Notes and Comment

Single-system models and interference in category learning: Commentary on Waldron and Ashby (2001)

ROBERT M. NOSOFSKY and JOHN K. KRUSCHKE
Indiana University, Bloomington, Indiana

In a recent article, Waldron and Ashby (2001) observed that performing a concurrent task caused greater interference in learning a simple one-dimensional categorization rule than in learning a complex three-dimensional one. They argued that this result was incompatible with all existing single-system models of category learning but was as predicted by the multiple-system COVIS model (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). In contrast to Waldron and Ashby's argument, we demonstrate that the single-system ALCOVE model (Kruschke, 1992) naturally predicts the result by assuming that its selective-attention learning process is disrupted by the concurrent task.

An issue of major current interest in the field of categorization concerns the role of multiple systems in category learning. A number of theorists have recently developed quantitative models that assume multiple category-learning systems (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998). For example, according to Ashby et al.'s COVIS model, category learning takes place through a combination of explicit rule learning and a form of implicit procedural learning. By contrast, other researchers have questioned the need for positing such multiple-system models and have suggested that single-system models may be sufficient to account for a large number of the phenomena of interest. For example, Nosofsky and Johansen (2000) argued that selective-attention exemplar models, such as the generalized context model (Nosofsky, 1986) and ALCOVE (Kruschke, 1992), accounted in a natural way for a wide variety of such multiple-system phenomena.

In a recent article, Waldron and Ashby (2001) presented new experimental results that they argued provided further evidence for the role of separate category learning systems and that challenged single-system models such as ALCOVE. In this commentary, we dispute Waldron and Ashby's claim and argue that their experimental results are naturally predicted by ALCOVE. We

emphasize that our commentary does not present an argument concerning the merits of multiple-system versus single-system models in general. Rather, we argue that the specific form of evidence presented by Waldron and Ashby does not sharply distinguish these modeling approaches.

Waldron and Ashby (2001) tested subjects on two category structures. In both structures, the stimuli varied along four binary-valued dimensions. In the *simple* structure, only a single dimension was relevant. So, for example, objects with a value of 1 on Dimension 1 might belong to Category A, whereas objects with a value of 2 on this same dimension might belong to Category B. The remaining three dimensions were nondiagnostic and irrelevant to the task. In the *complex* structure, three dimensions were relevant, and the fourth was irrelevant. If a majority of the three relevant dimensions had a logical value of 1, the stimulus belonged to Category A; otherwise, the stimulus belonged to Category B.

In Waldron and Ashby's (2001) experiment, subjects learned these category structures under either standard single-task conditions (i.e., where the category-learning task was the only task being performed) or dual-task conditions, in which subjects performed a numerical Stroop task at the same time as they were learning the categories. The main results of interest are shown in Figure 1, which plots mean number of trials to criterion under each of the conditions. The most important result emphasized by Waldron and Ashby was that the concurrent Stroop task significantly impaired learning of the simple category structure involving the single-dimension rule but did not significantly impair learning of the complex category structure involving the three-dimensional rule.

Waldron and Ashby (2001) argued that these results were as predicted by their multiple-system COVIS model but were inconsistent with the predictions of single-system models, such as ALCOVE. They reasoned that "if ... there is only a single system (or processing resource) that operates on all rules, then the degree of concurrent task interference for the more complex rules should always be greater than or equal to the interference for the less complex rules" (p. 170). Because their results showed the opposite pattern, and because certain simulations that they conducted with the single-system ALCOVE model failed to predict their data, they concluded that "our results are incompatible with all *existing* single-system models of category learning" (p. 175).

We argue, however, that the single-system ALCOVE model very naturally predicts the pattern of results observed by Waldron and Ashby (2001). According to

This work was supported by Grants R01 MH48494-10 from the National Institute of Mental Health and BCS 9910720 from the National Science Foundation. Correspondence concerning this article should be addressed to R. M. Nosofsky, Department of Psychology, Indiana University, Bloomington, IN 47405 (email: nosofsky@indiana.edu).

Figure 1. Observed data from Waldron and Ashby (2001). The figure plots the mean trials to criterion for the single-dimension (1D) and three-dimensional (3D) rules under the control and concurrent task conditions.

