Neural Network Models of Discrimination Shifts

Sylvain Sirois (sirois@psych.mcgill.ca)
Department of Psychology; McGill University, 1205 Dr Penfield Avenue
Montréal, Qc H3A 1B1 Canada

Thomas R. Shultz (shultz@psych.mcgill.ca)
Department of Psychology; McGill University, 1205 Dr Penfield Avenue
Montréal, Qc H3A 1B1 Canada

Abstract

The importance of discrimination shifts to learning and developmental psychology is highlighted. Basic tasks used in continuous and total change paradigms are presented, and major theoretical accounts are briefly reviewed. The lack of a general and comprehensive interpretation of human shift learning is identified, and a recent model based on neural network research is described. This model suggests that human adult performance in discrimination shifts differs from preschool performance because of a process called spontaneous overtraining. This hypothesis has been previously used in neural network simulations to successfully capture developmental regularities in continuous discrimination shifts (e.g., reversal and nonreversal shifts). In the present paper, new simulations using this model are applied to total change discrimination shifts (e.g., intradimensional and extradimensional shifts). Several developmental regularities are successfully captured by the networks. The contribution of the spontaneous overtraining hypothesis is discussed.

Introduction

Discrimination learning involves learning to make different responses to particular stimulus conditions. Specifically, participants must learn to reliably identify among competing stimuli the one that exhibits a single, several, or combinations of attributes. The vast psychological literature on discrimination learning offers a substantial database of empirical regularities about learning and development (e.g., Esposito, 1975; Wolff, 1967). These findings on both child and adult learning are of significant importance for theories of human learning, cognition, and cognitive development. A rigorous way to develop and test such theories is through computational modeling. Because many neural networks use learning algorithms and have been favored for modeling a variety of perceptual, cognitive, and learning phenomena, discrimination learning tasks provide a useful benchmark for the adequacy of neural networks as models of human cognition.

Surprisingly, there are few studies that report connectionist modeling of discrimination shift learning. Of these, only one model successfully captures the human ontogeny of reversal shift and nonreversal shift performance (Sirois & Shultz, submitted).

In the present paper, we extend the application of this model to intradimensional and extradimensional shifts. The first section presents the discrimination shift tasks that this research addresses, as well as the psychological regularities associated with them. In the second section, we review previous theoretical interpretations of discrimination learning. The third section presents our cascade-correlation model of discrimination learning and new simulations that model intradimensional and extradimensional shifts. The discussion focuses on the implications of our model for a general theory of discrimination learning.

Discrimination Shifts

The discrimination shift tasks we consider involve the pairwise presentation of stimuli with varying attributes on three binary dimensions (e.g., shape, color, and position). In each pair, the stimuli are constrained such that they exhibit mutually exclusive combinations of the attributes on all three dimensions. Figure 1 presents four stimulus pairs that exhaust such a combination of shape, color, and position.

Participants in these tasks are required to consistently identify in each pair the stimulus that exhibits the attribute targeted by the experimenter (e.g., square). They are repeatedly presented with the pairs of stimuli, and they are provided with reinforcement feedback on each trial. Learning continues until the participant reliably identifies the target stimulus, typically on eight out of ten consecutive

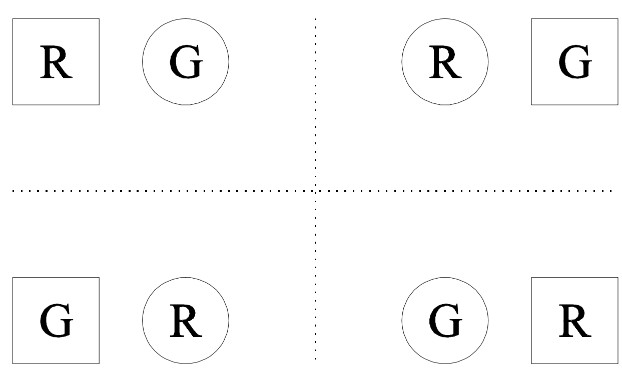

Figure 1: Pairs of stimuli that exhaust the mutually exclusive combinations of shape, color, and position. R denotes red, and G denotes green (from Sirois and Shultz, submitted).

trials. When this initial criterion is reached, shifts in reward contingencies may be introduced. Performance on such shifts often reveals important features of the leaning. For example, shift performance might reveal whether the learner is using concepts or attentional responses to mediate the associations between stimuli and responses. We elaborate this later in a presentation of different theoretical interpretations of shift learning. Participants are not usually told explicitly about the introduction of a shift.

