Unsupervised Learning in an Animal Model

Edward A. Wasserman (ed-wasserman@uiowa.edu)
Leyre Castro (leyre-castroruiz@uiowa.edu)
Department of Psychological and Brain Sciences, The University of Iowa
301 East Jefferson Street, Iowa City, IA 52242, USA

Vladimir M. Sloutsky (sloutsky.1@osu.edu)
Department of Psychology, The Ohio State University
1835 Neil Avenue, Columbus, OH 43210, USA

Keywords: unsupervised learning; supervised learning; associative learning; animal model; pigeons

Introduction

The learning of rich associative networks is foundational to language and other advanced cognitive competencies. Such learning is commonly encouraged by direct supervision entailing explicit feedback for correct and incorrect responses. The power of explicit feedback has recently been demonstrated in an animal model (Wasserman, Brooks, & McMurray, 2015), in which pigeons successfully learned to categorize 128 stimuli into 16 human language categories via both associative strengthening and weakening processes (Roembke, Wasserman, & McMurray, 2016).

The use of supervised training strongly implies "cajoling" an organism to respond correctly. Most supervised learning tasks for animals *differentially* reinforce behavior, so that reward is given when the animal's response is correct, but not when the animal's response is incorrect. Arranging unsupervised associative learning tasks in an animal model is decidedly more difficult, as doing so requires that experimenters *nondifferentially* reinforce behavior.

Nonetheless, unsupervised learning may materially participate in the acquisition of rich associative networks. Given that supervision can dramatically change how stimuli are learned and represented (Love, 2002), a key and as yet unmet challenge is to devise behavioral tasks that demonstrate associative learning in animals and that do not involve explicit supervision. Extending work in infant behavior (Sloutsky & Robinson, 2013), we devised and deployed a promising new paradigm to assess unsupervised learning in a pigeon model.

Experiment

Two groups of four pigeons each were shown eight object images and eight color patterns on a touch-sensitive computer screen. In the Consistent Pairings group, each object was paired with a particular pattern (e.g., Object A was always paired with Pattern 1, etc.), so that these birds could learn eight specific object-pattern pairs. In the Random Pairings group, each of the eight objects was presented an equal number of times with each of the eight color patterns, so that these birds could not learn any consistent object-pattern pairs. All birds were first trained without supervision and later trained with supervision, yielding two different ways to assess associative learning.

Unsupervised Phase

Daily sessions comprised 128 trials, in which each object image was shown 16 times: always followed by the same color pattern in the Consistent Pairings group or randomly followed by each of the color patterns in the Random Pairings group. Birds simply had to peck each of the images a fixed number of times (gradually increased to 10). After completing that requirement—first to the object and next to the color pattern—food was always given. There were no correct or incorrect responses; so, no differential feedback was ever provided. Under these unsupervised training conditions, were the Consistent Pairings birds learning the statistical relations between each object and color pattern?

To find out, after 5 sessions of training, we gave 1 testing session in which—in addition to the 128 training trials—we included 28 testing trials. On testing trials, after pecking the object, two color patterns (rather than one) appeared to the left and right of the object. In the Consistent Pairings group, one color pattern was the same one that had consistently followed the object in training sessions, whereas the other was one of the remaining color patterns, randomly chosen. In the Random Pairings group, the two color patterns were randomly chosen. We repeated this sequence of 5 sessions of training followed by 1 session of testing 8 times. Food always followed any choice response the bird made.

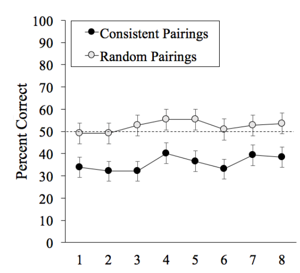

Figure 1: Mean percent correct for Consistent and Random Pairings groups on test trials during the unsupervised phase.

Results on testing trials are shown in Figure 1. The percent of correct responses reflects the tendency of birds in the Consistent Pairings group to choose the color pattern that had been paired with the object. There was no correct choice possible for the birds in the Random Pairings group, so their choice was bound to lie near the 50% chance level (*M* = 52%). Surprisingly, birds in the Consistent Pairings group actually preferred to peck the *other* color pattern, not the one that had been paired with the object; their choice performance was significantly below chance (*M* = 35%).

