Cost of Attention as an Indicator of Category Learning

Hyungwook Yim (yim.31@osu.edu)
Department of Psychology & Center for Cognitive Science, The Ohio State University
209C Ohio Stadium East, 1961 Tuttle Park Place
Columbus, OH 43210, USA

Catherine A. Best (best.140@osu.edu)
Department of Psychology & Center for Cognitive Science, The Ohio State University
208G Ohio Stadium East, 1961 Tuttle Park Place
Columbus, OH 43210, USA

Vladimir M. Sloutsky (sloutsky.1@osu.edu)
Department of Psychology & Center for Cognitive Science, The Ohio State University
208C Ohio Stadium East, 1961 Tuttle Park Place
Columbus, OH 43210, USA

Abstract

Category learning often involves selective attention to category relevant information, which may result in learned inattention to category irrelevant information. This learned inattention is a cost of selective attention. In the current research, the cost of attention was used as an indicator of category learning. Participants were given a category learning task, and the amount of supervision given to them was manipulated. Along with behavioral data, recorded eye movements during the task showed signature patterns of learning via a cost of attention. In addition, a simple neural network (perceptron) was able to use these eye-tracking data to predict success in learning. Thus, the observed attentional pattern – the cost of selective attention – was proposed as an indicator of category learning.

Keywords: category learning, cost of attention, eye tracking, supervised learning, sparse category, classifier, perceptron, neural network

Introduction

Attention plays a central role in many models of category learning (Kruschke, 1992; Nosofsky, 1986). During category learning, the ability to selectively attend to category-relevant cues while ignoring category-irrelevant cues allows for more efficient category learning. However, attention should also be flexible to enable learning of new categories. Consider learning to discriminate plums from nectarines. The most efficient way to distinguish them visually would be focusing on color as a cue rather than shape. However, when encountering a new category like lemons versus bananas, the once useful color cue no longer helps, while the previously unhelpful cue of shape becomes a good dimension to efficiently learn the categories. The process of ignoring the shape cue in the first learning instance often results in learned inattention to this cue (Kruschke & Blair, 2000). Learned inattention to a previously irrelevant dimension creates a deficit in future learning. This deficit constitutes a cost of attention (Hoffman & Rehder, 2010). For example, Hoffman &

Rehder (2010) recorded eye movements during a supervised category learning task and found evidence for a cost of attention. If category learning involves selective attention, then a cost of attention could function as an indicator of category learning.

Overview of Current Experiments

Two eye-tracking experiments were conducted with identical stimuli. Experiment 1 tested participants in a two-phase supervised category learning task that should promote learning. Experiment 2 tested participants in a two-phase unsupervised category learning task that should prevent learning. Critically, the second phase of each learning task relied on previously irrelevant cues to learn a new category. Based on previous research, it was predicted that supervision would facilitate category learning in Experiment 1, compared to unsupervised learning in Experiment 2. However, there is also evidence with adults that supervision only facilitates category learning when a category does not have much structure (Kloos & Sloutsky, 2008). Category structure could be measured as category density or "a ratio of variance relevant for category membership to the total variance across members and nonmembers of the category" (Sloutsky, 2010). Therefore, categories that have many features in common and those features are not shared with non-members are statistically dense. On the other hand, categories that have few features in common while having many features that are common with non-members are statistically sparse. Thus, in the current research an artificial sparse category was used to manipulate learning via supervision.

The first aim of the current research was to replicate a condition with a cost of attention during category learning and a condition with no cost of attention during category learning by manipulating the amount of supervision provided to participants. The second aim was to use the demonstrated cost and lack of cost to classify adults into

learners and non-learners so that the signature patterns of attention during category learning could be systematically examined. To achieve this aim, a classifier using a simple neural network (perceptron) was used to predict individual learning data on the basis of eye gaze patterns during category learning.

Experiment 1

Methods

Participants Twenty-five adults with normal or corrected to normal vision participated in the experiment. Participants were undergraduate students at The Ohio State University participating for course credit. An additional 17 participants were excluded where 5 of them had problems tracking their eye movements, 4 missed the catcher trials (see Procedure), and 8 did not learn the category.