In a reversal shift (RS), the stimuli that exhibit the other attribute of the initial dimension are now associated with reward (e.g., circle instead of square). Participants must consequently change their responses on all pairs. This is shown in the first row of Figure 2.

A nonreversal shift (NS) involves a shift to an attribute of a previously irrelevant dimension (e.g., red instead of square). In this task, only half of the responses must be changed. In our example, half of the square stimuli are also red. The second row of Figure 2 presents a NS task.

Because RS and NS use the same attributes in both learning phases, they are referred to as continuous paradigm tasks. Two other discrimination shift tasks also involve a complete or a partial shift in reinforcement contingencies. These are the intradimensional (IDS) and extradimensional (EDS) shifts. They are known as total change tasks, because new attributes of the initial dimensions are introduced at the onset of the shift (Esposito, 1975).

In an IDS, stimuli that exhibit a specific novel attribute of the previously relevant dimension are associated with reinforcement (e.g., from square to diamond). On the other hand, an EDS involves a shift to a novel attribute of a previously irrelevant dimension (e.g., from square to yellow). Both tasks are shown in the bottom rows of Figure 2.

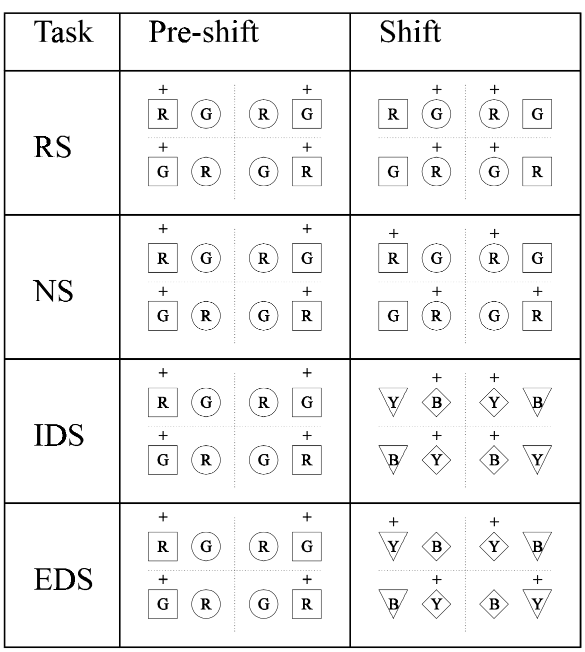

Figure 2: Examples of RS, NS, IDS and EDS. Plus signs identify reinforcement. R, G, B, and Y denote red, green, blue, and yellow, respectively.

Continuous and total change paradigms are highly similar. RS and IDS tasks involve a shift within the initially relevant dimension, and NS and EDS tasks involve a shift to a previously irrelevant dimension. The important distinction is that total change paradigms introduce new stimuli at the onset of the shift, which makes the shift more obvious to participants.

Decades of research have identified robust psychological regularities within the continuous and total change paradigms. Of importance, the ease of executing a shift, as measured by the number of trials to criterion, has been shown to vary between tasks and age groups. Children above the age of 10 years and adults reach the shift criterion in a RS quicker than in a NS, and reach the shift criterion quicker for IDS than for EDS (Esposito, 1975; Wolff, 1967). Preschool children also execute an IDS faster than an EDS, but they execute a NS as quickly as a RS (Esposito, 1975; Wolff, 1967).[1] Between the ages of 4 and 10 years, RS becomes easier than NS (Esposito, 1975). Although some studies report comparisons between continuous and total change paradigms (Esposito, 1975), we have failed to find well replicated, unequivocal regularities in such comparisons.

A final psychological regularity is that when trained several trials beyond the usual success criterion in the initial learning phase, most preschoolers then execute a RS faster than an NS, as do adults. This is called the overtraining effect, and is significant for the design of our neural network model.