ALCOVE, people represent categories by storing individual exemplars in memory. Classification decisions are determined by how similar an item is to the stored exemplars and by the degree to which the exemplars are associated with the alternative categories. A crucial component assumption in the model is that similarities among exemplars are influenced by an attention-learning process in which the component dimensions of the exemplars are differentially weighted. In general, ALCOVE learns to fine-tune the dimension weights in an adaptive manner by attending selectively to those dimensions that are relevant for solving a classification problem and by ignoring those dimensions that are irrelevant. As has been demonstrated in numerous previous studies, this attention-learning process can have a dramatic influence on the difficulty of solving alternative classification problems (e.g., Kruschke, 1992, 1993, 1996; Nosofsky, 1984, 1987; Nosofsky, Gluck, Palmeri, McKinley, & Glauthier, 1994; Nosofsky & Palmeri, 1996).

A highly likely consequence of the concurrent Stroop task manipulation used by Waldron and Ashby (2001) is that it would interfere with subjects' ability to learn to selectively attend to the relevant dimensions. Indeed, in explaining how the multiple-system COVIS model would account for their results, Waldron and Ashby themselves argued that the concurrent Stroop task would interfere with an observer's ability to selectively attend to a single stimulus dimension, thereby disrupting the explicit rule-learning component of the multiple-system COVIS model.

From the perspective of ALCOVE, it is intuitively clear that disrupting the attention-learning process can cause major interference in learning the simple one-dimensional category structure, without necessarily having a major impact on learning the complex three-dimensional structure. First, if the subject fails completely to attend to the single relevant dimension in the simple structure, his or her performance will be at chance. Second, if the subject does attend partially to the relevant dimension but also spreads attention to the remaining three irrelevant dimensions, he or she will be wasting an enormous amount of processing capacity. By contrast, a failure of attention learning in the more complex structure will tend to have less severe repercussions. First, because three of the four dimensions are relevant in the complex structure, a wide variety of attentional distributions will lead to above-chance and even reasonably good performance on this structure. Second, because there is only a single irrelevant dimension, little processing capacity will be wasted by spreading attention to this irrelevant dimension.

To verify these intuitions, we conducted extensive model investigations by simulating ALCOVE on the simple and complex problems tested by Waldron and Ashby (2001). (Details of the simulation procedures are reported in a brief appendix to this commentary.) ALCOVE has four free parameters: c , ϕ , λ_w , and λ_a . The c parameter is an overall sensitivity parameter that measures the overall discriminability of the stimuli; ϕ is a response-scaling parameter that influences the extent to which responding is probabilistic versus deterministic; λ_w determines the

rate at which exemplar-category associations are learned; and λ_a governs the attention-learning process—that is, the rate at which the system learns the attention weights. The critical question that we pursued in our simulations was the following: Assuming that the concurrent task reduced the attentional learning rate λ_a , would ALCOVE predict greater interference in learning the simple one-dimensional rule than in learning the complex three-dimensional rule? In a nutshell, across a wide variety of parameter settings, we observed precisely such a qualitative pattern of predictions.

We document this point in Figure 2, where we plot ALCOVE's predictions of mean trials to criterion in the simple and complex tasks as a function of variations in c , λ_w , and λ_a (for simplicity, we set $\phi = 4.5$ in all of these examples, but the same pattern of results holds for a wide range of values of ϕ). In each panel of Figure 2, c and λ_w are held fixed while λ_a varies. The figure reveals that, in general, as long as the sensitivity parameter c is not too high, ALCOVE predicts far greater interference on the simple one-dimensional task than on the complex three-dimensional task as the attentional learning rate λ_a is reduced toward zero. Contrary to Waldron and Ashby's (2001) claim, the pattern of interference observed in their category-learning tasks is precisely as predicted by the single-system ALCOVE model when it is assumed that the concurrent task interferes with attentional learning.¹

Waldron and Ashby (2001) did conduct certain modeling analyses to corroborate their claim that ALCOVE failed to predict greater concurrent-task interference for the one-dimensional rule than for the three-dimensional one. In particular, they first found parameter values that allowed ALCOVE to fit the data from the control conditions in which the concurrent Stroop task was not performed. Next, to test whether ALCOVE could predict the pattern of concurrent-task interference, they considered the predictions from ALCOVE if all previously estimated parameters were held fixed, with only a single free parameter allowed to vary. They reported that regardless of whether c , ϕ , λ_w , or λ_a was allowed to vary, ALCOVE failed to capture the qualitative pattern of results in their study. Obviously, given the results of the simulation analyses that we have reported in this commentary, in which only the λ_a parameter varied across the control- and concurrent-task conditions, Waldron and Ashby's claim about the behavior of ALCOVE is not a very general one. In part, Waldron and Ashby may have missed the potential of ALCOVE to account for their patterns of concurrent-task interference because they used a relatively large value of the sensitivity parameter c in their simulations. As we have illustrated in Figure 2, the predicted interaction effect that results from reducing λ_a is diminished for higher values of sensitivity.² In addition, inspection of Waldron and Ashby's Figure 2D reveals that they did not plot ALCOVE's predictions for values of λ_a less than approximately .01, but it is at these values of the attention-learning parameter where much of the "action" takes place (see Figure 2 of