Stimuli For the practice phase, two artificial categories that look like molecules were used on a black background (see Figure 1). Molecules consisted of five circles where four changed colors in a binary fashion and one remained constant throughout the practice phase, serving as a category relevant feature. There were 16 category members for each molecule category where the only difference between the categories was the color of the constant feature. The category relevant feature was always the upper left circle, which was purple in one category and yellow in the other category. For the main experiment, four artificial categories that looked like flowers were used (see Figure 1). Presented on a black background, each exemplar had a dark gray hexagon in the middle that was constant for every category with six colored shapes on every side. Among the six colored shapes, five changed their color/shape in a binary fashion; whereas, one was constant, serving as a category relevant feature. Therefore, there were 32 unique stimuli for each category with two categories having the relevant feature on the right-bottom side of the hexagon (i.e., category A: purple triangle, category B: blue semi-circle) and two categories having the relevant feature on the left side of the hexagon (i.e., category C: orange square, category D: yellow pentagon). For the catcher trials (see Procedure), four flower-like colored shapes were used. These colored shapes were similar to the four artificial categories but had unique shapes and colors so that it was easier to distinguish them from the four artificial categories.

Procedure The experiment was controlled by E-prime version 2.0, and a Tobii T60 eye tracker was used to collect eye gaze with the sampling rate of 60Hz. Following a practice phase, the main experiment had two learning phases with each phase including 4 blocks where each block consisted of 8 learning trials followed by 4 test trials. All learning trials were presented for 1.5 seconds. Moreover, before each learning and test trial, participants fixated on a randomly placed fixation point (e.g., red cross) appearing on a random-dot background. The participants were told to look at the fixation to proceed with the experiment.
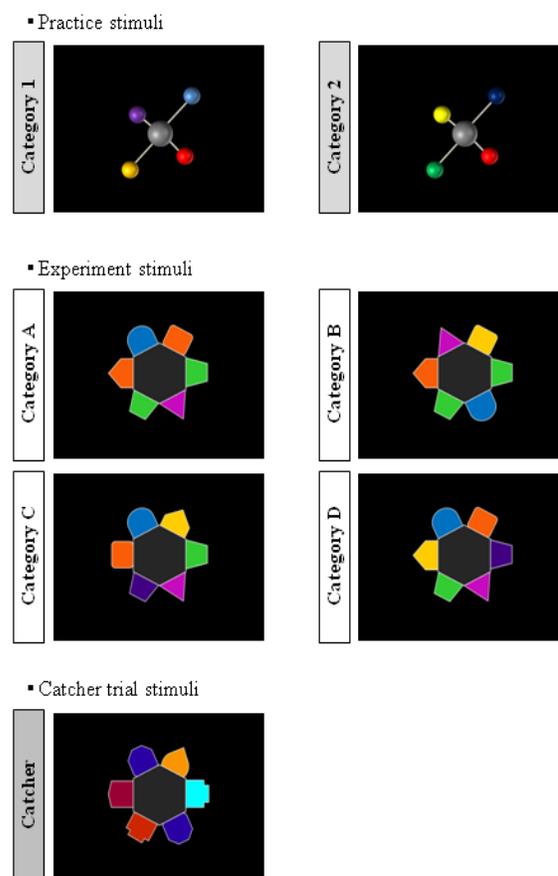

Figure 1. Stimuli used in the experiment.

In the practice phase, participants were told that they would see a set of molecules that have one common atom, and that they should try to find the common atom. After the practice learning trials, the participants were tested with two molecules side by side and were asked to choose the molecule that they had just seen by pressing a corresponding button on a keypad. The left/right side of the correct response was counterbalanced and onscreen feedback was provided as to whether the answer was correct or incorrect after every practice test trial. Moreover, a verbal explanation accompanied with actual stimuli explaining the relevant category was given after the four practice test trials. Instructions were also given before the four practice test trials to instruct participants to respond as fast and accurately as they could and to guess if they did not know the answer.