Theoretical Interpretations

Three major accounts of discrimination learning have been presented over the years to account for a variety of findings (Sirois & Shultz, submitted). These are the Kendlers' two-stage theory (Kendler & Kendler, 1975; Kendler, 1983), Zeaman and House's attentional theory (Zeaman & House, 1963), and the Tighes' perceptual differentiation theory (Tighe & Tighe, 1966).

The Kendlers worked primarily within the continuous paradigm, and their model fairs poorly when applied to total change data (Sirois & Shultz, submitted). They argue that older children and adults use covert categorical responses to mediate between the stimuli and overt behavior (Kendler & Kendler, 1975; Kendler, 1983). These responses represent the specific attributes of the stimuli. Participants thus learn to respond to the relevant attribute and use this covert response to produce overt behavior. For the Kendlers, a RS is easier than a NS because only the link between the covert and overt responses needs to be changed. In a NS, the previous covert response has to be extinguished, and a new covert response must be learned, as well as the appropriate

[1] This last finding is in contradiction with a pervasive belief that preschoolers execute a NS faster than a RS (Kendler, 1983). An extended discussion of this controversy, in our paper on continuous shift paradigms (Sirois & Shultz, submitted), shows that the confusion stems from misinterpretations of results that are confounded with certain methodological variations.

response to the categorical response. This implies more learning than required for a NS. Because the covert responses represent the discrete attributes of the stimuli, though, their mediational model does not account for the regularities observed in total change tasks (Sirois & Shultz, submitted; Wolff, 1967).

The Kendlers suggested that preschoolers, unlike older children and adults, behaved according to an associative model (Kendler, 1983). Mere associations between stimulus and response are involved in their shift performance. This model predicts that a NS is easier than a RS, because more responses must be changed in the latter. The prediction has little support in the literature (Sirois & Shultz, submitted).

Zeaman and House, on the other hand, developed a model within the total change paradigm (Zeaman & House, 1963). They suggested that the stimuli were associated with an attentional response, which in turn was associated with an overt response. The attentional responses involve the different dimensions of variation, and not the discrete values of these dimensions. Following the same logic as that found in the Kendlers' model, this model predicts that an IDS will be easier than an EDS. Indeed, only the link between the attentional response and overt behavior needs to be changed in an IDS. In an EDS, both a new attentional response and the appropriate association with overt behavior must be learned. It also predicts that a RS will be faster than a NS, for the same reasons.

What is novel in Zeaman and House's model is that mediation is also involved in preschoolers, and different learning parameter values would explain their distinct behavior. Unfortunately, it also predicts easier RS over NS in preschoolers, which is not supported by the literature.

Finally, Tighe and Tighe's perceptual differentiation model does not imply mediation of stimuli into covert responses (Tighe & Tighe, 1966). Rather, the compound stimuli are differentiated in specific overt responses as a function of reinforcement. Older children and adults, compared to preschoolers, perform differently due to their larger amount of perceptual experience. This enables them to better focus on the relevant dimension and ignore the irrelevant one. The model predicts that RS and IDS are easier than NS and EDS, respectively, because participants are attending to the dimension that remains relevant.

Preschoolers, on the other hand, poorly differentiate the stimuli into relevant dimensions and instead respond to the compound properties of the stimuli (Tighe & Tighe, 1966, 1978). They associate the appropriate response with the poorly differentiated stimulus pair. As such, they cannot use information about the relevant dimension as an advantage in shifts within this dimension using the same stimuli (i.e., RS). And because a RS requires more relearning than a NS, this model also wrongly predicts faster NS. In total change paradigms, though, the model suggests a faster IDS than EDS (Tighe & Tighe, 1978). This is because the new stimuli prevent an interference of the previous associations, which were bound to the initial stimuli. The minimal differentiation acquired during preshift learning then becomes helpful for a shift within the same dimension (i.e., IDS).

All three theories have their own limitations, which prevented the formulation of a general and comprehensive account of discrimination learning (Esposito, 1975; Sirois & Shultz, submitted). This in spite of the fact that work in discrimination learning began over sixty years ago (Kendler, 1983). Because of the important issues of learning and development raised by this literature, the formulation of a comprehensive theoretical account of shift learning is still relevant today.

