Cognition 173 (2018) 28–33

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Brief article

Unsupervised learning of complex associations in an animal model

Leyre Castro*, Edward A. Wasserman, Marisol Lauffer

The University of Iowa, United States

ARTICLE INFO

Keywords:
Unsupervised learning
Supervised learning
Statistical learning
Novelty preference

ABSTRACT

Supervised learning results from explicit corrective feedback, whereas unsupervised learning results from statistical co-occurrence. In an initial training phase, we gave pigeons an unsupervised learning task to see if mere pairing could establish associations between multiple pairs of visual images. To assess learning, we administered occasional testing trials in which pigeons were shown an object and had to choose between previously paired and unpaired tokens. Learning was evidenced by preferential choice of the previously *unpaired* token. In a subsequent supervised training phase, learning was facilitated if the object and token had previously been *paired*. These results document unsupervised learning in pigeons and resemble statistical learning in infants, suggesting an important parallel between human and animal cognition.

1. Introduction

Most studies of learning involve supervision; the learner is given a task and is appropriately provided with explicit contingent feedback about the correct or incorrect response. In this way, the learner is *directed* to make that specific correct response. In unsupervised learning, however, the learner must extract the underlying structure of the information presented, without being provided with any corrective feedback about that information.

This distinction between supervised and unsupervised learning is encountered in many realms of cognitive science: from machine learning (Huang, Huang, Song, & You, 2015), to speech perception (Goudbeek, Cutler, & Smits, 2008), to visual categorization (Colreavy & Lewandowsky, 2008; Love, 2002). Although it is entirely reasonable to equate supervised learning with operant conditioning (where the prevailing contingencies of reinforcement depend on an organism's behavior; see Staddon & Cerutti, 2003) and unsupervised learning with Pavlovian conditioning (where the contingencies of reinforcement prevail independently of an organism's behavior; see Domjan, 2005), these equations are not commonly made. Perhaps ironically, unsupervised learning in humans is believed to participate in complex linguistic and cognitive behaviors (Aslin & Newport, 2012), whereas Pavlovian conditioning in animals is typically relegated to behaviors of stark simplicity (Wasserman & Miller, 1997).

In a recent supervised learning study, pigeons were taught to sort 128 photographs of objects into 16 categories by associating each photograph with 1 of 16 visual tokens ('pexigrams') that could be pecked to report category membership (Wasserman, Brooks, & McMurray, 2015). On each trial, pigeons saw a single object plus two pexigrams: the correct pexigram for the presented category and a randomly selected pexigram from one of the other 15 categories. Pigeons received food for pecking the correct pexigram, but they received no food and were given correction trials for pecking the incorrect pexigram; thus, the experimenters only reinforced the pigeons' correct responses. Pigeons accordingly mastered the category learning task: They learned virtually all 16 training categories, they evidenced reliable generalization to new category exemplars, and they exhibited a high degree of coherence in their responding to stimuli within specific categories.

In Wasserman et al.'s (2015) task, category exemplars and correct pexigrams co-occurred; yet, as just noted, the experimenters additionally administered explicit corrective feedback. So, it is not possible with that paradigm to separately assess the contributions of supervision and statistical co-occurrence to pigeons' category learning. Because statistical co-occurrence is believed to be the sole route to association formation in unsupervised learning, we sought to develop and validate an unsupervised animal learning paradigm that is capable of supporting a rich set of associations acquired without explicit corrective feedback, much as infants and youngsters learn statistical regularities in their environment (e.g., Saffran, 2003).

In the current project, we explored if simply pairing each of eight objects with each of eight pexigrams is sufficient to establish associations between them. Because the mere pairing of stimuli under unsupervised training would not yield informative performance, we included occasional probe trials in which each object was presented with both paired and unpaired pexigrams, and pigeons had to select one of

* Corresponding author at: Department of Psychological and Brain Sciences, 301 East Jefferson Street, The University of Iowa, Iowa City, IA 52242, United States.
E-mail address: leyre-castroruiz@uiowa.edu (L. Castro).

https://doi.org/10.1016/j.cognition.2017.12.015
Received 19 October 2017; Received in revised form 20 December 2017; Accepted 21 December 2017
0010-0277/

them. This choice test allowed us to determine whether learning on unsupervised trials had taken place. Specifically, we could see if pigeons preferred to choose: (a) the pexigram that had previously been paired with that object or (b) a pexigram that had never been paired with that object, but that was equally familiar because it had been paired with one of the other objects. Option (a) was arguably the most straightforward outcome, as it would represent elementary excitatory conditioning; however, Option (b) was also a plausible outcome as it would parallel the frequent finding of infants' fondness for novelty after habituation training (Oakes, 2010). Either option would attest to associative learning having taken place during unsupervised training. If so, then we could discount the requirement of explicit corrective feedback for learning to occur, because no errors were possible under our training procedure. We would then have succeeded in unequivocally demonstrating unsupervised learning of multiple associations in animals.

A second training phase followed during which supervised training was arranged. Now, only choice of the pexigram that had previously been paired with a specific object on unsupervised training trials was followed by food. The aim here was to see if prior unsupervised training affected subsequent supervised training by hastening, or perhaps delaying, acquisition under supervised conditions.