In the main experiment, each block was similar to the practice phase except that the stimuli were different and at the start of every block participants were told that they would see flowers that have a common feature they had to find. Moreover, unknown to the participants, the categories presented in the first four blocks (first learning phase) differed from the categories presented in the last four blocks (second learning phase). The categories differed by the position of the relevant feature. Therefore if the first half of

the blocks were presented with category A, the second half of the blocks were presented with category C or D. The category switching manipulation was intended to promote a cost of attention in the second half of the blocks, and the switch was never indicated to the participants. At test, two stimuli that had the relevant features at the same spatial location were presented side by side just as in the practice test. Therefore, if category A was learned, category B was presented with category A at test and same as for category C and D. Additionally, there was no feedback related to the test trials, and at the end of the last test trial of the second learning phase, there were four catcher trials with more obvious new category features on the novel item that were implemented to ensure the participants were paying attention during the entire experiment with continued motivation to participate. Moreover, participants who did not respond correctly to all four catcher trials were also excluded.

Results

To determine whether a cost was incurred, accuracy, reaction time, and eye gaze data were analyzed by block with a particular focus on the blocks before and after the unknown category switch, namely block 4 and block 5. Participants who were not accurate on the final four test trials in the last block of learning phase 1 were classified as non-learners and were excluded from the final analyses.

The mean accuracy for all test trials was 92.75% (learning 1: $M = 94\%$, $SD = 15.94\%$, learning 2: $M = 91.5\%$, $SD = 17.12\%$), which was significantly higher than chance performance, $p < .001$, for all blocks indicating learning occurred. Results of a 2×4 (Learning $\times$ Block) within-subjects ANOVA conducted on accuracy scores at test showed a main effect for Block, $F(2.29, 55.01) = 9.96$, $p < 0.001$, indicating that accuracy differed by block. Moreover, a significant cost of attention was demonstrated between the last block of learning phase 1 (block 4) and the first block of learning phase 2 (block 5) with a significant decrease in accuracy from block 4 to block 5, $t(24) = 4.27$, $p < .001$ (see Figure 2).

The mean reaction time for all test trials was 1032.47 ms ($SD = 566.66$ ms) (learning 1: $M = 991.46$, $SD = 538.11$, learning 2: $M = 1071.00$, $SD = 592.36$). A 2×4 (Learning $\times$ Block) within-subjects ANOVA conducted on mean reaction times showed a main effect for Block, $F(1.99, 37.84) = 15.20$, $p < .001$, but there was no significant main effect for Learning or a Block X Learning interaction. Statistical difference between block 4 and block 5 were also found, $t(23) = 4.36$, $p < .001$, demonstrating a cost of attention (see Figure 3).

Eye gaze data were also analyzed for each block by calculating the weighted proportion of looking to the relevant features. This value was calculated by taking the proportion of looking to the relevant features divided by the proportion of looking to the irrelevant and relevant features combined. However, since there was greater spatial area for irrelevant features (5 shapes) than the relevant features (1

shape), the proportion of looking to the relevant features was multiplied by five to equate the spatial area. Note that a value 0.50 in the analysis represents an equal amount of looking to the relevant and irrelevant features at a given time. Using the same method, the cost of attention was calculated by comparing the last four blocks of learning phase 1 to the first four blocks of learning phase 2 (see Figure 3).

As shown in Figure 4, the first 250 ms exhibit a random pattern of looking around 0.50 that reflected the first saccade away from a randomly moving fixation cross to the stimulus. In each learning phase, the proportion of looking to the relevant features increased across blocks, and in the last two blocks the proportion of looking to the relevant feature shows an asymptote. A 2×4 (Learning $\times$ Block) between-subject ANOVA by the group data only showed a main effect for Block, $F(3, 712) = 114.47$, $p < .001$. However, comparing the first block and the 5th block, where the 5th block is the start of the second learning phase, we observed a lower proportion of looking in the 5th block, which indicates a cost of attention, $t(178) = 2.68$, $p < .01$. A cost of attention was also found by comparing the last block of learning phase 1 ($M = .84$, $SD = .19$) to the first block of learning phase 2 ($M = .56$, $SD = .07$), $t(178) = 13.16$, $p < .001$. Moreover, the cost was more dramatic when comparing the last 4 trials of the first learning phase with the first 4 trials of the second learning phase, $t(178) = 16.4$, $p < .001$.