A Cascade-correlation Model

We have recently applied the cascade-correlation neural network algorithm to continuous paradigm tasks (Sirois & Shultz, submitted). The algorithm successfully captured the psychological regularities in preschool and adult behavior. Networks parameterized as adults executed a RS faster than a NS, whereas networks parameterized as preschoolers executed both shifts equally quickly. The networks also captured other developmental regularities associated with the continuous paradigm that are not discussed here (i.e., trial-by-trial behavior in shift learning and optional shifts). But the success of the model is of limited interest if it cannot capture the regularities in total change paradigms as well.

Central to our model is the suggestion that older children and adults, compared to preschoolers, submit themselves to extended processing of stimuli and reinforcement through a process similar to rehearsal (Sirois & Shultz, submitted). This had already been suggested in the discrimination learning literature by Levine (1975). By using a lower score threshold, neural networks are submitted to additional training trials. This is because, in cascade-correlation, training continues until all output activations are within score threshold of their targets.[2] We used a score threshold of 0.01 to model adult performance, and the default value of 0.4 to model preschool performance. Using a lower score threshold allows networks to learn a problem to greater depth and precision. All other parameters in the algorithm were set to the default values of cascade-correlation (Fahlman & Lebiere, 1990).

This adjustment of the score threshold parameter resulted in capturing the developmental effects in continuous paradigms with networks that lack hidden (i.e., mediational) units, because cascade-correlation will not install any for such linear problems. It is consistent with developmental changes in spontaneous rehearsal, as well as with the overtraining effect (Sirois & Shultz, submitted). Adult-like performance in continuous discrimination shifts can be achieved through extended training with a low score-threshold. We now report the application of this model to total change tasks.

Simulation of Continuous and Total Change Tasks

In these simulations, we used networks with eight input units. The first two units coded shape of the left stimulus, the next two units color of the left stimulus, the following

[2] The same thing could be accomplished in backpropation networks by lowering the error criterion.

two units the shape of the right stimulus, and the final two units color of the right stimulus. These units were connected to two output units, with initially random connection values. Target output was [0.5, -0.5] when the correct stimulus was on the left, and [-0.5, 0.5] when it was on the right.

We used two input units for each attribute in each stimulus because we needed to represent four attributes in the total change paradigms (two initial attributes, and two new attributes at the onset of the shift). The four possible attributes of each dimension were represented by a combination of -1 and 1 values. For example, four possible shapes were coded as [-1, -1], [-1, 1], [1, -1], and [1, 1]. Because each attribute is a combination of values of the same magnitude (all 1s), pre- and post-shift attributes have equivalent salience.

Networks parameterized as adults are expected to replicate our previous finding that a RS is performed faster than a NS. They should also execute an IDS quicker than an EDS. Networks parameterized as preschoolers should also perform the IDS quicker than the EDS. As we have previously observed, though, RS and NS should be learned at equivalent rates.

Method Two hundred and eighty adult networks were used in this simulation, with a score threshold of 0.01. One hundred and forty were initially trained on one attribute of color, and one hundred and forty on another attribute of color. When performance reached threshold on all problems of the initial discrimination, training was shifted to another attribute. In each subset of one hundred and forty networks, training was shifted to the other attribute of color for twenty networks (RS). For forty networks, it was shifted to an attribute of the previously irrelevant dimension (NS, n = 20 per attribute). For the remaining eighty networks, new dimensional attributes were introduced. Forty networks were trained on a new attribute from the previously relevant dimension (IDS, n = 20 per attribute), and the remaining forty were trained on a new attribute from the previously irrelevant dimension (EDS, n = 20 per attribute). Learning continued until criterion was reached in the shift training phase (i.e., output activations were within threshold of the target values for both output units on all problems). Two hundred and eighty networks parameterized as preschoolers were used under the same conditions, with the score-threshold set at 0.4.

To assess how quickly networks perform each task, we recorded the number of epochs required to reach criterion in shift learning. An epoch is, in this case, a block of four trials, one with each stimulus pair. For control purposes, we also recorded epochs to criterion for the initial learning phase.

