

Competing Strategies in Categorization: Expediency and Resistance to Knowledge Restructuring

Stephan Lewandowsky, Mike Kalish, and Thomas L. Griffiths
University of Western Australia

The authors investigated people's ability to restructure their knowledge when additional information about a categorization task is revealed. In 2 experiments, people first learned to rely on a fairly accurate (but imperfect) predictor. At various points in training, a complex relationship between 2 other predictors was revealed in a schematic diagram that could support perfect performance. In Experiment 1, people adopted the complex strategy when it was revealed at the outset but were unable to restructure their knowledge after the expedient predictor had been learned. In Experiment 2, expedient knowledge persisted even with an adaptive display. The persistence of expedient knowledge is explained by associative blocking of potential alternative cues. A 3rd experiment analyzed the strategies people use with and without the diagram. The study confirmed that the diagram, when presented at the outset, significantly alters people's approach to the task.

Knowledge restructuring is often observed when people develop new skills, both in the real world and in the laboratory; paradoxically, expertise at these skills in some cases also seems to prevent restructuring of knowledge. Little is known about the processes involved in knowledge restructuring, and our investigation formed a first step toward describing those processes in order to better understand when restructuring occurs and when it is resisted.

Specifically, we examined people's ability to adapt and restructure their knowledge when previously concealed information about a categorization task was revealed halfway through training. The task could be performed either by relying on an expedient single predictor that permitted fairly accurate (but imperfect) performance or by learning one of several complex (but potentially perfect) strategies involving two predictors. It was known from previous research that people can use a complex strategy for this task if the relationship between predictors is revealed in a schematic diagram before training (Lewandowsky, Dunn, Kirsner, & Randell, 1997). Here, we asked two questions: First, can people use the diagram to restructure their knowledge after having learned to rely on the expedient predictor, and what are the processes underlying that restructuring? Identification of those processes requires exploration of the second question, namely what exact strategy or rule people used with and without the diagram. We report three experiments in which the diagram was presented at various points in training.

To foreshadow our principal conclusions, when the expedient strategy was learned first, people subsequently could not be encouraged to adopt a more complex strategy, regardless of whether the underlying relationship between predictors was revealed in a static diagram (Experiment 1) or through an adaptive display (Experiment 2). It was only when the static diagram was presented at the outset that people adopted a complex strategy involving two predictors (Experiments 1 and 3). In contrast to the persistence of expediency, the complex strategy was, after extended practice, spontaneously abandoned in favor of the expedient approach (Experiment 1). We explain this asymmetry of knowledge restructuring by appealing to two associative learning principles, blocking and retrospective revaluation, that underlie the persistence of expediency and the gradual disposal of the complex alternative, respectively. Our results provide a first glimpse into the processes underlying knowledge restructuring.

Knowledge Restructuring

The acquisition of skill and expertise often entails the repeated restructuring of knowledge, defined here as a shift to a new set of rules, strategies, or declarative knowledge to perform a task without overt changes to the stimuli or the learning environment. A review of the literature reveals an unresolved conflict between, on the one hand, reports of flexible and successful knowledge restructuring and, on the other hand, instances in which people resisted restructuring.

Successful Knowledge Restructuring

In an early report, spontaneous knowledge restructuring was implied by discontinuities between periods of learning. Bryan and Harter (1899, cited in Shiffrin, 1996) observed the fluency of Morse telegraph operators across extended practice and found that operators seemingly first learned to send and receive individual letters, then words, and then entire sequences of words. These periods of learning were separated by plateaus during which per-

Stephan Lewandowsky, Mike Kalish, and Thomas L. Griffiths, Department of Psychology, University of Western Australia, Nedlands, Australia.

Preparation of this article was facilitated by Large Research Grants A79600016 and A79941108 from the Australian Research Council. We thank Janet Phang and Lih Yi Soong for assistance during data collection and analysis, and Mike Mundy and Christina Chandler for help during preparation of this article.

Correspondence concerning this article should be addressed to Stephan Lewandowsky, Department of Psychology, University of Western Australia, Nedlands, W.A. 6907, Australia. Electronic mail may be sent to lewan@psy.uwa.edu.au.

formance changed little, suggesting that existing knowledge had to be consolidated and restructured for further learning to occur.

More recently, in a longitudinal study involving undergraduate science students, Pearsall, Skipper, and Mintzes (1997) examined the repeated restructuring of emerging knowledge in biology. Extensive restructuring tended to occur early in the semester, followed by periods of primarily incremental learning. Davies (1994) provided a similar picture of the development of computer programming expertise.

Turning to controlled laboratory experiments, Staszewski (1988) trained a participant to become expert on mental multiplication problems involving multidigit operands (e.g., 52×44) using an explicitly provided strategy. After 500 sessions of practice, the expert was taught another, computationally more efficient strategy. The new technique almost immediately gave rise to considerably faster solution times, suggesting rapid restructuring of knowledge. A quantitative reexploration of these data by Delaney, Reder, Staszewski, and Ritter (1998) showed that performance on each strategy was described by its own learning function, confirming that knowledge restructuring was both rapid and extensive.

Resistance to Knowledge Restructuring

The preceding results stand in conflict to reports of resistance to knowledge restructuring during the acquisition of expertise, for example as observed by Lesgold et al. (1988) with radiologists. Radiological diagnosis can be viewed either as a perceptual pattern recognition task, in which some complex set of features must be detected to enable identification of a disease, or as a cognitive inference task, in which it must be determined which disease is consistent with an identified set of features. Lesgold et al. identified perceptual learning as occurring early during expertise acquisition, whereas cognitive learning, which relies on successful perceptual feature extraction, necessarily occurs during later stages of training. Accordingly, Lesgold et al. observed a nonmonotonic relationship between diagnostic accuracy and level of expertise; accuracy suffered for some time after knowledge had been restructured from a perceptual to a cognitive strategy. According to Lesgold et al., the decline in accuracy occurred because the new conceptual approach had to compete with strong prior perceptual learning. Lesgold et al. suggested that those "points of impasse" (p. 340) at which earlier learning stands in conflict to new and improved ways to approach a task are common during the acquisition of expertise.

In another case, prolonged resistance to knowledge restructuring was evident when experts with upward of 15 years experience were found to use an expedient but imperfect approach to a highly domain-relevant task (Lewandowsky & Kirsner, 2000). The task used in that study, prediction of fire spread by expert bush-fire commanders, was adapted for the present experiments and thus deserves to be introduced in some detail.

The spread of bush fires is primarily determined by two physical variables, wind and slope of the terrain (Albini, 1984). All other variables being equal, fires tend to spread with the wind and uphill. Typically, with moderate to strong winds, the slope of the terrain either accelerates the wind-driven fire (when slopes are uphill with respect to wind direction) or slows its spread (downhill slopes). Psychologically interesting complexities arise with light downhill

winds, in which case the two predictor variables are in competition. The outcome—whether fire spreads downhill with the wind or uphill against the wind—then depends on the relative strengths of the predictors.

Lewandowsky and Kirsner (2000) found that under those circumstances the experts' predictions depended on an additional variable, the physically irrelevant problem context. When a fire was presented as one that had to be brought under control, experts uniformly expected it to spread with the wind. When an identical fire was presented as a "back burn," experts predicted the reverse, namely that the fire would spread uphill and into wind. Back burns are fires that are lit by firefighters in the path of the advancing main fire to starve it of fuel. Back burns differ from to-be-controlled fires in several ways; being lit by firefighters they rarely roar out of control, they tend not to be used in strong wind conditions, and they tend to be lit in proximity of natural fire breaks. Nonetheless, back burns obey the same laws of physics as to-be-controlled fires.

The finding by Lewandowsky and Kirsner (2000) that, depending on the problem context, experts make two opposing predictions under physically identical conditions, arguably reflected an expedient strategy because the frequency of the two opposing fire directions differs with context: Most fires that need to be brought under control are wind-driven overall (because fires tend to cover areas that include slopes of various orientations), whereas back burns are more likely to be slope-driven (because ignition points are selected to maximize the likelihood of a back burn advancing toward the wind-driven to-be-controlled fire). Hence, on a purely statistical basis, an expedient focus on context would predict fire direction with reasonable accuracy and without consideration of the more complex interaction between wind and slope.

What might have caused this persistent expert expediency? Given the strong emphasis on the physics of fire during training, why did the expert fire commanders rely on context? One possibility is that the interaction between wind and slope is simply too complex to be learned, even after much training and extensive experience. This possibility was ruled out by Lewandowsky et al. (1997), who showed that novice participants can learn to rely on slope and wind when their relationship was visualized at the outset of training. It is more likely, therefore, that the expert bush-fire commanders relied on context because initial learning of that expedient but imperfect rule prevented subsequent knowledge restructuring.

Observations in other domains support this assertion. For example, Lightfoot and Shiffrin (1992), using a visual search task, trained participants to develop unitized representations of novel visual stimuli comprising several features (horizontal and slanted lines). When the task was subsequently simplified to permit target identification on the basis of a single feature match, no immediate improvement in performance was observed. Instead, people showed gradual further improvement along the same initial learning curve. Similarly, in a probabilistic prediction task, Edgell and Morrissey (1987) observed the persistence of expedient initial learning. In their task, a single cue was initially predictive of the outcome, and people continued to rely on that cue despite subsequent opportunities for knowledge restructuring when the configuration of two cues became predictive.

An Unresolved Conflict

Overall, there is clear evidence that knowledge restructuring occurs during skill acquisition. Paradoxically, there is also clear evidence to the contrary, that people find restructuring of their knowledge difficult or resist it altogether. Nothing is known about the cognitive processes underlying restructuring or resistance against it. In particular, it is currently not possible to anticipate, for any given task or paradigm, whether knowledge restructuring will occur, either spontaneously or in response to experimental encouragement. As a first step toward understanding knowledge restructuring, and to seek a possible reconciliation between those opposing outcomes, we chose to study skill acquisition in a categorization learning paradigm using a variant of the fire-prediction task.

Categorization Learning

In categorization learning, people learn to classify stimuli into one of several categories on the basis of several diagnostic features or predictors. For example, people may learn to classify cartoon faces into the MacDonalds and the Campbells categories on the basis of features such as nose length, curvature of mouth, and so on.

Categorization is one of the most studied of all cognitive skills and, for several reasons, is ideally suited for an examination of knowledge restructuring. First, categorization learning is closely linked to the acquisition of expertise (e.g., Medin & Edelson, 1988). Many expert tasks, such as dermatological diagnosis (Brooks, Norman, & Allen, 1991), can be thought of as variants of categorization. Second, categorization tasks can be designed so that they are solvable by several alternative strategies, which in turn permits the use of experimental cues to encourage knowledge restructuring without overt alteration of the stimuli. Third, and most important, it is possible empirically to identify the strategies that people use to categorize a given set of stimuli.