the present article). Even with the parameter values that they considered, however, Waldron and Ashby's own simulation results do not fully support their conclusions. As is revealed by inspecting Waldron and Ashby's Figure 2D, given their assumed parameter values, when the magnitude of λ_a is reduced, the predicted trials to criterion for the single-dimension task increases, whereas the predicted trials to criterion for the three-dimensional task actually decreases. Thus, although the quantitative effect may not have been large, the same qualitative pattern of results that is evident from our simulations is present in Waldron and Ashby's simulations as well.

Furthermore, it does not seem sensible to us to draw very general conclusions about the behavior of ALCOVE (or any model) by fitting it to the control data and then constraining all of its parameters except one. First, because the control data consist of only two data points, whereas ALCOVE has four free parameters, it is likely that a large number of different parameter settings could allow ALCOVE to fit these data. Choosing one particular parameter configuration from this large number of possibilities does not allow for very general conclusions. Second, there is no guarantee that a particular experimental manipulation will have its effect by influencing just a single model parameter. Although our results demonstrate that allowing just the λ_a parameter to vary across conditions is sufficient to account for the patterns of concurrent-task interference, it does not follow logically that only a single model parameter could change. Waldron and Ashby (2001, p. 173) introduced this constraint on the grounds that the "concurrent task data contained only two degrees of freedom." This state of affairs, however, is a result of a design feature of their experiment; it is not the model's fault that the data have only two degrees of freedom. It seems advisable to test alternative models by using richer sets of parametric data, rather than relying on qualitative patterns of results involving just two data points. Before ruling out an entire class of models, one should conduct an investigation over a reasonably large region of the models' available parameter space.

An interesting issue is raised as a result of the modeling demonstrations reported in this article. Note that ALCOVE incorporates four free parameters corresponding to different processes (i.e., separate free parameters corresponding to response selection, memory sensitivity, exemplar-category association learning, and dimensional attention learning). In a sense, therefore, ALCOVE, too, can be considered a "multiple-system" model, and so its success may be viewed as not contradicting Waldron and Ashby's (2001) conclusions. We believe, however, that rendering this interpretation makes the debate between "single-system" versus "multiple-system" models empty and uninteresting. Essentially, all formal models designed to account for complex behavioral phenomena require multiple parameters corresponding to distinct processes, and so, in this loose sense, all models are "multiple-system" models. The key issue

Figure 2. ALCOVE's predicted mean trials to criterion for the single- and three-dimensional rules as a function of variations in c , λ_w , and λ_a . In each individual panel of the figure, c and λ_w are held fixed, with only λ_a allowed to vary. Each row of panels corresponds to a fixed value of c , whereas each column of panels corresponds to a fixed value of λ_w . Within each panel, the bars to the left correspond to ALCOVE's predictions for the single-dimension (1D) rule, whereas the bars to the right correspond to the three-dimensional (3D) rule. Values of λ_a within each panel are equally spaced on a logarithmic scale except for the value $\lambda_a = 0$.

that distinguishes modern single-system and multiple-system models, we believe, is whether distinct *representational* systems are involved. For example, in Ashby et al.'s (1998) COVIS model, one system represents categories in terms of easily verbalizable, explicit rules, whereas a second system represents categories by using nonverbalizable, "implicit" decision boundaries. By contrast, in ALCOVE, a single representational system based on the storage of individual exemplars is assumed to mediate category learning. Waldron and Ashby seem to have had something similar in mind, because they too considered ALCOVE to be a prime representative of the class of "single-system" categorization models.