In sum, both behavioral and eye gaze patterns indicated a cost of attention for participants who learned the first category. Even though there were 8 non-learners, supervision was sufficient enough to help participant learn a sparse category, which is by definition harder to learn than a dense category (Kloos & Sloutsky, 2008).
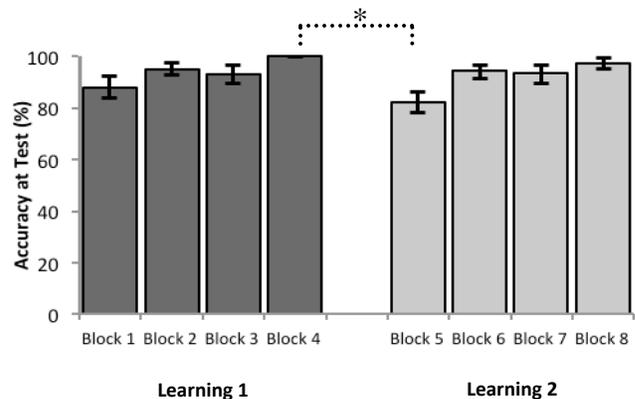

Figure 2. Accuracy at test in Experiment 1. Error bars represent +/- one standard error. *$p < .001$

Figure 3. Reaction time at test in Experiment 1. Error bars represent +/- one standard error. *$p < .001$

Figure 4. Eye gaze analysis for Experiment 1. Shaded area represents +/- one standard error. Weighted proportion of looking for learning phase 1 *(top)*. Weighted proportion of looking for learning phase 2 *(middle)*. Weighted proportion of looking before and after category switch (block 4 vs. block 5) *(bottom)*. Please see online version for colored line graphs.

Experiment 2

Methods

Participants Fourteen adults with normal or corrected to normal vision participated in the experiment. Participants were undergraduate students at The Ohio State University participating for course credit and none of the students participated in Experiment 1. An additional 13 participants were excluded where 8 of them had problems tracking their eye movements, 3 missed the catcher trials, and 2 learned the category.

Procedure The procedure was identical to Experiment 1 except there was no supervision given in the practice phase or main experiment. For the practice phase, participants were instructed that they would see different molecules one at a time but did not receive any feedback during or after the practice test trials. In the main experiment, the participants were told that they would merely see flowers one at a time. As in Experiment 1, participants who did not respond correctly to all four catcher trials were excluded from the analyses.

Results

To ensure no learning in the data, participants who performed perfectly on the four test trial in last block of learning were excluded from the analysis.

The overall accuracy for the test trials was 55.8% (learning 1: $M = 56.25\%$, $SD = 26.22\%$, learning 2: $M = 55.36\%$, $SD = 25.10\%$), where only block 5 was significantly different from chance, $p < .05$. However, block 5 still had lower accuracy compared to blocks in Experiment 1 ($M = 0.62$, $SD = 0.19$). A 2×4 (Learning × Block) within-subjects ANOVA conducted on accuracy scores at test showed no main effect of Block, $p = 0.41$, Learning, $p = .84$, or interactions, $p = .81$. Moreover, a comparison of the last block of learning phase 1 (block 4) and the first block of learning phase 2 (block 5) was not statistically different, $t(13) = 1.73$, $p = .12$, indicating there was no cost of attention (see Figure 5). The mean reaction time for all test trials was 1032.47 ms, $SD = 566.66$ ms (learning 1: $M = 991.46$ ms, $SD = 538.11$ ms, learning 2: $M = 1071.00$ ms, $SD = 592.36$ ms). A 2 X 4 (Learning × Block) within-subjects ANOVA conducted on mean reaction times revealed no significant results. Furthermore, no evidence for a cost of attention was found by comparing block 4 and block 5, $t(12) = .66$ $p = .52$ (see Figure 6).