Rules and Associative Learning

Although there are many competing theories of categorization, including some that view performance as relying entirely on representations of previously encountered instances (e.g., Kruschke, 1992; Nosofsky, 1986), there is widespread agreement that performance is at least sometimes driven by rules (Ashby & Gott, 1988; Erickson & Kruschke, 1998; Nosofsky, Clark, & Shin, 1989; Nosofsky, Palmeri, & McKinley, 1994). Nosofsky et al. (1989) showed that rules are particularly likely to be used when people are made aware of their existence before training, whereas performance is more likely to rely on instances if the same stimuli are presented without mention of categorization rules.

Particularly relevant here is that rule-based categorization has been implemented within an associative learning approach in the ATRIUM (attention to rules in a unified model) theory of Erickson and Kruschke (e.g., 1998, 1999). On this view, rules are not considered as static and rigidly declarative entities, but as decision boundaries that are continually updated and refined. The emphasis on continued learning is obviously essential in the context of knowledge restructuring. Moreover, associative learning principles are plausible candidates for an account of resistance to knowledge

restructuring, as they are also known to explain the absence of learning in certain situations (e.g., blocking; Kamin, 1969; Shanks, 1991).

Rule-based approaches to categorization typically assume that rule use is accompanied by memorization of occasional exceptions to the rule during training. Thus, a full characterization of a learner's strategy involves memory for instances that violate an applied otherwise valid rule. This rule-plus-exception approach is supported by much empirical evidence (e.g., Erickson & Kruschke, 1998; Nosofsky et al., 1994; Palmeri & Nosofsky, 1995). Here, we focused primarily on the empirical identification of rules, which has been shown to have a stronger influence on people's generalizations than memorized exceptions do (Erickson & Kruschke, 2000).

Rule Identification and Knowledge Restructuring in Categorization

For most categorization tasks, there are numerous strategies that can provide partial solutions, and different people may adopt different strategies (e.g., Nosofsky et al., 1994). Critically, on the rule-based view of categorization adopted here, those strategies are empirically identifiable, which renders the presence of knowledge restructuring apparent when a learner switches from one strategy to another.

In a study by Medin, Altom, Edelson, and Freko (1982), participants were able to classify stimuli on the basis of any one of two features with 75% accuracy, whereas the correlation between two other binary-valued features could support perfect performance. Specifically, a stimulus belonged to one category when two of its features had the same value and to another category when the values on those features were different from each other (i.e., an "XOR" relationship). Medin et al. found that after completion of training, people were sensitive to the XOR structure in the stimuli when classifying novel items. In a replication of that study, McKinley and Nosofsky (1993, reported in Nosofsky et al., 1994) examined the emergence of knowledge of the XOR relationship during training. At various points in training, participants completed a transfer test that included novel critical items that formed exceptions to the single dimensional rule but conformed to the XOR rule. Early in training, people were found to apply the expedient single dimensional rule, as revealed by low performance on the critical exception items. It was only on completion of training that people were found to have induced the more powerful XOR rule, as revealed by accurate categorization of the exceptions to the expedient rule. The results of McKinley and Nosofsky suggest that people can spontaneously acquire a more complex rule, after initially using a simpler one, if given enough learning opportunity.

Fire Prediction and Categorization

Fire prediction can be reexpressed as a categorization task in which two predictors (wind and slope) and the problem context (firefighting or back burning) permit categorization (uphill vs. downhill fire spread). Akin to the stimuli of Medin et al. (1982) and McKinley and Nosofsky (1993, reported in Nosofsky et al., 1994), the predictors wind and slope interact—albeit in a fairly complex manner—to determine fire direction. Again like in Medin

et al. (1982), a single feature, in this case context, provides an expedient but imperfect predictor.

The availability of competing expedient and complex rules renders the firefighting task suitable to an investigation of knowledge restructuring. To experimentally induce knowledge restructuring, the present experiments relied on a schematic diagram.

Schematic Diagrams and the Acquisition of Complex Knowledge

Diagrams have long been known to play a significant role in the acquisition of complex knowledge (e.g., Larkin & Simon, 1987; Mayer & Gallini, 1990). For example, Mayer and Gallini (1990) showed that problem solving involving mechanical devices (e.g., car brakes) was facilitated when a previously shown explanation of those devices was accompanied by a diagram of the component parts and their functions. Likewise, Biederman and Shiffrar (1987) showed that novices, after minimal instruction with a few simple diagrams, could match the performance of highly skilled experts on a perceptual learning task. The task involved the sexing of day-old chicks, a commercially relevant skill that had previously defied pictorial representation or training and that previously could only be learned through extensive practice (up to 1,000,000 trials). Here, we relied on a schematic diagram to encourage participants to take a different approach to a problem. The diagram was chosen atheoretically as a tool whose utility had been established previously, and without any commitment to specific perceptual or cognitive processes.

In the fire-prediction task, Lewandowsky et al. (1997) showed that a schematic diagram can improve subsequent training performance and can prevent reliance on an expedient imperfect rule. The diagram, shown in Figure 1, used the angle between the flames and the ground as the common referent for both wind- and slope-driven fires. According to the diagram, fire spreads whichever way the flames are pointing, whether blown by the wind (top panel of Figure 1), met by rising ground (bottom panel), or both.¹ The diagram altered performance by enhancing sensitivity to slope, thus leading to more accurate predictions overall.

Summary

There is strong evidence that skill acquisition entails the repeated restructuring of knowledge. However, there is also considerable evidence for resistance to knowledge restructuring during training and among fully trained experts. Nothing is known about the processes underlying knowledge restructuring, and it is currently impossible to predict whether people will restructure their knowledge in a given situation.

The present experiments use a categorization paradigm that is known to permit empirical identification of rules or strategies (e.g., Nosofsky et al., 1994). It is also known that knowledge restructuring can occur in categorization (McKinley & Nosofsky, 1993, reported in Nosofsky et al., 1994; Medin et al., 1982). The particular task chosen for the present studies is a categorization analog of a task on which experts arguably resist knowledge restructuring and continue to use an expedient strategy (Lewandowsky & Kirsner, 2000). Novices, by default, also adopt an expedient strategy, unless they are instructed with the aid of a diagram that




Figure 1. Schematic diagram of the effect of wind and slope on the angle between the flames and the ground. In the top panel, wind tilts the flames, thus forming an acute angle with the ground, whereas in the bottom panel, the flames remain vertical but nonetheless form an acute angle with the ground, owing to the rising slope.

visualizes a complex alternative strategy (Lewandowsky et al., 1997).

What remains to be resolved by the present experiments is whether people can be experimentally encouraged to restructure their knowledge by abandoning an expedient rule after initial learning in favor of a complex alternative. To help identify the processes underlying restructuring, the experiments also sought to identify the exact nature of the various strategies that people acquire in this task.

Experiment 1

The first experiment extended the results of Lewandowsky et al. (1997) to a categorization paradigm, using a diagram to encourage participants to adopt a complex strategy. The experiment also used the diagram to test whether knowledge restructuring could be induced after extended uninformed learning.

The experiment included three conditions: In the *control condition*, the diagram was never shown; in the *diagram throughout condition*, participants had access to the diagram throughout; and in the *reveal condition*, the diagram was presented after an initial phase of uninformed training. Based on previous results (Lewandowsky et al., 1997), the control and diagram throughout groups were expected to differ after the first phase of training, with the control condition relying on the expedient context rule and the diagram throughout condition adopting a more complex strategy.

¹ The true physical role of the angle between flames and ground is unknown and irrelevant here. What is important is that the angle provides a plausible integrative visualization of the two predictors.

Of greatest interest in this experiment was the performance of the reveal condition during the second phase of training: If performance resembled that in the diagram throughout condition, this would be indicative of successful knowledge restructuring. Conversely, if performance in the reveal condition were to resemble the control condition, this would be indicative of resistance to knowledge restructuring.

Method

Participants and apparatus. The participants were 60 undergraduate volunteers from the University of Western Australia who received partial course credit. The experiment was controlled by a computer that presented all stimuli and collected and scored all responses.

Design. An equal number of participants were randomly assigned to each of the three conditions (control, diagram throughout, and reveal). Regardless of condition, participants completed two phases, each involving 152 categorization learning trials followed by a transfer test. Phases were separated by a short break.

Stimuli. Stimuli presented for categorization consisted of written information about the three predictors wind speed, slope gradient, and context (back-burning vs. firefighting). Wind direction was always down slope, thus creating the conflict between potentially wind-driven and slope-driven fires. To visualize the opposition between wind (W) and slope (S), each stimulus was accompanied by the same schematic representation of a fire on a slope and an arrow that designated (downhill) wind direction. Participants had to classify each fire as slope-driven or wind-driven.

In each phase, 16 stimuli were used during training, and an additional eight were used during the transfer test. Predictor values for all stimuli are shown in Table 1. Half of the stimuli consisted of slope-driven fires and the

other half of wind-driven fires. All items obeyed an arbitrary, but physically plausible, nonlinear two-dimensional category boundary that related wind speed and slope gradient to fire direction. The best two-dimensional linear rule, by contrast, would incorrectly classify two of the training stimuli. All analyses used the true nonlinear boundary to determine when a stimulus was correctly classified.

Additionally, an imperfect one-dimensional rule was present, such that most back burns were slope-driven, whereas most to-be-controlled fires were wind-driven. Exceptions to this imperfect context rule are identified by asterisks in Table 1. The table also includes a column (β), relevant in Experiment 2, that represents the relation between wind and slope reexpressed as the angle (in degrees) between the flames and the ground. A representation of the stimuli and the category boundary in the space defined by wind speed and slope gradient appears in Figure 2.

Critically, the context rule was the only one-dimensional rule by which more than 50% of the stimuli could be correctly categorized, whereas there were several alternative strategies involving more than one dimension. For example, in addition to the correct nonlinear boundary, a complex rule such as "apply the context rule unless the slope is 13, in which case the fire is slope-driven, or unless the slope is 6, in which case the fire is wind-driven" could correctly classify all training items. For the purposes of this experiment, we contrasted the context rule with all possible multidimensional alternatives without seeking to identify which of those complex alternatives participants adapted when not relying on context. Experiment 3 takes up the issue of alternative complex rules in more detail.

Procedure. During the first training phase, participants learned to categorize stimuli through 10 blocks of 12 or 16 trials each. Each training stimulus occurred once per block, with a different random order of presentation in each block. During the first two blocks, the four exceptions (S7–S8, W7–W8) were withheld to encourage participants in the control condition to adopt the context rule (cf. Palmeri & Nosofsky, 1995).