Results There were no significant differences between any of the groups in the preshift phase, for networks parameterized as adults ($F(3,276) = 2.56$, n.s.). There were no significant differences either in the initial phase for networks parameterized as preschoolers ($F(3,276) = 1.06$, n.s.). Table 1 presents the mean number of epochs required to learn each type of shift, for both types of networks.

Table 1: Mean number of epochs to criterion for shift training.

	Network Parameterization	
Task	Adult	Preschool
RS	6.48	4.25
NS	10.08	4.28
IDS	8.43	2.93
EDS	9.69	3.59

For adult networks, the result of a One-Way analysis of variance show a significant difference between the groups ($F(3,276) = 136.09$, $p < .001$). We performed Scheffé post-hoc comparisons on the data. Table 2 presents the significant differences between the groups (significant test values are between 21.08 and 114.26).

In child networks, the results of a One-Way analysis of variance also show a significant difference between the groups ($F(3,276) = 52.189$, $p < .001$). We report significant Scheffé post-hoc comparisons in Table 2 as well (significant test values range between 6.98 and 43.44).

Discussion The simulation results reported in the first two rows of Table 2 are consistent with the psychological regularities reported for both adults and preschoolers. Networks parameterized as adults execute a RS faster than a NS, and an IDS quicker than an EDS, as normal adults do. And networks parameterized as preschoolers execute an IDS quicker than an EDS, even though they perform equally fast on RS and NS, like preschoolers do.

Network analyses in our previous simulations suggested that extensive training in adult networks yields finely tuned discriminations that focus on the relevant dimension and ignore the irrelevant one, which enable faster RS over NS (Sirois & Shultz, submitted). The adult networks in these new simulations also executed the RS faster than the NS. Because such discriminations remain relevant in IDS (the shift is within the initial dimension), this also enables the networks to execute an IDS faster than the EDS.

For networks parameterized as preschool children, though, we have argued that their behavior was a function of the compound properties of the pair of stimuli (Sirois & Shultz, submitted). That is, the minimal amount of

Table 2: Post-hoc comparisons between groups. Significant differences are represented by < (the group on the left took less time than the one on the right) or > (the group on the left took more time than the one on the right). The equal (=) sign indicates no significant difference between the two groups.

	Network Parameterization	
Comparison	Adult	Preschool
RS vs. NS	<	=
IDS vs. EDS	<	<
RS vs. IDS	<	>
RS vs. EDS	<	>
IDS vs. NS	<	<
NS vs. EDS	=	>

processing they perform does not allow fine discriminations between the dimensions to be made. This prevents an advantage of RS over NS. But in the case of total change tasks, the introduction of new attributes may remove the influence of the initial stimuli and allow the minimal abstraction achieved in the initial phase to favor an IDS over an EDS. These generalizations would otherwise be masked by the influence of the initial attributes in continuous tasks, as Tighe and Tighe (1978) have argued. Further network analyses are required before general conclusions are drawn.

The results reported in rows 3-6 in Table 2 can be taken as predictions made by the model. As we have noted previously, these comparisons have either not been performed, or have led to equivocal findings. Our simulation results should be evaluated in light of any new evidence from such psychological comparisons.

Finally, the reader may notice that preschool networks take fewer epochs to learn the tasks than adult networks do. This is in contradiction with human data (e.g., Wolff, 1967). Our assumption, though, is that older children and adults spontaneously submit themselves to extended training. The number of epochs we report is an index of the amount of processing needed to learn the tasks, and should not be equated with the actual number of trials in humans. In particular, at least some of the epochs taken by networks parameterized as adults represent rehearsal of the patterns rather than discrete trials. Consequently, valid main effect comparison of epochs to learn between child adult networks cannot be made.

General Discussion

Our previous simulations of discrimination learning were based on the assumption that older children and adults spontaneously train themselves to a greater extent than preschoolers do (Sirois & Shultz, submitted). This hypothesis was derived from the overtraining literature, and is consistent with the development of active rehearsal in children. Indeed, changes in discrimination learning and spontaneous rehearsal overlap between the ages of 4 and 10 (Sirois & Shultz, submitted). More training results in better discrimination of the relevant dimension, which in turn allows shifts within the same dimension to be performed more quickly than shifts to previously ignored dimensions. In networks parameterized as preschoolers, nonselective encoding of the input prevents an advantage of shifts within the previously relevant dimension.