Eye gaze data were analyzed in the same way as in Experiment 1. The same random looking pattern around 0.5 at the first 250 ms was also observed as in Experiment 1. There was a significant main effect of Learning, $F(1, 712) = 196.60$, $p < .001$, and of Block, $F(3, 712) = 16.79$, and a Learning × Block interaction, $F(3, 712) = 15.698$, $p < .001$. However, there were no significant theoretical patterns between the learning phases or blocks such as learning optimization (Blair, Watson, & Meier, 2009) or the cost of attention, block 4 ($M = 0.18$, $SD = 0.16$) versus block 5 ($M = 0.43$, $SD = 0.17$), $p < 0.01$ (see Figure 7).
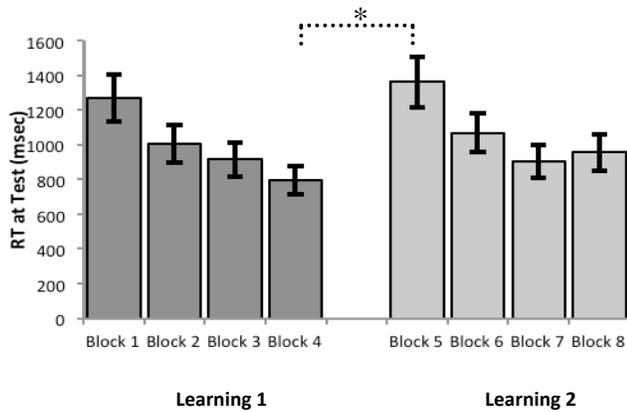
Figure 5. Accuracy at test in Experiment 2. Error bars represent +/- one standard error.

Figure 6. Reaction time at test in Experiment 2. Error bars represent +/- one standard error.
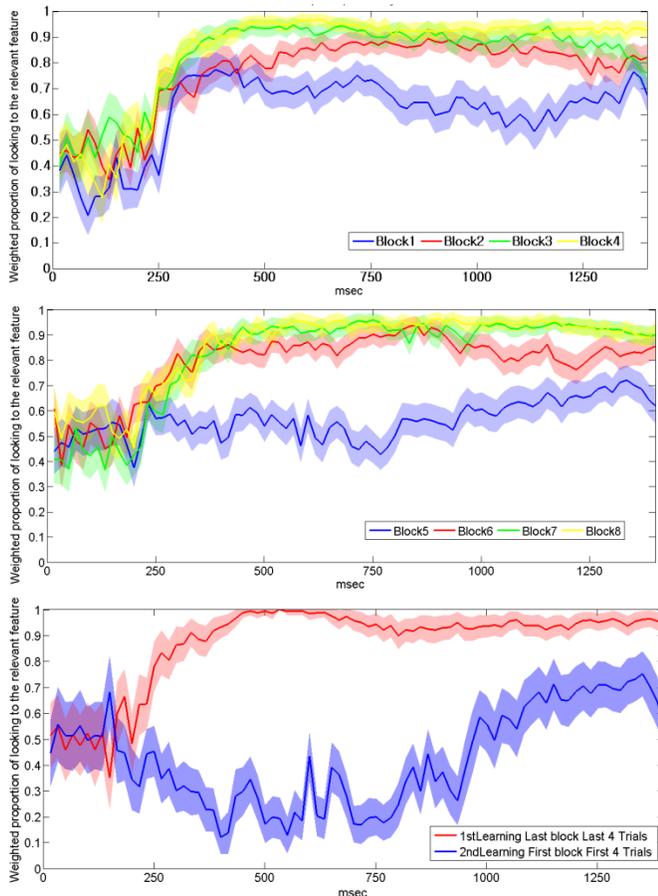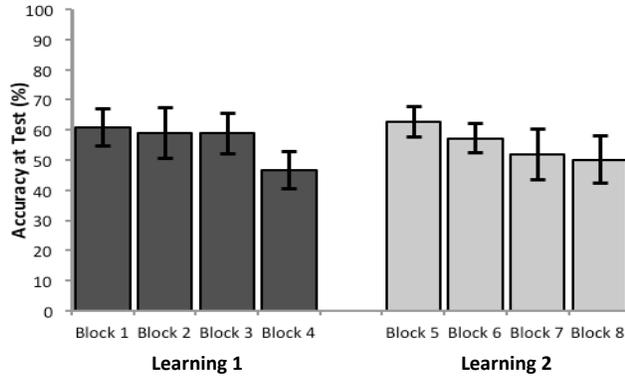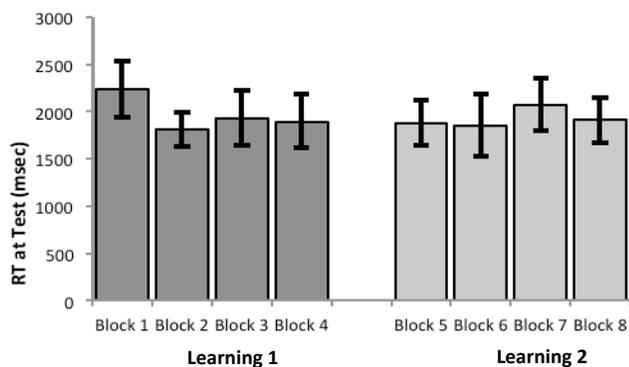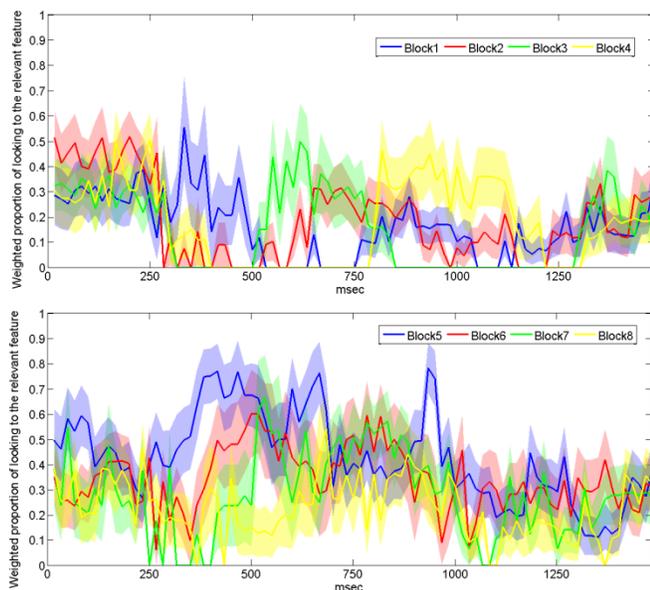