On each trial, the stimulus remained visible until participants pressed the appropriate arrow key to indicate the predicted direction of the fire. Corrective feedback was provided for 2 s after a response, followed 1 s later by the next stimulus.

After completion of each training phase, a transfer test was administered. The transfer test included eight novel stimuli (S9–S12, W9–W12) in addition to the 16 training stimuli. The transfer test also included eight single-predictor stimuli, in which only context, slope gradient, or wind speed was provided. On each of the 32 randomly ordered test trials, participants again used the arrow keys to predict the direction of fire spread, but, in contrast to the training phases, no corrective feedback was given.

Participants in the diagram throughout condition were familiarized with the underlying relation between slope and wind at the outset. They were instructed that wind and slope can act in opposition to drive a fire, with the stronger of the two forces winning out. A diagram similar to Figure 1 remained visible throughout the experiment. In the reveal condition, these instructions and the diagram were administered during the break between training phases. In the control condition, instructions did not touch on the physics of fire, and the diagram was never presented.

General instructions about the task were given at the outset of the experiment and were identical between conditions. Participants were instructed to judge whether the fire represented on each trial would spread with the slope or with the wind, based on the values of the three predictors.

Results and Discussion

Analyses focused on the relative strength of the expedient context rule across conditions and phases of training. The primary index of the strength of this rule was the difference in performance between rule-compliant stimuli (S1–S6, W1–W6) and items that were exceptions to the rule (S7–S8, W7–W8). We first consider performance during training before reporting the transfer data. We

Table 1
Training Stimuli Used in Experiments 1 and 2

Category	Wind (km/h)	Slope (degree)	β (degree)	Item
Slope-driven				
Training/back-burning	5	8	88	S1
	10	8	89	S2
	5	10	86	S3
	10	10	87	S4
	10	13	84	S5
	15	13	87	S6
Training/firefighting	5	13	83	*S7
	20	13	89	*S8
Transfer/back-burning	7	9	89	S9
	12	10	88	S10
	14	15	84	S11
Transfer/firefighting	18	12	89	*S12
Wind-driven				
Training/back-burning	5	6	91	*W7
	20	6	97	*W8
Training/firefighting	10	6	92	W1
	15	6	95	W2
	15	8	93	W3
	20	8	95	W4
	15	10	91	W5
	20	10	93	W6
Transfer/back-burning	7	5	92	*W9
Transfer/firefighting	14	5	95	W10
	19	7	96	W11
	17	9	93	W12

Note. Exceptions to the simple Context rule are identified by asterisks. S = slope; W = wind.




Figure 2. Category structure used in Experiments 1 and 2. The curved solid line represents the true category boundary separating slope-driven from wind-driven fires. Training items are represented by circles, and new transfer test items by squares. Filled symbols represent to-be-controlled fires, whereas unfilled symbols represent back burns; this context rule allows 75% correct categorization.

then introduce a categorical analysis to further clarify which strategies people applied to the task.

Training. The categorization accuracy across blocks of training in both phases is shown in Figure 3 averaged across the two categories (wind-driven and slope-driven), with the context rule-consistent and exception items presented separately.

The proportion correct categorizations for rule-consistent items and exceptions were entered into a 3 (condition) \times 2 (phase) \times 2 (stimulus type: rule-consistent or exception) between-within analysis of variance (ANOVA), conducted on the last eight blocks of each phase. The main effect of type, $F(1, 57) = 189.16$, $MSE = 0.023$, $p < .01$, was significant, reflecting the fact that rule-consistent items were classified far more accurately overall than exceptions. The significant Type \times Phase interaction, $F(1, 57) = 43.53$, $MSE = 0.009$, $p < .01$, shows that this difference was greater in Phase 1 than in Phase 2. Practice also increased the level of performance overall, as shown by a significant effect for phase, $F(1, 57) = 142.33$, $MSE = 0.010$, $p < .01$. The remaining two-way Condition \times Phase interaction was not significant, $F(2, 57) = 1.13$, and neither was the main effect of condition, $F(2, 57) = 2.27$, $MSE = 0.050$, $p = .11$.

Of greatest interest is the significant three-way interaction, $F(2, 57) = 4.16$, $MSE = 0.009$, $p < .05$, which, together with the significant Condition \times Type interaction, $F(1, 57) = 6.57$, $MSE = 0.023$, $p < .01$, confirmed that performance on rule-consistent items and exceptions differed between conditions and across phases. In particular, planned comparisons revealed that

exceptions were classified more accurately in the diagram throughout condition than the other two conditions combined during Phase 1, $F(1, 57) = 15.36$, $MSE = 0.030$, $p < .01$, whereas accuracy for the rule-consistent items did not differ between conditions, $F(1, 57) = 1.36$. The selective benefit for exceptions to the context rule suggests that people in the diagram throughout condition were better able to understand the relation between slope and wind than were people who did not commence training with the diagram. This confirms the utility of the diagram and extends the finding by Lewandowsky et al. (1997) to a categorization learning paradigm.

In Phase 2, all three conditions were similar to each other in terms of both rule-consistent items and exceptions, with the largest difference (between the diagram throughout and control conditions on exceptions) falling far short of significance, $F(1, 57) = 1.47$. This confirms the pattern, apparent in Figure 3, that participants in all groups had mastered the training stimuli by the end of the experiment.

Transfer test. The training data alone cannot fully reveal what people have learned because categorization can be based on memory for the training instances. Thus, the seemingly identical training performance between conditions at the end of Phase 2 can obscure important differences in rule usage. To discover what rules people use, the new transfer items were analyzed by forming the difference between new rule-consistent items and new exceptions. A large difference would be indicative of the context rule being used, whereas a small difference would reveal the use of other knowledge, such as the correct physical relation between slope and wind.

The mean differences between new-rule and new-exception performance are shown in Table 2. The corresponding 3 (condition) \times 2 (phase) between-within ANOVA revealed a marginal Phase \times Condition interaction, $F(2, 57) = 2.96$, $MSE = .209$, $p = .06$, along with a marginal main effect of condition, $F(2, 57) = 2.91$, $MSE = .091$, $p = .06$, but no significant effect of phase, $F(1, 57) < 1$.²

The overarching interaction between condition and phase arose from several effects. First, performance in the diagram throughout condition varied with phase. In Phase 1, participants in that condition exhibited a much smaller difference between rule-consistent items and exceptions than did participants in the other two conditions, as confirmed by a planned comparison, $F(1, 57) = 18.57$, $MSE = .209$, $p < .01$. This paralleled the pattern in the training data and further suggested that people in the diagram throughout condition ignored the context rule in favor of a more complex strategy.

Second, the diagram only benefited performance when it was presented at the outset of training. Thus, the reveal condition in Phase 2 showed significantly greater difference scores—and hence continued reliance on the context rule—than did the diagram throughout condition in Phase 1, $F(1, 38) = 5.24$, $MSE = .111$, $p < .05$. Correspondingly, the reveal and control conditions did not differ significantly in Phase 2, $F(1, 38) < 1$. Taken together, this

² A parallel analysis used difference scores divided by the performance on the new exceptions to adjust for absolute level of accuracy. Because this measure yielded qualitatively similar results, only the simpler difference score analysis is reported here.

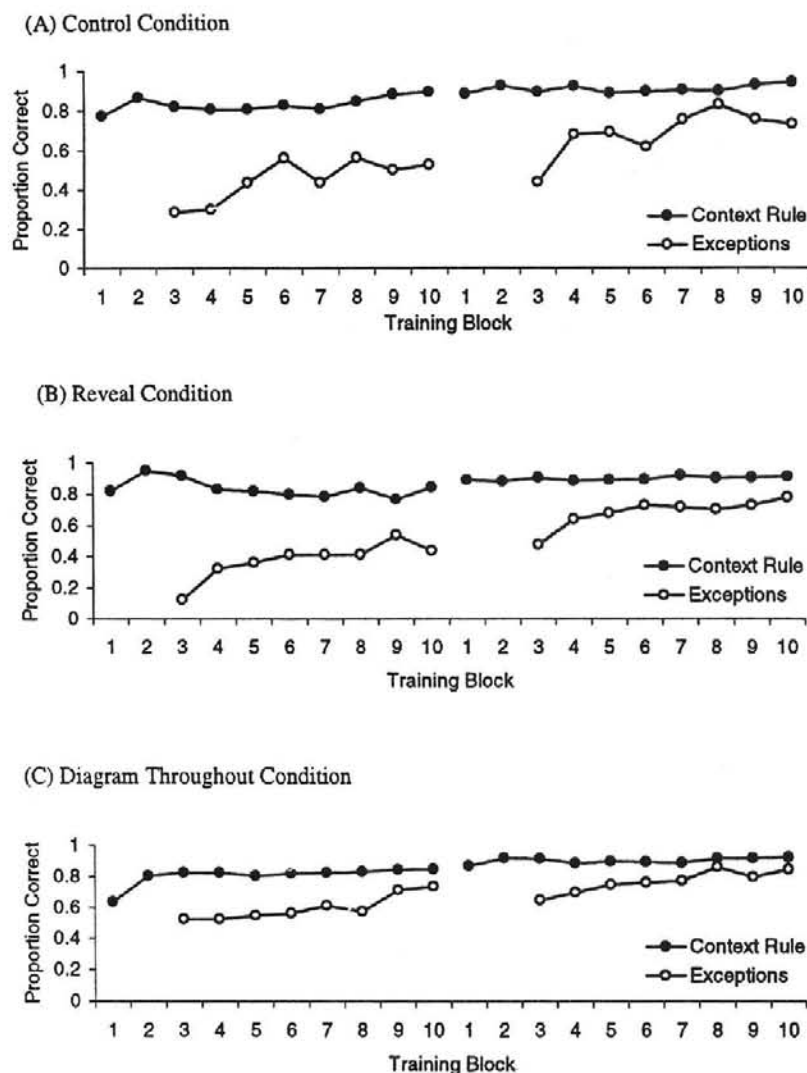


Figure 3. Categorization performance during training in Experiment 1.

indicates that once an expedient rule has been learned it is not readily discarded even when additional critical information becomes available.

Finally, the three conditions became more similar toward the end of training. The rule-minus-exception measure increased across phases for the diagram throughout condition and slightly decreased for the other two conditions. After Phase 2, the diagram throughout condition no longer differed significantly from the average of the other two, $F(1, 57) < 1$, suggesting that participants may have reverted to the expedient context rule even though they initially used a more complex strategy. Memorization of exceptions during Phase 2 would have enabled people to do so without loss of accuracy during training.

Categorical analysis. One difficulty with the rule-minus-exception measure is that, when aggregating across participants, only extreme values are diagnostic of the rules people are using. Specifically, near-zero performance on exceptions clearly indicates use of the context rule, and near-perfect performance is indicative of application of a more complex two-dimensional rule.

