; although the two types of fire differ with respect to their origin, both obey the same rules of physics.) Under the knowledge partitioning view, this finding can be explained by assuming that different parcels of knowledge were accessed and used in each of the two contexts.

Similar instances of contradictory behaviors have been reported in the domain of so-called street mathematics (e.g., Carraher, Carraher, & Schliemann, 1985; Nunes, Schliemann, & Carraher, 1993). Participants in those studies—Brazilian street vendors, carpenters, and cooks with minimal formal schooling—were sometimes shown to use different and contradictory ways to solve an otherwise identical problem in different contexts (Schliemann & Carraher, 1993).

Further direct experimental support for knowledge partitioning was provided by Lewandowsky et al. (2002) and Kalish et al. (2001) in the function learning studies discussed earlier. The present experiments, then, provided a third setting in which the fractionation of knowledge—and the contradictions it entails—has been observed.

At a theoretical level, one implication of the only existing computational model of knowledge partitioning (POLE: Kalish et al., 2001) is particularly intriguing. We noted above that POLE can accommodate a wide variety of results from function learning, including knowledge partitioning, by assuming that candidate modules compete for a response. The implication of this mechanism is that knowledge partitioning formed the basis of *all* function learning in POLE. That is, each candidate module consisted of a linear function of a given slope and intercept that was associated during learning to a set of stimuli (similar to ATRIUM; Erickson & Kruschke, 2001). It follows that POLE partitions *any* but the simplest linear function by associating different subsets of stimuli to different partial functions; for example, a quadratic function is learned by piecing together a number of linear modules that are associated with different sets of stimulus magnitudes. POLE thus considers knowledge partitioning not as a curious exception but as a fundamental attribute of function learning. Likewise, the gated-modules explanation for our data sketched above considers knowledge partitioning to be fundamental to all aspects of category learning.

Conclusion

This article makes several empirical and theoretical contributions: Empirically, we showed that (a) people can be sensitive to the correlation between a context cue and other predictor(s), even if the other predictors alone permit perfect classification of training items; (b) a significant proportion of people use context to partition their knowledge; and (c) when people partition their knowledge, each knowledge parcel is used as though others do not exist.

At a theoretical level, (a) we extended the knowledge partitioning framework from expertise and function learning to categorization; (b) we showed that our results challenge a number of current theories of categorization; and (c) to account for the data, we proposed a mixture-of-experts approach that represents partial knowledge in multiple modules. Modules may contain mutually contradictory information and, once chosen, are applied exclusively and without consideration of knowledge in other modules.

References

Aha, D. W., & Goldstone, R. L. (1992). Concept learning and flexible weighting. In J. K. Kruschke (Ed.), *Proceedings of the fourteenth annual conference of the Cognitive Science Society* (pp. 534–539). Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Fincham, J. M. (1996). Categorization and sensitivity to correlation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 259–277.

Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 33–53.

Ashby, F. G., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception & Performance, 16,* 598–612.

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 50–71.

Busemeyer, J. R., Byun, E., Delosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts, and categories* (pp. 405–435). Cambridge, England: Psychology Press.

Carraher, T. N., Carraher, D. W., & Schliemann, A. D. (1985). Mathematics in the streets and in schools. *British Journal of Developmental Psychology, 3,* 21–29.

Delosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 968–986.

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

Erickson, M. A., & Kruschke, J. K. (2001). *Multiple representations in inductive category learning: Evidence of stimulus- and time-dependent representation.* Manuscript submitted for publication.

Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. *Psychological Science, 2,* 50–55.

Gluck, M. A. (1992). Stimulus sampling and distributed representation in adaptive network theories of learning. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *Essays in honor of William K. Estes: Vol. 2. From learning processes to cognitive processes* (pp. 169–199). Hillsdale, NJ: Erlbaum.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General, 117,* 227–247.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 418–439.

Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2001). *Population of linear experts: Knowledge partitioning and function learning.* Manuscript submitted for publication.

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

Kruschke, J. K. (1996). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 3–26.

Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 1083–1119.

Lamberts, K. (1994). Flexible tuning of similarity in exemplar-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1003–1021.

Lewandowsky, S., Kalish, M., & Griffiths, T. L. (2000). Competing strategies in categorization: Expediency and resistance to knowledge restructuring. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1666–1684.

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). Expert knowledge is not always integrated: A case of cognitive partition. *Memory & Cognition, 28,* 295–305.

Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 37–50.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20,* 158–190.

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

Nosofsky, R. M., Clark, S. E., & Shin, J. H. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 282–304.

Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review, 7,* 375–402.

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

Nunes, T., Schliemann, A. D., & Carraher, D. W. (1993). *Street mathematics and school mathematics.* Cambridge, England: Cambridge University Press.

Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15,* 346–378.

Schliemann, A. D., & Carraher, D. W. (1993). Proportional reasoning in and out of school. In P. Light & G. Butterworth (Eds.), *Context and cognition: Ways of learning and knowing* (pp. 47–73). Hillsdale, NJ: Erlbaum.

Shanks, D. R. (1992). Connectionist accounts of the inverse base-rate effect in categorization. *Journal of Neural Computing, Artificial Intelligence & Cognitive Research, 4,* 3–18.

Shanks, D. R., Charles, D., Darby, R. J., & Azmi, A. (1998). Configural processes in human associative learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1353–1378.

Shanks, D. R., Darby, R. J., & Charles, D. (1998). Resistance to interference in human associative learning: Evidence of configural processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 136–150.

Shanks, D. R., & Gluck, M. A. (1994). Tests of an adaptive network model for the identification and categorization of continuous-dimension stimuli. *Connection Science, 6,* 59–89.

Thomas, R. D. (1998). Learning correlations in categorization tasks using large, ill-defined categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 119–143.

Vandierendonck, A. (1995). A parallel rule activation and rule synthesis model for generalization in category learning. *Psychonomic Bulletin & Review, 2,* 442–459.

Wasserman, E. A., & Berglan, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgment: The role of within-compound associations. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 51*(B), 121–138.

Wickens, T. D. (1989). *Multiway contingency tables analysis for the social sciences.* Hillsdale, NJ: Erlbaum.

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 694–709.

Yang, L.-X., & Lewandowsky, S. (2002). *Knowledge partitioning in categorization: Constraints on exemplar models.* Manuscript in preparation.

Received January 8, 2002
Revision received December 17, 2002
Accepted December 17, 2002 ■