In summary, Waldron and Ashby (2001) observed greater concurrent-task interference on a simple one-dimensional category structure than on a complex three-dimensional one. They interpreted this result as being inconsistent with all existing single-system models of category learning. In this commentary, however, we showed that the result is naturally predicted by the single-system ALCOVE model by assuming that the concurrent task interfered with ALCOVE's attention-learning process. Indeed, our simulation analyses revealed that, given a disrupted attention-learning process, ALCOVE is highly constrained to predict the pattern of interference effects reported by Waldron and Ashby. We think it is likely that numerous other single-system models that place emphasis on the role of attention learning in categorization would also predict the result. We reemphasize that we are not claiming that single-system models can account for all phenomena to which modern multiple-system models have been applied. Rather, we claim that the specific form of evidence reported by Waldron and Ashby does not distinguish between the modern single-system and multiple-system approaches.

REFERENCES

- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN A. U., & WALDRON, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, **105**, 442-481.
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**, 107-140.
- KRUSCHKE, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22-44.
- KRUSCHKE, J. K. (1993). Human category learning: Implications for backpropagation models. *Connection Science*, **5**, 3-36.
- KRUSCHKE, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science*, **8**, 201-223.
- KRUSCHKE, J. K., & JOHANSEN, M. J. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 1083-1119.
- NOSOFKY, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **10**, 104-114.
- NOSOFKY, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**, 39-57.
- NOSOFKY, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 87-109.

- NOSOFKY, R. M., GLUCK, M. A., PALMERI, T. J., MCKINLEY, S. C., & GLAUTHIER, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*, **22**, 352-369.
- NOSOFKY, R. M., & JOHANSEN, M. J. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review*, **7**, 375-402.
- NOSOFKY, R. M., & KRUSCHKE, J. K. (1992). Investigations of an exemplar-based connectionist model of category learning. *Psychology of Learning & Motivation*, **28**, 207-250.
- NOSOFKY, R. M., & PALMERI, T. J. (1996). Learning to classify integral-dimension stimuli. *Psychonomic Bulletin & Review*, **3**, 222-226.
- WALDRON, E. M., & ASHBY, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, **8**, 168-176.

NOTES

1. Moreover, various parameter settings from Figure 2 yield good quantitative fits to Waldron and Ashby's (2001) trials-to-criterion data. Because there are only four degrees of freedom in their data, however, we dispense with quantitative model fitting in the present case and focus instead on the broad qualitative pattern of predictions made by ALCOVE.

2. The reason is that high values of c tend to isolate individual exemplars from one another in the psychological similarity space. Because there is little resulting interaction among exemplars in ALCOVE's error-driven learning process, effects of attention to dimensions are thereby weakened.

APPENDIX

In each simulation of ALCOVE, a unique random sequence of 200 learning trials was generated, and ALCOVE was used to predict the probability of a Category A or B response for each trial. Let p denote the probability of a predicted Category A response. A random number in the interval (0,1) was selected. If the random number was less than or equal to p , a Category A response was chosen; otherwise, a Category B response was chosen. The learning process continued until the model achieved a learning criterion of eight consecutive correct responses, which was the trials-to-criterion measure used by Waldron and Ashby (2001). The results shown in Figure 2 are based on averages across 1,000 such simulations. Following Waldron and Ashby, individual simulations in which the model did not achieve the learning criterion were not included in generating these predictions.

The specific version of ALCOVE that we report in the Figure 2 simulations followed the one described in Kruschke's (1992) article in all respects except the following. First, in Kruschke's (1992) modeling, a uniform distribution of initial attention weights was used (i.e., all of the individual attention weights were set equal to one another at the start of training). By contrast, in the modeling reported in our Figure 2, we followed Waldron and Ashby's (2001) procedure of setting the initial attention weights at random values, subject to the constraint that they be nonnegative and sum to one. We emphasize, however, that we conducted our simulation analyses with both procedures and found that both led to the same qualitative pattern of predictions of interference effects in the present paradigm.

A second difference from Kruschke's (1992) original procedure was that the learning algorithm imposed the constraint that, on each attention-weight update, all attention weights be

nonnegative and that the weights sum to one. These attention-weight constraints form a fundamental assumption of Nosofsky's (1984, 1986) GCM, which is ALCOVE's direct ancestor, and have been imposed as well in most subsequent applications of ALCOVE (e.g., Kruschke & Johansen, 1999; Nosof-

sky et al., 1994; Nosofsky & Kruschke, 1992; Nosofsky & Palmeri, 1996). Without imposing these constraints on the attention-weight parameters, ALCOVE still predicts the same pattern of interaction effects of the concurrent task on category learning, but the quantitative differences are not as pronounced.

(Manuscript received July 11, 2000;
revision accepted for publication February 15, 2001.)