However, performance on exceptions could be near 0.5 either because all participants were responding at chance for these items or because a condition included individuals with different knowledge (e.g., some using the context rule, others using a more complex alternative) who showed complementary patterns of responding.

These problems were resolved by an additional categorical analysis that focused on the overall pattern of responses by each participant across all conditions. The eight new transfer stimuli were separated into the six items consistent with the context rule and the two exceptions. For each phase, the pattern of responding by each participant was then summarized by two scores, namely the number of correct responses to rule-consistent items (i.e., range 0–6) and the number of correctly classified exceptions (range 0–2). The 120 composite scores thus derived were entered into a k -means cluster analysis. This analysis classifies scores into k clusters, iteratively setting the cluster centroids so as to minimize the distance between scores within each cluster, while maximizing the distance between scores in different clusters. Once clusters are

Table 2
Mean Proportions Correct (and Standard Errors) During
Transfer Tests in Experiment 1

Condition and stimulus type	Old-new	Phase	
		1	2
Reveal			
Difference	New	.76 (.08)	.63 (.08)
Rule	New	.91 (.03)	.91 (.03)
Exception	New	.15 (.07)	.28 (.07)
Rule	Old	.87 (.03)	.94 (.03)
Exception	Old	.55 (.06)	.80 (.05)
Diagram throughout			
Difference	New	.39 (.10)	.56 (.08)
Rule	New	.84 (.04)	.93 (.03)
Exception	New	.45 (.09)	.38 (.09)
Rule	Old	.88 (.03)	.98 (.01)
Exception	Old	.71 (.06)	.94 (.04)
Control			
Difference	New	.73 (.07)	.63 (.10)
Rule	New	.91 (.03)	.91 (.03)
Exception	New	.18 (.05)	.28 (.08)
Rule	Old	.93 (.02)	.96 (.02)
Exception	Old	.74 (.06)	.84 (.04)

Note. Values marked in boldface were used in the analysis.

obtained, the number of participants in each cluster can be compared between conditions and phases.

On the basis of the outcome of the earlier ANOVA, we expected that response patterns could be classified into three clusters, corresponding to use of the context rule, use of the correct two-dimensional rule, and intermediate responses not consistent with either of these rules. The cluster analysis was thus performed with k set to 3, relying on the algorithm to determine optimal cluster centroids. The obtained cluster centroids based on all 120 patterns (i.e., without differentiating between conditions and phases) are given in Table 3, along with the labels chosen as best descriptors of each cluster. The *Correct* cluster contained patterns in which a large number of correct responses were given for both the rule-consistent items and the exceptions. The *Rule* cluster featured high performance on rule-consistent items but not on exceptions. The *Partial* cluster included patterns of responding that suggested partial acquisition of the task or, possibly, chance performance on some items.

Because cluster labels are necessarily subjective, independent support of their descriptive accuracy is desirable. Accordingly, we examined responses to the single-predictor transfer items for participants within each cluster separately, with the mean proportion of "wind-driven" responses for each of the predictors shown in Figure 4.

Table 3
Results of k -Means Analysis on Composite Scores
in Experiment 1

Cluster name	Position of cluster centroid		No. of scores
	Correct exceptions	Correct rule items	
Correct	1.54	5.25	28
Rule	0.23	5.76	80
Partial	0.58	3.42	12




Figure 4. Sensitivity to single-dimensional stimuli for participants in different clusters in Experiment 1. The figure shows the proportion of responses on which a single-dimensional stimulus was expected to go with the wind.

Sensitivity to context, wind, or slope is apparent if the proportion of wind-driven responses differs across levels of each predictor. Consistent with the labeling of clusters, it is evident that participants in the *Rule* cluster displayed the greatest sensitivity to context, whereas those in the *Correct* cluster were most sensitive to the effects of wind speed and slope gradient. The individuals in the *Partial* cluster showed some awareness of the effects of wind and slope, but showed less sensitivity overall than people in the other clusters.

Categorical tests of transfer results. The number of participants in each cluster was then determined for each condition and phase. These values are shown in Figure 5, expressed as proportions of the total number of participants in each condition. To examine whether these proportions differed with condition, chi-square tests were carried out with $N = 60$ for each phase separately, yielding a significant result in Phase 1, $\chi^2(4, N = 60) = 11.97, p < .05$, but not in Phase 2, $\chi^2(4, N = 60) = 1.64$. Examination of Figure 5 shows that the significant effect in Phase 1 resulted from the fact that in the diagram throughout condition, more participants were assigned to the *Correct* cluster

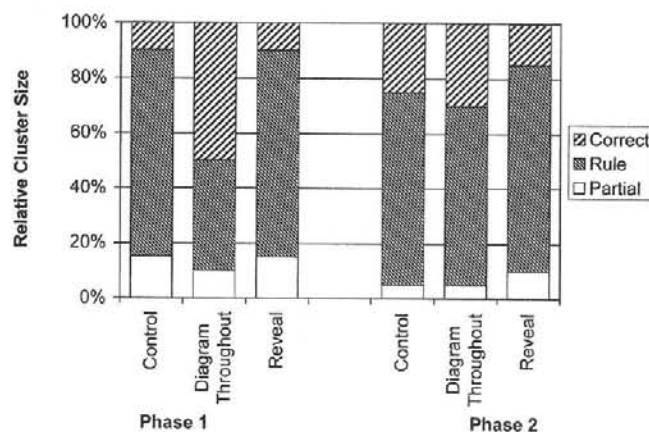


Figure 5. Relative proportion of participants assigned to clusters as a function of condition and phase for Experiment 1.

and fewer to the Rule cluster than in the other two conditions. This parallels the earlier analysis of difference scores and further supports the notion that the diagram, when presented at the outset of training, encourages participants to forego the expedient context rule. Unlike the difference scores, the cluster analysis identified participants (in the Partial cluster) who may have responded at chance for some items.

Turning to Phase 2, the absence of any differences between conditions mirrors the earlier ANOVA in two ways. First, the similarity of the reveal and control conditions adds to the evidence that the diagram could not induce a more complex strategy once an expedient rule had been acquired. Second, the similarity of the control and diagram throughout conditions suggests that people in the latter condition, who initially used a complex categorization strategy, eventually reverted to the expedient context rule. The details of our preferred explanation for this finding are provided in the General Discussion section. Briefly, the explanation assumes that participants remained in a state of learning regardless of which rule they applied to the task. Because accuracy in Phase 1 never reached 100%, the perfect performance of the context rule during the initial exception-free blocks in Phase 2 likely convinced participants to revisit this expedient alternative to a partially learned complex strategy. Additionally, a strategy that combined the expedient rule with memorization of exceptions would have permitted perfect performance on the training items, thus further reducing the attractiveness of a partially learned complex alternative.

The tendency of our participants to adopt the context rule despite complex initial knowledge may parallel the real-world behavior of expert firefighters observed by Lewandowsky and Kirsner (2000). Firefighters are initially trained with a complex physical model of fire spread and may use that information to guide their decision making. Years of experience in the field, however, may offer an expedient alternative: The statistical coincidence of fire contexts with direction, combined with a likely degree of error when experts apply the correct physical model, may cause the emergence of a context-based prediction of fire spread.

Summary. The results of Experiment 1 are readily summarized: (a) When several rules compete during categorization, people by default select an expedient rule augmented by memorization of exceptions. (b) People can be encouraged to use a more complex alternative when minimal instruction and a visualization of the alternative is provided at the outset. (c) Task performance differs qualitatively between those two conditions, with great differences in sensitivity to individual predictors and differences in overall response patterns. (d) Once an expedient rule has been acquired, visualization and instruction are insufficient to induce more complex processing. (e) Conversely, even if people commence the task with a complex strategy, after prolonged training they tend to adapt the expedient context rule.

The asymmetry of knowledge restructuring that is implied by the latter two results forms the core contribution of Experiment 1. We discuss possible explanations for the asymmetry after presenting two additional studies that explored a possible enhancement to the diagram's utility for knowledge restructuring (Experiment 2) and that identified the exact nature of the knowledge that is learned at the outset with and without the diagram (Experiment 3).

Experiment 2

The purpose of the second study was to test a more explicit and continuously available technique to encourage expediently trained participants to shift to a complex strategy. To this end, Experiment 2 included an adaptive representation of the diagram that reexpressed the particular wind-slope combination on each trial as an exact angle between the flames and the ground. The adaptive display, when shown during a training trial, reliably and unambiguously predicted fire spread as a function of displayed flame-ground angle.

The adaptive display was available at varying points during training: In the adaptive removal condition, participants commenced training with the adaptive display before it was removed in the second phase. Conversely, in the reveal adaptive condition, the adaptive display was introduced for the second training phase. Finally, the control condition was identical to that of Experiment 1 and received neither the adaptive display nor instructions relating to flame-ground angle.

As in Experiment 1, we focused on performance of the reveal adaptive condition during Phase 2. If performance resembled that in the control condition in Phase 2, people would again have been shown to resist knowledge restructuring. If performance instead resembled that in Phase 1 of the adaptive removal condition, knowledge restructuring would have occurred. The latter finding would be expected if the adaptive display was more powerful than the static diagram of Experiment 1.

Method

Participants, apparatus, and stimuli. A different 60 participants were recruited from the University of Western Australia community to participate in Experiment 2. An equal number of participants were randomly assigned to each of the three conditions. The apparatus and category structure were identical to those used in the first study.

Procedure. The procedure was identical to that of Experiment 1, with the exception of the presence of the adaptive display on some training trials. In the adaptive removal condition, participants were familiarized with the underlying relationship between slope and wind at the outset, and the first phase of training was accompanied by the adaptive display. In the reveal adaptive condition, instructions about flame-ground angle were given during the break between phases, and the second phase of training was accompanied by the adaptive display.

When the adaptive display was present, the stimulus was accompanied by an abstract visual rendition of the angle between the flames and the ground (β), consisting of a horizontal line with an oriented line segment joined to its center point. On each trial, the line segment was oriented with angle β and was presented below the written information about predictor values in a frame of about 5×5 cm. The value of β for each stimulus was given earlier in Table 1 and was calculated using an arbitrarily designed equation that described a feasible nonlinear category boundary: $\beta = 94.75 - (\text{slope}) + .21 (\text{wind}) + .01 (\text{wind})^2$. Whenever β was less than 90° , the fire was slope-driven, and when β exceeded 90° , it was wind-driven. The value of β was printed within the frame next to the visual rendition of the angle; see Figure 6 for a sample adaptive display.

Regardless of condition or phase, transfer trials did not include the adaptive display.

