Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2009 May ; 35(3): 731–741. doi:10.1037/a0015005.

Learning mode and exemplar sequencing in unsupervised category learning

Dagmar Zeithamova¹ and W. Todd Maddox

Department of Psychology Institute for Neuroscience University of Texas at Austin

Abstract

Exemplar sequencing effects in incidental and intentional unsupervised category learning were investigated to illuminate how people form categories without an external teacher. Stimuli were perfectly separable into two categories based on one of two dimensions of variation. Sequencing of the first twenty training stimuli was manipulated. In the blocked condition, 10 category A stimuli were followed by 10 category B stimuli. In the intermixed condition, these 20 stimuli were ordered randomly. Experiment 1 revealed an interaction between learning mode and sequence, with better intentional learning for intermixed sequences, but better incidental learning for blocked sequences. Experiment 2 showed that manipulating trial-to-trial variability along each dimension can impact intentional learning. Training sequences that emphasized variation along the category-relevant dimension resulted in better performance than sequences that emphasized variation along the category-irrelevant dimension. The results suggest that unsupervised category learning is influenced by the mode of learning and the order and nature of encountered exemplars.

Keywords

unsupervised learning; classification; sequence effects

Introduction

Categorization is an integral part of cognition. Most category learning research focuses on supervised categorization, i.e. categorization that involves external feedback regarding the correct category label. However, unsupervised category learning – the ability to discover or invent categories in the environment without category labels or feedback from an external tutor – may be equally prevalent in everyday life (e.g. Billman & Heit, 1988).

Unlike supervised categorization that is always intentional, unsupervised categorization can be intentional or incidental. For instance, a person who wants to organize her email inbox may decide to create various folders (categories) and come up with a criterion (categorization rule) by which she sorts the incoming messages into these folders. In this case, categorization is unsupervised, but intentional – the person decides to create categories that make the most sense to her and best reflect the structure of her incoming emails. On the other hand, a person may also notice or invent categories incidentally, without a priori intention to do so. A child playing outdoors may notice the large variability among trees and perhaps start to distinguish trees into

Correspondence should be addressed to Dasa Zeithamova (E-mail: zeithamova@mail.utexas.edu) or Todd Maddox (E-mail: maddox@psy.utexas.edu).

¹This research was supported in part by NIH grants MH59196 and MH077708 to WTM and a University continuing fellowship to DZ. We thank Brad Love for many helpful discussions and Stephan Lewandowsky, Dorrit Billman and 3 anonymous reviewers for helpful comments on earlier versions of the manuscript.

those with needles and those with leaves. In these cases, categorization is unsupervised and incidental, the persons created separate categories because they happen to notice some structure in the environment even when they were not explicitly looking for any. In this article, we investigate the under-studied topic of unsupervised category learning, with a primary focus on the mechanism of intentional unsupervised learning.

Unsupervised categorization has been studied using two complementary paradigms. In one tradition, participants are instructed to sort a set of stimuli into categories of their own choosing, with no right or wrong answer; the research question is what category sorts are preferred by the participants under various circumstances (Colreavy & Lewandowsky, 2008; Medin, Wattenmaker, & Hampson, 1987; Pothos & Chater, 2002; Regehr & Brooks, 1995). In a second tradition, an underlying category structure exists and the research question is whether participants can discover or match the underlying category structure in unsupervised learning (Ashby, Queller, & Berretty, 1999; Love, 2003). The research presented here is in line with the second tradition and focuses on how sequencing of exemplars can promote or hinder participants' ability to match their categorization choices to the structure of the environment.

Exemplar Sequencing Effects in Unsupervised Category Learning

A useful tool for understanding category learning and constraining its models are studies on the effect of exemplar sequencing. Sequencing (order) effects have been shown to affect both supervised and (incidental) unsupervised learning. While existing studies of supervised category learning addressed primarily the question of how contrast of successive stimuli affects the category assignment for the current stimulus with respect to its usual category assignment (Jones, Love, & Maddox, 2006; Jones & Sieck, 2003; Stewart, Brown, & Chater, 2002), studies of unsupervised learning addressed the question of whether the underlying category structure itself can be learned under particular exemplar sequences (Clapper, 2006; Clapper & Bower, 1994, 2002). As we also aim to address the question of category learnability under unsupervised learning, we first turn to the studies of unsupervised learning

An important set of studies that demonstrated large sequence effects in unsupervised learning was conducted by Clapper and Bower (Clapper, 2006; Clapper & Bower, 1994, 2002). They presented participants with exemplars from two feature-correlation based conceptual categories. The stimuli varied along multiple binary dimensions, were highly correlated within a category, and formed two perfectly separable categories. Participants did not notice the underlying category structure when presented with category exemplars intermixed from the two categories. However, when presented with a block of exemplars from one category followed by exemplars from the other category, the category structure became obvious.

The results of these experiments supported a theory that postulated a surprise-driven category invention mechanism in unsupervised learning. The main idea is that a sequence of stimuli that belong to one category leads a participant to build a representation of the stimulus space that matches that of the category. When the first stimulus from a second category is presented, its values violate participants' expectations and this "surprise" leads to the construction of a second category.