2. Method

2.1. Subjects

The subjects were 16 male and female homing pigeons (*Columba livia*) maintained at 85% of their free-feeding weights. These pigeons had served in unrelated studies prior to the present project. The pigeons were randomly placed into four groups: Paired Differential Effort (DE), Paired Nondifferential Effort (NDE), Random Pairing, and Supervised. All procedures were approved by the Institutional Animal Care and Use Committee at the University of Iowa.

2.2. Apparatus

We used 16 computer-controlled conditioning chambers detailed by Gibson, Wasserman, Frei, and Miller (2004); each was equipped with an LCD monitor located behind a resistive touchscreen. Pecks to the touchscreen were processed by a serial controller board. A rotary dispenser delivered 45-mg pigeon pellets into a food cup on the wall opposite the touchscreen. Programs were developed in MATLAB® with Psychtoolbox-3 extensions (Brainard, 1997; http://psychtoolbox.org/).

2.3. Stimuli and procedure

Eight colored objects were presented within 6.3 × 6.3 cm squares and were paired with multicolored 5 × 5 cm squares—pexigrams (Fig. 1).

2.3.1. Acclimation

All 16 pigeons were exposed to all of the images (8 objects and 8 pexigrams) for one session. Each image was presented 10 times, for a total of 160 trials. At the beginning of each trial, pigeons were shown an orienting stimulus: a white square (3 × 3 cm) in the middle of the computer screen. After one peck to this stimulus, one of the images appeared on the screen and pigeons had to peck at it twice, after which food was delivered. This phase allowed the birds to acclimate to the new stimuli.

2.3.2. Unsupervised phase

Daily training sessions comprised 128 trials. Paired DE, Paired NDE, and Random Pairing groups were trained in this phase, but not the Supervised group. A trial started with a white square in the middle of the screen. Following one peck at this orienting stimulus, one of the objects was displayed. The pigeons had to satisfy an observing response requirement to the object (gradually increased from 3 to 10 pecks, on a daily basis). On completing this requirement, one of the pexigrams randomly appeared 2 cm to the left or right of the object. In groups Paired DE and Paired NDE, the same object was consistently paired with the same pexigram (half of the time in the left or right location), so that an object-pexigram association could be learned; however, in the Random Pairing group, each of the objects was equally often followed by each of the pexigrams, so that no associations could be learned. The pigeons had to peck at the pexigrams the same number of times as the objects. Once this requirement was fulfilled, food was delivered and pigeons proceeded to the next trial. After five sessions of unsupervised training, one testing session was given. This training/testing sequence was repeated eight times to see if and how responding changed with additional experience.

2.3.3. Testing

Testing sessions comprised 156 trials. Following 16 training trials, 28 testing trials were randomly interspersed among 112 additional training trials. Training trials were the same as in unsupervised training sessions. On testing trials, once the object was presented in the center of the screen and the pecking requirement was completed, *two* pexigrams appeared, to the left and right of the object, and the pigeons had to peck *one* of them. In groups Paired DE and Paired NDE, one option was always the pexigram that had been paired with that object during the unsupervised phase, whereas the other option was one of the seven remaining pexigrams, each presented equally often as possible choices and equally often in left and right locations.

On unsupervised *training* trials, the pigeons in all three groups were required to peck at the different objects and pexigrams the same number of times; however, that was not the case on unsupervised *testing* trials. On testing trials, after pecking at the object the required number of times, a choice response could be made by just one peck (in the Paired DE group) or by several pecks (the same number of pecks as were made at the object in the Paired NDE group). So, in the Paired DE group, there was a difference between the effort (number of pecks) expended on unsupervised training trials and on testing trials; but, there was no such difference in the Paired NDE group. We deemed it necessary to include both types of testing trials in these two groups because of possible long-term effects of this response requirement disparity. We did not expect this differential effort to have any effect on pigeons' test responding in the Random Pairing group because no associations could have been learned in this group; so, only one peck was required on testing trials.

Regardless of the number of pecks, any choice response was reinforced (*nondifferential reinforcement*), in all three groups, so that no particular associations could be learned on testing trials. After the eighth testing session, supervised training began.

2.3.4. Supervised phase

All three earlier groups, plus the Supervised group, were trained in this phase. The Supervised group was given one acclimation session, as the other three groups, and then moved directly to the supervised phase, which was identical for all four groups. Now, each of the objects was consistently paired with one of the pexigrams; these were the same pairings given to the Paired DE and NDE groups in the unsupervised phase. On every trial, after one peck to the orienting stimulus, one of the objects was displayed in the center of the screen and pigeons had to complete the observing response requirement. Then, two pexigrams appeared to the left and right of the object, randomly located from trial to trial; one was always the assigned pexigram, whereas the other was one of the seven other pexigrams, randomly chosen on each trial. Pigeons had to peck the assigned pexigram. All of the groups were required to make just one choice peck. In order to encourage learning of the correct response, we used *differential reinforcement*: If the choice response was correct, then food was delivered; if the choice response