Results and Discussion

As in Experiment 1, we analyzed the training and transfer data separately and with a focus on use of the context rule. We again used a categorical analysis to further identify rule use. To foreshadow the




Figure 6. Example of the adaptive display used in Experiment 2. When present, this display was shown below the textual information for the stimulus.

main result, the adaptive β display, despite drastically altering training performance, failed to cause knowledge restructuring.

Training. Categorization accuracy during training is shown in Figure 7. Not surprisingly, categorization was virtually perfect when the adaptive display was present, with near 95% correct

performance during Phase 1 for the adaptive removal condition and during Phase 2 for the reveal adaptive condition. The training data were examined using a 3 (condition) \times 2 (phase) \times 2 (stimulus type) between-within ANOVA, excluding the first two blocks of each phase. All interactions and main effects were significant: condition, $F(2, 57) = 3.52$, $MSE = 0.02$, $p < .05$; phase, $F(1, 57) = 61.24$, $MSE = 0.01$, $p < .001$; type, $F(1, 57) = 197.64$, $MSE = 0.01$, $p < .001$; Condition \times Phase, $F(2, 57) = 196.54$, $MSE = 0.01$, $p < .001$; Condition \times Type, $F(2, 57) = 6.30$, $MSE = 0.14$, $p < .01$; Phase \times Type, $F(1, 57) = 25.92$, $MSE = 0.01$, $p < .001$; and the three-way interaction, $F(2, 57) = 72.56$, $MSE = 0.01$, $p < .001$.

The exceptional performance with the adaptive display was responsible for most of the observed effects. The adaptive removal condition did better than the other two conditions in Phase 1, and the reveal adaptive condition did better than the others in Phase 2. The three-way interaction is also indicative of the clear superiority of learning with the adaptive display in terms of accuracy on

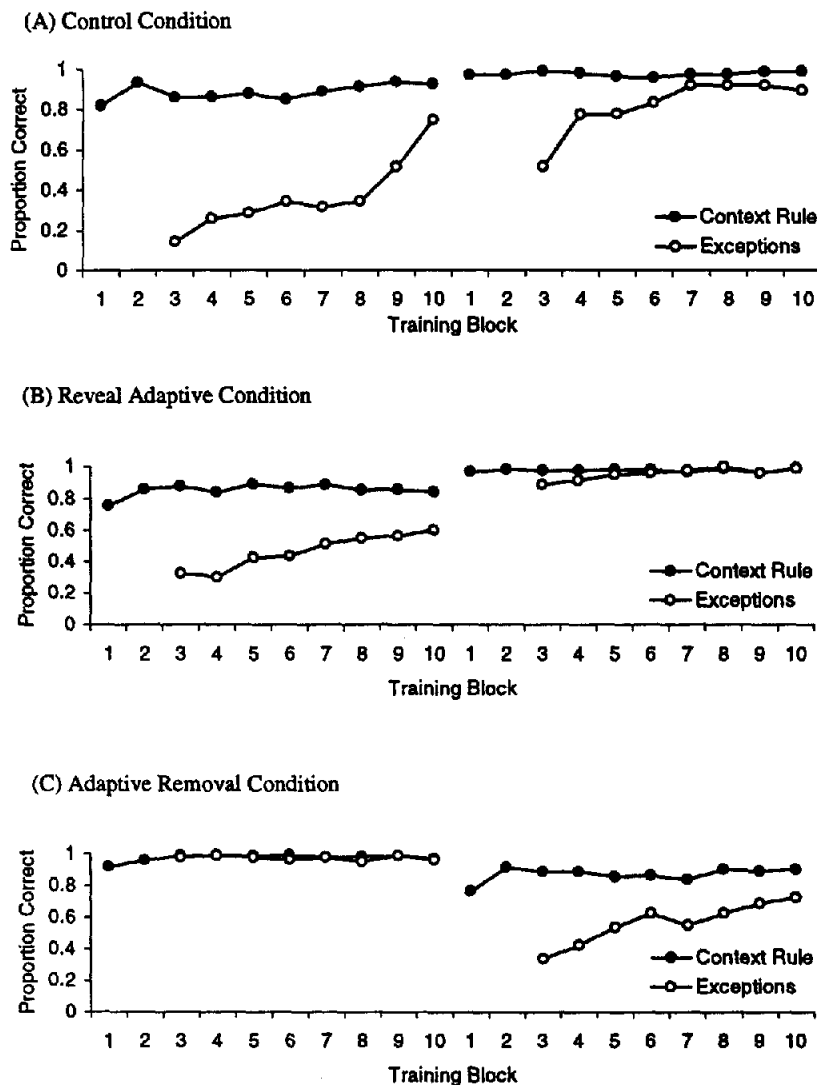


Figure 7. Categorization performance during training in Experiment 2.

exceptions relative to the control condition. Whenever the adaptive display was present, there was no reliable difference between rule and exception stimuli. Finally, the main effect of phase represented improvements due to practice.

Transfer test. The transfer data were analyzed using the same measure of expedient rule use (new rule-consistent items correct – new exceptions correct) as in Experiment 1. The results are shown in Table 4. The corresponding ANOVA revealed a Condition \times Phase interaction, $F(2, 57) = 7.05$, $MSE = 0.12$, $p < .01$, which primarily arose from the extremely low difference score for the adaptive removal condition in Phase 1. The additional main effects of condition, $F(2, 57) = 7.10$, $MSE = 0.180$, $p < .01$, and phase, $F(2, 57) = 3.22$, $MSE = 0.120$, $p < .05$, also reflected the impact of the adaptive display during Phase 1.

Two aspects of the transfer data merit discussion. First, the principal finding of Experiment 1 was replicated, because a single phase of training was again sufficient to block any reorganization of knowledge on presentation of the adaptive display. Specifically, the rule-exception difference for the reveal adaptive condition did not differ significantly between phases (.68 vs. .64).

Second, the adaptive display had a major impact on performance when presented at the outset, as shown by the strikingly small rule-exception difference in Phase 1 for the adaptive removal condition. This constitutes clear evidence that people did not rely on the expedient context rule under these circumstances. However, the small rule-exception difference was accompanied by absolute performance levels that were perilously close to chance even for rule-consistent items (.53 and .71 for old and new, respectively). It is therefore likely that the adaptive display prevented the learning of any display-independent rule, whether expedient or complex, rather than encouraging people to learn a more complex one. Instead, people seemingly learned only to judge whether β was less than or greater than 90. This was confirmed by the categorical analysis.

Table 4
Mean Proportions Correct (and Standard Errors) During Transfer Tests in Experiment 2

Condition and stimulus type	Old-new	Phase	
		1	2
Adaptive removal			
Difference	New	.16 (.10)	.61 (.09)
Rule	New	.71 (.05)	.88 (.03)
Exception	New	.55 (.09)	.28 (.08)
Rule	Old	.66 (.04)	.94 (.02)
Exception	Old	.53 (.04)	.81 (.05)
Reveal adaptive			
Difference	New	.68 (.08)	.64 (.09)
Rule	New	.91 (.02)	.94 (.02)
Exception	New	.23 (.08)	.30 (.08)
Rule	Old	.88 (.03)	.89 (.03)
Exception	Old	.71 (.06)	.69 (.07)
Control			
Difference	New	.75 (.06)	.68 (.09)
Rule	New	.98 (.02)	.96 (.02)
Exception	New	.23 (.06)	.28 (.08)
Rule	Old	.93 (.02)	.96 (.02)
Exception	Old	.74 (.06)	.84 (.04)

Note. Values marked in boldface were used in the analysis.

Categorical analysis. As in the previous study, a k -means cluster analysis was conducted that classified response patterns into three clusters. The three centroids are shown in Table 5. These clusters were consistent with those found in Experiment 1, and were thus given identical labels. To confirm that labels were appropriate, the responses to single-predictor stimuli were again examined separately for participants in each of the clusters. Figure 8 shows the mean proportion of “wind-driven” responses given to each of the single-predictor stimuli. As in Experiment 1, participants in the Rule cluster displayed the greatest sensitivity to context, whereas those in the Correct cluster were most sensitive to the effects of wind and slope. Participants in the Partial cluster showed some awareness of the effects of wind and slope, but less sensitivity than observed in the other clusters.

Separate chi-square tests were conducted for each phase to compare the relative proportions of participants assigned to each cluster across conditions. The data submitted to this analysis are shown in Figure 9. The figure shows that more participants were classified as Correct overall than in the previous experiment, which is the result of a small change in the location of cluster centroids. As in Experiment 1, the chi-square test yielded a significant result in Phase 1, $\chi^2(4, N = 60) = 22.49$, $p < .001$, but not in Phase 2, $\chi^2(4, N = 60) = 5.24$, $p > .25$. The significant effect for Phase 1 reflects the greater number of participants who were assigned to the Partial cluster, combined with a smaller number in the Rule cluster, in the adaptive removal condition than the other two conditions. The smaller number of participants in the Rule cluster confirms the conclusion of the earlier ANOVA, that the adaptive display decreased reliance on the context rule. However, the greater number of participants in the Partial cluster clarifies that this decreased reliance resulted from reduced learning overall, rather than from redirecting learning toward a more complex rule.

Turning to Phase 2, the number of Partial response profiles in the adaptive removal condition returned to a level consistent with that of the control condition, which in turn did not differ much from the reveal adaptive condition. The similar clustering observed in the control and reveal adaptive condition is consistent with the conclusion of the earlier ANOVA that the adaptive display had no effect on transfer performance once people had learned the context rule.

Summary. Experiment 2 yielded three principal findings: (a) In replication of the first study, presentation of an enhanced, adaptive version of the diagram also failed to induce knowledge restructuring after the expedient context rule had been acquired. (b) When presented at the outset, the adaptive display prevented significant learning of any rule, whether expedient or complex, presumably because it reduced the task to one of classifying a single number, β , as being above or below 90. (c) When the adaptive display was removed after it had prevented display-independent learning, people acquired the expedient context rule.

The experiment thus supported the attractiveness of the expedient context rule in two ways: Its use persisted even when the adaptive diagram was presented during Phase 2, and it was acquired after the adaptive display had obviated any display-independent learning during Phase 1.

Table 5
Results of *k*-Means Analysis on Composite Scores
in Experiment 2

Cluster name	Position of cluster center		Number of scores
	Correct exceptions	Correct rule items	
Correct	1.30	5.54	46
Rule	0.00	5.79	57
Partial	0.82	3.53	17

Experiment 3

The first two experiments revealed that people resisted knowledge restructuring after they had mastered the task in an expedient manner. In addition, Experiment 1 showed that people spontaneously abandoned a complex strategy in favor of an expedient one after extensive training. However, the experiments did not fully explore all possible strategies available to participants. In particular, it was assumed by exclusion that participants had learned the correct category boundary if they did not appear to use the context rule. The purpose of Experiment 3 was to further dissect the strategies used in the first two experiments, and to identify how presentation of the diagram at the outset in Experiment 1 affected people's strategies. Identification of those strategies is essential for the development of a theoretical account for the results of the first two experiments.