Figure 7. Eye gaze analysis for Experiment 2. Shaded area represents +/- one standard error. Weighted proportion of looking for learning phase 1 *(top)*. Weighted proportion of looking for learning phase 2 *(bottom)*. Please see online version for colored line graphs.

As was predicted, the non-learners in Experiment 2 did not show a cost of attention, as indicated by both behavioral and eye gaze patterns. Comparing these results to those of Experiment 1, it is notable that the only difference in the experiment procedure was the presence or absence of supervision. Therefore, the results showed that in sparse categories, learning could be manipulated by the amount of supervision, and that those participants who learned incurred a cost of attention.

Classifying Individual Learners

Taken together, both experiments indicated that a cost of attention could be one of the unique patterns of category learning. Specifically, the cost of attention during category learning was distinctively captured by eye gaze data. In this section a classifier using a simple neural network examined the predictability of the cost of attention for classifying individual learners from non-learners.

To classify learners from non-learners, a classical perceptron was used (Minsky & Papert, 1969). The input structure was constructed from individual data. Since the eye-tracker had a refresh rate of 60Hz, every individual had 90 sequential counts of whether they looked at the relevant or irrelevant features during every 1.5 sec of a trial. Moreover, since the cost of attention could be calculated by the difference in the eye gaze pattern between block 4 and block 5, individual eye gaze data from block 4 and block 5 were used for input. Therefore, the input structure consisted of 180 units where the first 90 units were from block 4 and the later 90 units were from block 5. The value of each unit was an average of 4 trials that consisted of the whole block. For each trial, relevant features were coded by 5 and irrelevant features were coded by -1. The weighted values for the relevant features equated the spatial coverage as was done in the eye gaze analysis. Thereafter, the trials were averaged by blocks, resulting in 180 input units for each individual. Output structures had one unit where the learners had a value of 1 and non-learners had a value of 0, with learners defined as subjects who were perfectly accurate on the last 4 test trial of the first learning phase.

Learning was conducted using a traditional delta-rule with a total of 49 individual data for training, 33 from Experiment 1 (supervised condition, 25 learners and 8 non-learners) and 16 from Experiment 2 (unsupervised condition, 2 learners and 14 non-learners). After 5 epochs, the network was able to learn the data set without errors, suggesting that the classification was a linearly separable data set (Minsky & Papert, 1969).

To simulate the predictability of the network, a Leave-One-Out Cross-Validation (LOOCV) method was used. A subset of the total individual data set (n=48) was used as the training set and one data set was left out for validation. All sub-sets of the training data were perfectly learned after a mean epoch of 5.02. The error rate was 12.2% with 6 out of 49 individual data sets and was significantly greater than chance, $p < .0001$. There were no systematic patterns among the 6 mis-predicted instances.