The new simulations reported here provide further support for the adequacy of this model. Continuous paradigm tasks like RS and NS may mask a minimal amount of abstraction by networks with preschool parameters. The introduction of new stimuli at the onset of the shift removes the influence of the initial stimuli and allows faster shifts within the same dimension. This had previously been suggested by Tighe and Tighe (1978). They argued that the distinct nature of continuous and total change paradigms stress different processes, and that the latter is more sensitive to dimensional discrimination than the former (i.e., it better detects generalization). This is because total change tasks assess generalizations acquired in preshift learning, without the influence of the specific material used to acquire the discrimination. Previous responses to the specific stimuli may exert too large an influence on shift learning to yield generalization effects in continuous paradigms.

Our model does not implement mediation, which in neural networks would require hidden units (Sirois & Shultz, submitted). As such, our model is more consistent with Tighe and Tighe's (1966) perceptual differentiation model than with the Kendlers' and Zeaman and House's mediational models. Unlike all three other models, though, ours has been able thus far to capture all regularities it was applied to in continuous and total change paradigms. It has not shown the limitations identified in the other models, including Tighe and Tighe's (Sirois & Shultz, submitted).

Our spontaneous overtraining hypothesis of discrimination learning requires further research before general theoretical claims can be made. One important step will be to evaluate the discrimination shift performance of older children and adults in conditions that would prevent processing of the stimulus pair beyond its presentation (e.g., by using a distracter task simultaneously). Our model would predict equal ease of RS and NS in conditions that prevent rehearsal, yet IDS should still be faster than EDS. Their performance is expected to be like that of preschoolers in such conditions, because only the amount of processing distinguishes both groups in standard tasks under the spontaneous overtraining hypothesis.

There are also a variety of related tasks on which our model should be tested. These include dimensionless shifts (e.g., Goulet & Williams, 1970) and compound categorization (e.g., Kruschke, 1996), which would provide a good test of the generality of our model (Sirois & Shultz, submitted). Hopefully, further work will represent positive steps towards a comprehensive account of human shift learning.

Acknowledgments

This research was supported in part by a Natural Sciences and Engineering Research Council of Canada (NSERC) graduate fellowship to the first author, and an NSERC operating grant to the second author.

References

Esposito, N. J. (1975). Review of discrimination shift learning in young children. *Psychological Bulletin*, *82*, 432-455.

Fahlman, S. E., & Lebiere, C. (1990). *The cascade-correlation learning architecture* (Tech. Rep. CMU-CS-90-100). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.

Goulet, L. R., & Williams, K. G. (1970). Children's shift performance in the absence of dimensionality and a learned representational response. *Journal of Experimental Child Psychology*, *10*, 287-294.

Kendler, H. H., & Kendler, T. S. (1975). From discrimination learning to cognitive development: A neobehavioristic odyssey. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 1).

Hillsdale, NJ: Lawrence Erlbaum Ass.
Kendler, T. S. (1983). Labeling, overtraining and levels of function. In T. J. Tighe, & B. E. Shepp (Eds.), *Perception, cognition, and development: interactional analysis*. Hillsdale, NJ: Lawrence Erlbaum.
Kruschke, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science*, *8*, 201-223.
Levine, M. (1975). *A cognitive theory of learning: Research on hypothesis testing*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Sirois, S., & Shultz, T. R. (submitted). *Neural network modeling of developmental effects in discrimination shifts*.
Tighe, L. S., & Tighe, T. J. (1966). Discrimination learning: Two views in historical perspective. *Psychological Bulletin*, *5*, 353-370.
Tighe, T. J., & Tighe, L. S. (1978). A perceptual view of conceptual development. In R. D. Walk & H. L. Pick, Jr (Eds.), *Perception and experience*. New-York: Plenum.
Wolff, J. L. (1967). Concept-shift and discrimination-reversal learning in humans. *Psychological Bulletin*, *68*, 369-408.
Zeaman, D., & House, B. J. (1963). The role of attention in retardate discrimination learning. In N.R. Ellis (Ed.), *Handbook of mental deficiency*. New York, NY: McGraw-Hill.