The categorical analysis of Experiments 1 and 2 focused on two principal strategies available to participants. The *Correct* strategy relied on the complex nonlinear rule involving wind and slope. The *Context* strategy predicted fire spread on the basis of whether it was a back burn or a to-be-controlled fire. Examination of the training stimuli used in those experiments (see Table 1) revealed two additional strategies that participants might potentially use: the *Greater-Than-Two* strategy and the *Thirteen-and-Six* rule.

The *Greater-Than-Two* strategy is based on a rule that can be stated as "if wind speed is larger than slope gradient by a number greater than two, then the fire is wind-driven," or, formally:




Figure 8. Sensitivity to single-dimensional stimuli for participants in different clusters in Experiment 2. The graph shows the proportion of responses on which a single-dimensional stimulus was expected to go with the wind.




Figure 9. Relative proportion of participants assigned to clusters as a function of condition and phase for Experiment 2.

$$\text{Response} = \begin{cases} \text{Wind if } W - S > 2 \\ \text{Slope otherwise,} \end{cases} \quad (1)$$

where W is the wind speed and S is the slope gradient. Although this strategy nominally involves two predictors, it is best considered a one-dimensional rule because the participant can combine the simultaneously visible numeric wind and slope values into a single difference score, which is then compared with a criterion value. As such, the strategy is best considered another variant of an expedient rule, similar to the context rule. It can be visualized as a linear boundary in the stimulus space in Figure 2. The rule works for all but two of the training stimuli, so the *Greater-Than-Two* strategy would require memorization of two instances during training.

The *Thirteen-and-Six* strategy can be stated as "if the slope is 13 or greater, then the fire is slope-driven, if the slope is six or less, then the fire is wind-driven, otherwise apply the context rule," or formally:

$$\text{Response} = \begin{cases} \text{Slope if } S + 6C \geq 13 \\ \text{Wind otherwise,} \end{cases} \quad (2)$$

where C is the context, with "0" representing a to-be-controlled fire and "1" a back burn. The rule used in the *Thirteen-and-Six* strategy is clearly more complicated than the *Greater-Than-Two* and the *Context* rules, but allows perfect categorization of the training stimuli.

The transfer stimuli used in Experiments 1 and 2 were not distributed widely enough to discriminate between these four possible strategies. The aim of Experiment 3 was to replicate the effect of the diagram in the first phase of Experiment 1 and to use a richer set of transfer stimuli to provide insight into what rules participants learned. For comparison, an uninformed control condition was included that also received the richer transfer set. Because this investigation focused on the exact differences between strategies with and without the diagram presented at the outset, only one phase of training was conducted.

Based on the results so far, strategies were expected to differ between conditions. Usage of the *Correct* rule was expected to be greater in the diagram condition than in the control condition, with the reverse being true for the *Context* rule. Concerning the remain-

ing strategies, the Greater-Than-Two rule was thought to be more likely to arise in the control condition, because it offered an expedient solution once slope was subtracted from wind, whereas the complex Thirteen-and-Six rule was expected to occur in the diagram condition, if at all.

Method

Participants and apparatus. The participants were 60 undergraduate volunteers from the University of Western Australia who received partial course credit. The experiment was controlled by a computer that presented all stimuli and collected and scored all responses.

Design. An equal number of participants was randomly assigned to one of two conditions, control and diagram. The control condition used the same training procedure as the first phase of the control condition in Experiments 1 and 2 but included a new set of transfer stimuli. Participants in this condition received only basic instructions about the task. Participants in the diagram condition were provided with the diagram and additional instructions at the outset, matching the first phase of the diagram throughout condition in Experiment 1. Other than that, the diagram condition was identical to the control condition.

All participants completed a single training phase involving 152 categorization learning trials (two blocks of 12 context rule-consistent items followed by eight blocks incorporating all 16 training items). The transfer test that followed consisted of 40 trials.

Stimuli. The training stimuli were the same as those used in Experiments 1 and 2. The transfer stimuli included the original eight, as well as 32 new stimuli designed to be diagnostic of strategy use. The new stimuli extended the transfer set along the true category boundary, where the four strategies showed the greatest variation in predictions. Altogether, the set of transfer stimuli occupied 20 locations in wind-slope space, and unlike the first two experiments, each location was presented once in each of the two contexts. The locations of the new transfer stimuli, together with lines indicating the true category boundary (the Correct rule) and the response boundary predicted by the Greater-Than-Two rule, appear in Figure 10. The figure also shows in which context training items were presented.

Procedure. With the exception of the enhanced set of transfer trials, the procedure was identical to the first phase of Experiment 1.





Figure 10. Category structure used in Experiment 3. Circles are training stimuli, identified by the context in which they were presented. Crosses are new transfer stimuli, occurring in both contexts. The heavy line shows the true category boundary; the broken line indicates the category boundary for the Greater-Than-Two strategy.

(A) Control condition



(B) Diagram Condition




Figure 11. Categorization performance during training in Experiment 3.

Results and Discussion

Focus rested on replication of the diagram manipulation from Experiment 1. For this study, the categorical analysis used the four hypothesized strategies rather than inductively constructed clusters. To foreshadow the results, in replication of Experiment 1, participants were less influenced by context when trained in the diagram condition. Extending the findings of the first study, the analyses additionally clarified that participants in the diagram condition were most likely to learn the Correct strategy as opposed to any of the other two-dimensional alternatives (Greater-Than-Two and Thirteen-and-Six).

Training. In order to examine correspondences between the present results and those of Experiment 1, stimuli were again divided into context rule-consistent items and exceptions to that rule. Categorization accuracy during training is shown in Figure 11.

A two-way between-within ANOVA, examining the effects of condition (control vs. diagram) and stimulus type (rule vs. exception) conducted on the last eight blocks of training showed significant main effects of condition, $F(1, 58) = 4.16$, $MSE = 0.02$, $p < .05$, and type, $F(1, 58) = 114.96$, $MSE = 0.02$, $p < .001$, but no interaction, $F(1, 58) = 0.22$, $MSE = 0.02$, $p > .60$. The main effect of condition reflects better overall performance in the diagram condition, and the effect of type shows the expected superiority of rule-consistent items over exceptions.

On the basis of Experiment 1, the interaction between condition and type should have been significant, because the diagram should selectively enhance performance on the exceptions. In order to investigate this apparent discrepancy between experiments, a three-way between-within ANOVA was conducted, adding block

(3 to 10) as a factor. For this analysis, in addition to the earlier effects, significant effects emerged for block, $F(7, 406) = 8.70$, $MSE = 0.02$, $p < .001$, Condition \times Block, $F(7, 406) = 2.52$, $MSE = 0.02$, $p < .05$, Type \times Block, $F(7, 406) = 9.97$, $MSE = 0.03$, $p < .001$, and Condition \times Type \times Block, $F(7, 406) = 3.55$, $MSE = 0.03$, $p < .001$. The main effect of block was clearly a consequence of practice, whereas the two- and three-way interactions with block arose because of the predicted changing patterns of responses during the course of training.

Examination of Figure 11 shows that performance on exceptions to the context rule was far higher in the diagram condition than the control condition during the first few blocks of training with exceptions. This difference raised performance for exceptions in the diagram condition sufficiently to produce the two-way interactions, and its additional dependence on the amount of training resulted in the Condition \times Type \times Block interaction. Thus, the diagram did have a selective effect on performance on exception items, paralleling the outcome of Experiment 1. This effect was reduced by further practice in the control condition, presumably because participants began to memorize the exceptional instances.

Transfer test. The analyses of Experiments 1 and 2 used a rule-minus-exception measure to assess the performance of participants with respect to the expedient context rule. This measure was used here as well, although it was computed over the much larger set of new transfer items. A comparison of the two conditions by an independent groups t test yielded $t(58) = 3.48$, $p < .001$, with participants in the control condition showing a greater difference in performance between rule and exception items. The corresponding means and standard errors are given in Table 6. Thus, Experiment 3 clearly replicated the first study using a more extensive set of transfer items.

In this experiment, each combination of wind and slope was presented in each of the two contexts at transfer. This permits enumeration of the number of items that are expected to be treated differently between contexts for each strategy. For the Context rule, all 20 items would be treated differently. For the Thirteen-and-Six rule, only four of these 20 pairs would be treated differently between contexts. Finally, the Greater-Than-Two or Correct strategies would treat all these pairs the same.

The number of pairs categorized differently was tallied for each participant. The mean number of differently categorized pairs

was 8.0 for the control condition, and 5.2 for the diagram condition. Because tallies were nearly uniformly distributed in both conditions, a distribution-insensitive Wilcoxon test on the ranks of the participants was conducted, which showed that the reliance on context differed between conditions, $\chi^2(1, N = 60) = 4.06$, $p < .05$.

Analysis of rules. We further examined strategy use by constructing a vector of predicted responses across all test items for each strategy (arbitrarily coding wind-driven responses as 0, slope-driven ones as 1) and computing how much participants in each condition deviated from these vectors. The responses of each participant were converted to a vector by the same process, and the Hamming distance of each participant's response vector to each strategy vector was calculated. Each participant was then classified as using one of the four candidate rules by choosing the strategy whose vector of predicted responses was closest to the participant's response vector. As shown in Table 7, the diagram condition included more participants using the Correct strategy, and fewer using the Greater-Than-Two or Context strategies, relative to the control condition, $\chi^2(3, N = 60) = 8.30$, $p < 0.05$. This confirmed that the diagram encourages people to adopt the correct two-dimensional rule, as opposed to some other heuristic involving multiple predictors.

As expected, few participants relied on the Correct strategy in the control condition. Instead, most people relied either on the pure Context rule or, with roughly equal likelihood, on the Greater-Than-Two strategy. The latter strategy is particularly interesting because although it involves wind and slope, it is quite similar to the Context rule because it reduces wind and slope information to a single number (the difference between the two values). Thus, even when people learn to consider wind and slope in the control condition, they do so in the most expedient manner possible.

These results provide useful constraints for the theoretical account of the first two experiments that we offer in the General Discussion. In particular, the account must recognize that the diagram does more than simply shift people's attention away from context: The fact that roughly half the participants in the control condition considered wind and slope but without recognizing their complex relationship indicates that attention to those two predictors by itself is insufficient to induce the correct strategy. The further fact that most people in the diagram condition used the correct strategy confirms that the diagram provided unique information about the underlying relationship between wind and slope.