General Discussion

The current set of experiments manipulated supervision in the course of category learning. Manipulating supervision affected learning, and as a consequence of learning contrasting categories back to back, a cost of selective attention occurred.

The fact that supervision was required to learn sparse categories suggests that selective attention is necessary for learning a sparse category. Sparse categories have few relevant features and are therefore hard to learn. However, selective attention helps one to focus on relevant features while ignoring irrelevant features during learning. Additionally, relying on selective attention also results in learning to ignore irrelevant features (i.e., learned inattention), which can result in a cost of attention if a new category is introduced that requires shifting attention to previously irrelevant features.

In general, eye gaze data, especially the examination of participants' proportion of looking to the relevant features over time, provided a signature pattern of when learning occurred. Specifically, participants' looking to the category relevant feature during learning increased as their accuracy at test increased, indicating that participants were able to selectively attend to the category-relevant feature to learn the category. These results support previous research that adults will optimize their attention to category-relevant information for successful categorization (Hoffman & Rehder, 2010). Although eye-tracking confirmed participants' engagement in selective attention over time (i.e., greater proportion of looking to the relevant feature), the cost of attention immediately after the unknown category switch also confirmed their use of selective attention in the first learning phase.

The cost of attention as an indicator/predictor of learning was also examined using a neural network model. Using a section of eye gaze data that captured the cost of attention during the category switch, the network's prediction was quite accurate. The network's prediction is notable in that training was based on a relatively limited amount of samples. Moreover, the network's classification abilities were restricted to a simple linearly separable data set, which implies that the cost of attention is one of the strong and unique indicators for category learning. However, the results do not imply that learning is a consequence of a cost of attention. Instead, a cost of attention should be the result of learning, with the cost having strong links to the existence of learning.

Moreover, in future research it would be interesting to observe whether there are instances of category learning that are not accompanied by a cost of attention. Since a cost of attention is a consequence of selective attention, populations that have a relatively insufficiently developed prefrontal cortex may not rely on selective attention during category learning. It is known that infants and pigeons are capable of learning categories (Mareschal & Quinn, 2001; Soto & Wasserman, 2011), yet they arguably have immature prefrontal cortices, making for inept attentional control (e.g.,

selectively, inhibition). Therefore, it may be possible that populations of infants or animals with limited selective attention abilities would use alternative mechanisms to learn categories (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Sloutsky, 2010), and thus providing instances where a cost of attention does not indicate/predict learning.

References

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105*, 442-481.

Blair, M. R., Watson, M. R., & Meier, K. M. (2009). Errors, efficiency, and the interplay between attention and category learning. *Cognition, 112*(2), 330-336.

Hoffman, A. B., & Rehder, B. (2010). The Cost of Supervised Classification: The Effect of Learning Task on Conceptual Flexibility. *Journal of Experimental Psychology: General, 139*(2), 319-340.

Kloos, H., & Sloutsky, V. M. (2008). What's Behind Different Kinds of Kinds: Effects of Statistical Density on Learning and Representation of Categories. *Journal of Experimental Psychology: General, 137*(1), 52-72.

Kruschke, J. K. (1992). ALCOVE: An Exemplar-Based Connectionist Model of Category Learning. *Psychological Review, 99*(1), 22-44.

Kruschke, J. K., & Blair, M. (2000). Blocking and backward blocking involve learned inattention. *Psychological Bulletin & Review, 7*, 636-645.

Mareschal, D., & Quinn, P. C. (2001). Categorization in infancy. . *Trends in Cognitive Sciences, 5*, 443-450.

Minsky, M. L., & Papert, S. A. (1969). *Perceptrons.* Cambridge, MA: MIT Press.

Nosofsky, R. M. (1986). Attention, Similarity, and the Identification-Categorization Relationship. *Journal of Experimental Psychology: General, 115*(1), 39~57.

Sloutsky, V. M. (2010). From perceptual categories to concepts: What develops? *Cognitive Science, 34*, 1244-1286.

Soto, F. A., & Wasserman, E. A. (2011). Asymmetrical interactions in the perception of face identity and emotional expression are not unique to the primate visual system. *Journal of Vision, 11*(3), 1-18.