Thus, pigeons in the Consistent Pairings group did learn which specific color pattern was associated with each specific object; however, in choice tests, they displayed a preference for the other color pattern. This type of choice appears to parallel the preference for novelty in the classical children's preferential looking paradigm. Critically, all of the stimuli in our experiment were equally novel in both the Consistent Pairings group and in the Random Pairings group. Thus, the difference in performance here clearly implicates associative learning.

Supervised Phase

After eight cycles of unsupervised training and choice testing, we began the supervised training phase, which was the same for both groups of pigeons. Now, all trials presented one object followed by two color patterns. Choice of the correct color pattern was followed by food, whereas choice of the incorrect color pattern was not. Figure 2 shows the number of training sessions necessary to reach 65%, 75%, and 85% accuracy levels.

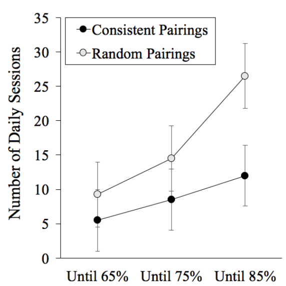

Figure 2: Mean number of days for Consistent and Random Pairings groups to reach 65%, 75%, and 85% accuracy levels in the supervised phase.

Mean accuracy on the first session of supervised training was 38% for the Consistent Pairings group and 51% for the Random pairings group, agreeing with their earlier behavior. Yet, despite this initial disadvantage, the Consistent Pairings group was actually faster to reach the 65%, 75%, and 85% correct levels than the Random Pairings group. Again, the Consistent Pairings group showed clear evidence of having learned the object-pattern pairs during the unsupervised training phase, here by learning faster than the Random Pairings group during the supervised learning phase.

Conclusions

These results clearly demonstrate *unsupervised* associative learning in pigeons using the same general stimuli and response options used in our earlier work in pigeons' *supervised* category learning. One must appreciate, however, that supervised learning tasks are not altogether free from the influence of unsupervised learning. When correct responses are made and positive feedback is given, a statistical regularity is enforced which can strengthen the stimulus-response bond on *correct-choice* trials. What most strikingly distinguishes supervised from unsupervised learning is therefore likely to be what happens on *incorrect-choice* trials; here, weakening or pruning of stimulus-response bonds can further direct responses to the correct choice option. Just such independent evidence was provided in the research of Roembke et al. (2016).

What remains to be determined is why—during the supervised phase—pigeons in the Consistent Pairings group were reluctant to peck the color pattern that had earlier been paired with the object image. Further work is underway to make that determination.

Acknowledgments

This research was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the NIH under Award Number P01HD080679.

Cited and Extended References

Couto, K. C., Navarro, V. M., Smith, T. R., & Wasserman, E. A. (2016). Concept learning without differential reinforcement in pigeons by means of contextual cueing. *Journal of Experimental Psychology: Animal Learning and Cognition, 42,* 221-227.

Love, B. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin & Review, 9,* 829-835.

Roembke, T. C., Wasserman, E. A., & McMurray, B. (2016). Learning in rich networks involves both positive and negative associations. *Journal of Experimental Psychology: General, 145,* 1062-1074.

Sloutsky, V. M. & Robinson, C. W. (2013). Redundancy matters: Flexible learning of multiple contingencies in infants. *Cognition, 126,* 156-164.

Wasserman, E. A. (2016). Conceptualization in pigeons: The evolution of a paradigm. *Behavioural Processes, 123,* 4-14.

Wasserman, E. A., Brooks, D. I., & McMurray, B. (2015). Pigeons acquire multiple categories in parallel via associative learning: A parallel to human word learning? *Cognition, 136,* 99-122.

Wasserman, E. A., Castro, L., & Freeman, J. H. (2012). Same-different categorization in rats. *Learning & Memory, 19,* 142-145.

Wasserman, E. A., & Young, M. E. (2010). Same-different discrimination: The keel and backbone of thought and reasoning. *Journal of Experimental Psychology: Animal Behavior Processes, 36,* 3-22.

Zentall, T. R., & Wasserman, E. A. (2012). *Oxford handbook of comparative cognition.* New York: Oxford University Press.