Consistency of performance and individual variability. There was considerable variability in the extent to which participants' responses were described by the strategy to which they were assigned. To illustrate this variability, and to further explore the differences between conditions, accuracy of responses to items that all four strategies classify identically (the 12 "common" items) were used to predict the fit of the strategy to which a participant was assigned. Accuracy on the common items is an index of the extent to which an individual's performance is rule-bound, as opposed to random, regardless of which particular strategy is being followed.

This analysis revealed a relationship between accuracy on the common items and consistency of performance on the remaining 28 strategy-unique items ($r^2 = 0.26$), $t(58) = 4.5$, $p < .01$, such that participants who classified the common items with greater accuracy, also more consistently applied whichever of the four strategies best described their performance. Figure 12 depicts this relationship between the level of misprediction on common

Table 6
Mean Proportions Correct (and Standard Errors) During
Transfer Test in Experiment 3

Condition and stimulus type	Old-new	Proportion
Diagram		
Difference	New	-.05 (.07)
Rule	New	.64 (.06)
Exception	New	.67 (.04)
Rule	Old	.89 (.02)
Exception	Old	.80 (.02)
Control		
Difference	New	.25 (.06)
Rule	New	.79 (.02)
Exception	New	.52 (.05)
Rule	Old	.85 (.03)
Exception	Old	.55 (.06)

Note. Values marked in boldface were used in the analysis.

Table 7
Best Fitting Strategies by Conditions in Experiment 3

Condition	Best-fitting strategy			
	Greater-Than-Two	Correct	Context	Thirteen-and-Six
Control	12	7	10	1
Diagram	7	17	4	2

items and the fit of the best strategy on the unique items for both conditions separately.

The figure also shows the overall level of variability in fit of the strategies. It is apparent that in the diagram condition, participants tended to conform more closely to one of the four identified strategies, regardless of which strategy best described a person. This reduction in variability was accompanied by a shift away from Context and toward increased use of the Correct strategy.

Participants in the control condition, by contrast, showed less consistency overall and were less well described by any of the four strategies. Nonetheless, the two expedient approaches (Context and Greater-Than-Two) are again seen to dominate responding.

General Discussion

Summary of Results

Participants in these experiments mastered the categorization task in one of two qualitatively different ways. They either applied an expedient rule (accompanied by memorization of exceptions) or they relied on a complex strategy that involved two predictors. In the absence of any training aids, participants strongly preferred an expedient approach.

The first experiment showed that people can favor the complex strategy over the expedient rule if an underlying relationship between predictors is revealed at the outset. However, when the relationship was visualized after the expedient rule had been demonstrably induced, people did not abandon the expedient strategy. The first experiment also showed that people tended to migrate toward the expedient strategy after extensive training, even if they had initially used a complex approach.

The second study used an adaptive display that reliably recoded the actual category boundary involving slope and wind into a single dimension—the angle between the flames and the ground—that permitted perfect categorization. Participants clearly relied on this recoded dimension during training, with performance being near ceiling on all items, rule and exception alike. However, during the diagram-free transfer phases, performance depended greatly on whether training had commenced with the adaptive display. People who received the adaptive display at the outset did not learn much about any rule, whether expedient or complex, as shown by the poor overall level of transfer performance. When the adaptive display was introduced after unaided initial training, people again showed no evidence of departing from the already-learned expedient rule. Finally, participants who initially received the adaptive display but then reverted to unaided training in the second phase succeeded in inducing the expedient rule during the second phase.

The third experiment further dissected people's strategies by directly testing different variants of expedient and complex rules.

In confirmation of the first study, the majority of participants were found to rely on the correct rule involving wind and slope when the diagram was presented at the outset. In the absence of the diagram, fewer than 25% of participants acquired the correct rule. Instead, most people used a variant of an expedient strategy that relied either exclusively on context, or on the simple linear difference between wind speed and slope.

Overall, the most compelling result that requires theoretical consideration is the asymmetry of knowledge restructuring observed in the first two experiments: On the one hand, once an expedient rule had been acquired, participants resisted a shift to greater complexity, whether it was revealed by a static diagram or an adaptive display. On the other hand, Experiment 1 showed that if participants initially used a more complex strategy, they tended to migrate toward an expedient rule after extended training. Thus, people may abandon complex strategies in favor of simpler ones, even if that entails the risk of a performance loss, whereas people resist shifting from a simple to a complex strategy, even if they are explicitly encouraged to do so.

We begin our discussion by considering relevant existing results and explore explanations that have been put forward for individual, related findings. These explanations are then integrated into the associative learning framework offered at the outset.

Prior Knowledge and New Learning

There is, of course, an abundance of evidence that old knowledge persists and interferes with the acquisition of new knowledge when novel task requirements are incompatible with those originally learned. Thus, learning of new verbal associations is inhibited by previously learned lists (e.g., Postman, 1971), performance on a visual search task is impaired for prolonged periods if targets and distractors are reassigned (Shiffrin & Schneider, 1977, Experiment 1), and in categorization learning, persistence of old knowledge is observed when a different set of features becomes predictive after initial training (e.g., Macho, 1997).

However, all these cases differ from the present studies in at least one critical respect: Here, the alternative ways of performing the task were not incompatible and coexisted throughout. Accordingly, aside from the presentation of the static diagram or adaptive display, overt task requirements were identical in all conditions and phases. Hence, any theoretical account of the observed asymmetry of knowledge restructuring must ultimately focus on people's cognitive representations of the task. Because we know of no precedents for a completely cognitive manipulation of a shift in strategies, let alone explanations, we begin by discussing four key aspects of our results in the context of procedures in which overt task requirements were changed. As noted at the outset, we framed our discussion within an associative learning approach.

Resistance to knowledge restructuring. Consider first the observed reluctance to shift to an asymptotically better, albeit more complex, strategy once an expedient solution had been learned. The first phase of training offered people a choice between learning of a single predictor, context, or a conjunction of predictors, such as wind and slope. On the assumption that people, by default, commence learning by giving equal consideration to all potential predictors, the emerging dominance of context is predicted by laws of associative learning. Specifically, Rescorla and Wagner's (1972) theory of learning, which is isomorphic to the error-driven

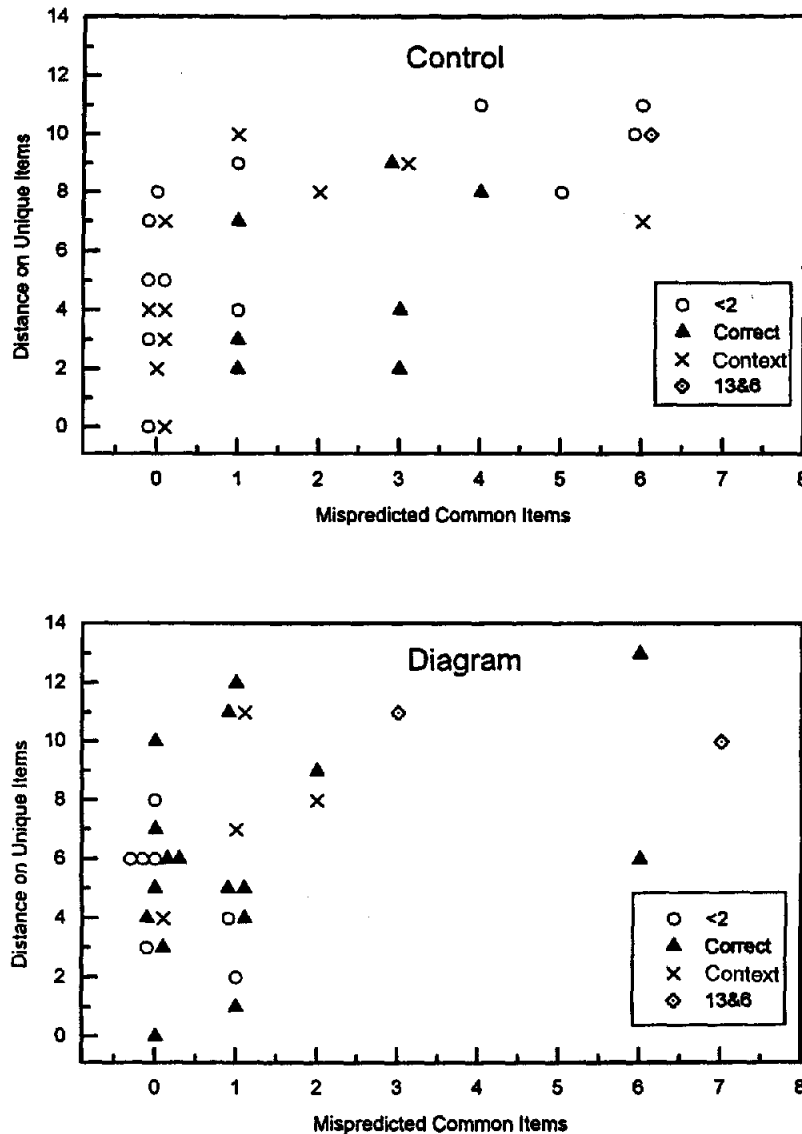


Figure 12. The number of mismatches ("distance") between each participant's responses and the predictions of the closest strategy on the unique items in Experiment 3 is shown as a function of the number of mispredicted responses on items for which all four strategies predict identical performance ("common items"). The control condition (top panel) shows less consistent responding than does the diagram condition (bottom panel). Context = Context rule.

learning embodied in most connectionist networks (e.g., Gluck & Bower, 1988), predicts that if a single cue and a pair of two others are equally predictive of an outcome, the single cue will acquire the stronger association. This is because when an incorrect prediction has been made, the entire error signal is focused on updating the predictive strength of a single cue, whereas if a pair of cues is present, the total signal is distributed across cues. Because the magnitude of the error signal is a determinant of the speed of learning, the weight of the single cue thus changes more rapidly than the weights of a pair of cues. It is important to note that this will only occur if the single predictor and the pair of potential predictors are of roughly equal validity.

In our experiments, wind and slope together predicted the outcome perfectly if they were combined nonlinearly, to form the

boundary shown in Figure 2. If wind and slope were instead combined linearly, their joint predictive value was roughly equal to that of context by itself, as can be ascertained by placing a linear boundary through the wind-slope space in Figure 2, and as shown for the Greater-Than-Two strategy in Experiment 3. Hence, if people commenced learning by considering all cues equally and linearly, an associative learning account would expect context to become dominant, exactly as observed during Phase 1 of the control and reveal conditions in the first two experiments. That emergence of context would be further facilitated by the first two blocks of training, during which it was a perfect predictor.

A further implication of associative learning laws is that, once dominant, context would be expected to remain dominant, even if additional information about the task becomes available. This

continued reliance on a predictor at the exclusion of additional learning is known as blocking (e.g., Kamin, 1969). *Blocking* refers to the fact that once a valid predictor has been associated with an outcome, a second stimulus that subsequently cooccurs with the predictor will not acquire much predictive power for that outcome on its own. For example, if A signals event O during initial training, and subsequently O is signaled by the combination of A and B, at a final test B on its own will not reliably elicit a response. Blocking is known to play a powerful role in human contingency judgments (e.g., Williams, Sagness, & McPhee, 1994) and, of greatest interest here, in categorization (e.g., Gluck & Bower, 1988; Shanks, 1991, 1993).

In our experiments, then, the diagram is best viewed as being functionally isomorphic to another predictor. Experiment 3 confirmed that the diagram does more than shift people's attention from context to wind and slope: Accordingly, we suggest that it rendered available an additional compound cue, consisting of the conjunction of wind and slope (WS). Because people demonstrably use the correct strategy when the diagram is shown at the outset (Experiment 3), it is likely that this compound cue is sensitive to the competition between wind and slope by combining the two predictors in a negatively correlated fashion. According to Rescorla and Wagner's (1972) original theory, learning rates and cue salience are free to vary between predictors, and it seems reasonable to assume that the diagram also selectively enhances the salience of wind and slope.

Nonetheless, an associative learning account would expect blocking of the new predictor (WS) if context has already been learned. This is because learning only occurs if a response deviates from the correct outcome: If responding is consistently correct, learning of additional cues or additional learning of existing cues is blocked, even if salience is selectively enhanced. In consequence, given the fairly high levels of performance achieved at the end of Phase 1, and the perfect performance the context rule would provide in the first two blocks of training during Phase 2, the diagram would not be expected to alter people's reliance on context when presented halfway through the experiment.

Inducing complex processing. Now consider the effects of introducing the diagram at the outset of learning. Continuing with the above account, this would again be functionally equivalent to the introduction of a compound cue (WS) combined with enhanced salience of the predictors W and S. Because learning is not blocked by prior reliance on context, people can be expected to rely on the compound cue to the extent that this is more salient and predictive than context alone.

The data from Experiments 1 and 3 confirm that people relied less on context when the diagram was presented at the outset. These results are analogous to the findings of Edgell (1983) and Edgell and Morissey (1987), who also showed a strong learning advantage for information that was presented at the outset compared with when it was presented a few trials after training commenced. In Edgell's (1983) study, a probability learning task either involved a single predictor or the conjunction of two cues. When the conjunction of cues was predictive from the outset, people relied on it more than when the conjunction became predictive after 40 learning trials with the single predictor.

Gradual shift toward expediency. A further finding of our first experiment was that people gravitated toward the expedient context rule even when they had started out using a more complex

strategy. At first glance, this presents a challenge to the associative learning framework, because it is not immediately clear why a proven predictor—in this case a complex strategy involving several cues—should be abandoned in favor of an expedient one-dimensional rule that, at best, can rival the complex strategy only by additional memorization of exceptional instances.

One possibility is that this finding presents an instance of retrospective revaluation (e.g., Shanks, 1985; Wasserman & Berglan, 1998). Retrospective revaluation occurs when a compound cue consisting of two distinguishable elements (e.g., AB) is first learned and is then followed by training of one of those elements in isolation (e.g., B). After the additional learning of B, people judge A to be less predictive than before, which is indicative of a retrospective revaluation of its role. In Experiment 1, this might have occurred if participants learned to rely on a compound consisting of all cues (i.e., the presumed cue WS and C) when the diagram was initially presented, and then used the first two blocks of the second phase, during which context was perfectly predictive on its own, to retrospectively reduce the perceived validity of WS.

Asymmetry of knowledge restructuring. We are aware of two instances of asymmetrical transfer effects that are similar to the asymmetry of knowledge restructuring observed here. Although those instances occurred in other domains and possibly reflect different processes, they deserve mention because any occurrence of asymmetry serves to place the present findings into a wider context.

Schyns and Rodet (1997) created a categorization task in which arbitrary visual shapes served as predictors that were presented embedded in other random "blobs." Two categories were created, each defined by one of two critical shapes (call those X and Y). A third category was defined by the concatenation of those two shapes into one (call that XY). The three categories were learned to perfection one at a time, and the critical manipulation was the sequence in which they were acquired. In one condition, participants learned categories in the order X, Y, and XY, whereas in another condition that sequence was reversed. Emphasis during the final transfer test was on a critical stimulus, labeled X-Y, in which both features (i.e., X and Y) were present but without being concatenated. Participants who learned categories in the order $X \rightarrow Y \rightarrow XY$ classified the critical X-Y item as belonging to category XY, suggesting that they had learned to rely exclusively on features X and Y, regardless of their adjacency, to categorize items in all three categories. Participants who learned the $XY \rightarrow X \rightarrow Y$ sequence, by contrast, classified the X-Y item as belonging to either category X or Y, but not XY, suggesting that these participants had additionally learned the configural feature XY.

The results of Schyns and Rodet (1997) showed that once single features (i.e., X and Y) have been learned, people are unable to acquire a separate representation for the compound cue XY. Conversely, learning of the compound cue XY does not prevent later acquisition of the elemental features X and Y. This asymmetric pattern mirrors the present results, although unlike the present experiments it involved changes to the learning environment (order of trials).

Finally, Edgell and Morissey (1987) examined the effect of additional information becoming relevant during learning of a probabilistic prediction task. In their study, participants had to predict an outcome from two cues whose validity was altered during the course of learning. As noted earlier, when a single cue

was predictive early in training, people did not detect the emerging validity of the conjunction of cues. Specifically, when a target color (red or green) could initially be predicted on the basis of shape alone (e.g., triangle → red vs. square → green), people did not learn the association between color and the conjunction of two predictors (shape and orientation of lines) when the configural validity of cues emerged later. Conversely, when training commenced with configural validity but one of the cues later became predictive on its own, people shifted to the new expedient strategy.

The overall pattern of results again mirrors the present findings, albeit with two critical differences in methodology. First, whereas Edgell and Morissey (1987) altered the relative validity of a configural cue and its constituent elements, the analogous manipulation in the present experiments involved a purely conceptual shift between a pair of cues (W and S) and a different single cue (C). Second, unlike the present studies, Edgell and Morissey changed the objective validity of cues across trials.

Conclusions

We reported two experiments that showed that people persist in using a previously acquired expedient rule in a categorization task even when an important underlying relationship among predictors is revealed. This occurred even when the relationship was visualized in an adaptive display. If people were informed about the underlying relation at the outset, they adopted a more complex strategy that provided an initial performance advantage. However, people gradually abandoned the complex strategy in favor of the expedient rule after extended training. A third experiment confirmed that people approach the task in qualitatively different ways, depending on whether they adopt an expedient or a complex strategy. This in turn implies that the diagram does more than simply shift attention to an alternative set of predictors.

We accounted for this asymmetry of knowledge restructuring with basic associative learning principles. Resistance to knowledge restructuring was seen to be a consequence of standard blocking, whereby an initially acquired expedient strategy prevented subsequent learning of an alternative. The gradual shift from a complex to a more expedient rule was seen to reflect retrospective reevaluation, whereby people reverted to the expedient rule when it was a perfect predictor early during the second training phase.

References

- Albini, F. A. (1984). Wildland fires. *American Scientist*, 72, 590-597.
- 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.
- Biederman, I., & Shiffrin, M. M. (1987). Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 640-645.
- Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, 120, 278-287.
- Davies, S. P. (1994). Knowledge restructuring and the acquisition of programming expertise. *International Journal of Human-Computer Studies*, 40, 703-726.
- Delaney, P. F., Reder, L. M., Staszewski, J. J., & Ritter, F. E. (1998). The strategy-specific nature of improvement: The power law applies by strategy within task. *Psychological Science*, 9, 1-7.
- Edgell, S. E. (1983). Delayed exposure to configural information in non-metric multiple-cue probability learning. *Organizational Behavior and Human Decision Processes*, 32, 55-65.
- Edgell, S. E., & Morissey, J. M. (1987). Delayed exposure to additional relevant information in nonmetric multiple-cue probability learning. *Organizational Behavior and Human Decision Processes*, 40, 22-38.
- 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. (1999). Rule and exemplar representation in rule-defined category structures. Manuscript submitted for publication.
- Erickson, M. A., & Kruschke, J. K. (2000). Rule-based extrapolation in perceptual categorization. Manuscript submitted for publication.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227-247.
- 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.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.
- Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing X-ray pictures. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise* (pp. 311-342). Hillsdale, NJ: Erlbaum.
- Lewandowsky, S., Dunn, J. C., Kirsner, K., & Randell, M. (1997). Expertise in the management of bush fires: Training and decision support. *The Australian Psychologist*, 32, 171-177.
- Lewandowsky, S., & Kirsner, K. (2000). Expert knowledge is not always integrated: A case of cognitive partition. *Memory & Cognition*, 28, 295-305.
- Lightfoot, N., & Shiffrin, R. M. (1992). On the unitization of novel, complex visual stimuli. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 277-282). Hillsdale, NJ: Erlbaum.
- Macho, S. (1997). Effects of relevance shifts in category acquisition: A test of neural networks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 30-53.
- Mayer, R. E., & Gallini, J. K. (1990). When is an illustration worth ten thousand words? *Journal of Educational Psychology*, 82, 715-716.
- 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., & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, 117, 68-85.
- 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., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.
- Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 548-568.
- Pearsall, N. R., Skipper, J. J., & Mintzes, J. J. (1997). Knowledge restructuring in the life sciences: A longitudinal study of conceptual change in biology. *Science Education*, 81, 193-215.

- Postman, L. (1971). Transfer, interference, and forgetting. In L. W. Kling & L. A. Riggs (Eds.), *Experimental psychology* (pp. 1019–1032). New York: Holt, Rinehart, and Winston.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 681–696.
- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology*, 37B, 1–21.
- Shanks, D. R. (1991). Categorization by a connectionist network. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 433–443.
- Shanks, D. R. (1993). Associative versus contingency accounts of category learning: Reply to Melz, Cheng, Holyoak, and Waldmann (1993). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 1411–1423.
- Shiffrin, R. M. (1996). Laboratory experimentation on the genesis of expertise. In K. A. Ericsson (Ed.), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports and games* (pp. 337–345). Hillsdale, NJ: Erlbaum.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127–190.
- Staszewski, J. J. (1988). Skilled memory and expert mental calculation. In M. Chi, R. Glaser, & M. Farr (Eds.), *The nature of expertise* (pp. 71–128). Hillsdale, NJ: Erlbaum.
- 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*, 51B, 121–138.
- 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.

Received June 14, 1999

Revision received April 13, 2000

Accepted April 20, 2000 ■