Clapper and Bower focused on incidental, spontaneous unsupervised category learning. However, the notion of surprise-driven learning is less meaningful when we consider intentional unsupervised learning. During incidental categorization, a participant has no reason to form multiple categories unless a large expectation violation - the "surprise" - alerts them to the possibility of separate categories. On the other hand, it is likely that a participant during intentional categorization actively searches for any meaningful distinction in the stimuli and thus expects some stimuli to be different from others. Interestingly, few studies have examined intentional unsupervised category learning, and none examined exemplar sequence effects. It

is reasonable to expect that the order in which objects are encountered influences intentional unsupervised learning. However, given that the mechanisms of incidental and intentional unsupervised learning may differ, it is questionable whether the sequence effects observed in incidental learning are generalizable to intentional learning.

Overview of Current Studies

The current studies examine the effects of exemplar sequencing on intentional unsupervised learning to determine how it may differ from incidental unsupervised learning. Experiments 1A and 1B examine the effects of blocked versus intermixed exemplar sequencing on intentional and incidental unsupervised category learning. To anticipate, we found that participants are more likely to categorize in agreement with the underlying structure when presented with blocked sequence during incidental learning, but interestingly, the opposite was true for intentional learning.

Experiment 2 extends previous work on stimulus contrast effects in supervised learning to unsupervised intentional learning situations. Stimulus contrast effects demonstrate that a participant's tendency to assign the current stimulus to one of two categories depends upon the category assignment of the previous stimulus, and the perceived dissimilarity between the current and previous stimulus (Jones, Love, & Maddox, 2006; Jones & Sieck, 2003; Stewart, Brown, & Chater, 2002). We proposed that exemplar sequences with different stimulus contrast characteristics will affect intentional unsupervised learning.

Experiment 1A

Experiment 1A tested directly whether a blocked exemplar presentation improves unsupervised learning under intentional learning conditions, as was previously found under incidental learning conditions.

Method

Participants—Seventy-one University of Texas at Austin students participated in the experiment in partial fulfillment of a course credit requirement, or for pay. All participants were tested for 20/20, or corrected to 20/20, vision.

Stimuli and apparatus—The stimuli were white lines varying in length and orientation, presented on a black background centered on a computer screen. These physical dimensions were chosen because their perceptual properties are well understood. For example, Nosofsky (1985, 1986) utilized multidimensional scaling (MDS) techniques and showed that the psychological scaling solutions mirrored the physical spacing for length and orientation. In addition, Nosofsky showed that the two dimensions are independent of one another and are generally considered to be separable [see also Shepard (1964)]. This ensures that the stimulus spacing in the perceptual space closely matches the physical spacing and that the value changes along one dimension do not alter the perception of the other dimension.

We followed the tradition of general recognition theory (Ashby, 1992; Ashby & Townsend, 1986) and represented categories as normally distributed clusters varying along two perceptual dimensions (see also Fried & Holyoak, 1984). The two categories were perfectly separable along one perceptual dimension that was bimodal, while the distribution of stimulus values along the second, category-irrelevant dimension was unimodal (see Figure 1). Stimulus value on the bimodal dimension was thus relevant for category membership while the value on the unimodal dimension was irrelevant for category membership. Which physical dimension (orientation or length of a line stimulus) was the bimodal, category-relevant dimension was

counterbalanced within each condition, with length serving as the bimodal dimension for half of the participants while orientation for the other half of the participants.

The selected category structure was convenient for two reasons. As people often prefer unidimensional sorts during unsupervised categorization (Ahn & Medin, 1992; Ashby, Queller, & Berretty, 1999; Medin, Wattenmaker, & Hampson, 1987; Regehr & Brooks, 1995; see however Colreavy & Lewandowsky, 2008), we aimed to use category structure that offered two possible unidimensional strategies, one matching the underlying category structure and the other being orthogonal to it. To evaluate learning, we determined whether each participant chose to separate stimuli along the bimodal dimension (matching the underlying category structure), the unimodal dimension, or using some other strategy. Second, previous research typically used a small number of stimuli varying along a few binary-valued dimensions, allowing for only a few testing trials. For example Clapper and Bower (1994) used 8, 24, and 13 testing trials in their Experiment 1, 2, and 3 respectively. However, it is possible that learning may continue with further exposure to the stimuli (Ashby, Queller, & Berretty, 1999). The category structure used here offers a large number of unique stimuli and allows for probing learning both early and after more extensive exposure.

The stimuli were generated and presented on a computer screen using Psychophysics toolbox (Brainard, 1997; Pelli, 1997) in Matlab (Mathworks, Natick, MA). The category distribution parameters in actual units are listed in Table 1. Line orientation was constrained to fall between 0 and 90 degrees to avoid crossing the highly salient cardinal orientations of 0 and 90 degrees (Zeithamova & Maddox, 2007).

Procedure—Participants were told that they would be presented with a series of lines that varied across trials in length and orientation and that each line was a member of one of two naturally distinct categories. Participants' task was to discover this natural distinction. The existence of the underlying category structure was emphasized by a cash bonus offered to participants who reach at least 75% accuracy in matching the underlying categories. Participants were told to press one button on the keyboard every time they saw a member of one category and another button every time they saw a member of the other category. Participants were told that it did not matter which button they associated with which category, but after they had figured out the nature of the two categories, they should remain consistent in which label they used for which category.

On each trial, a line was presented and the participant was prompted to categorize it either as an A or B stimulus by pressing the appropriate keyboard button. The stimulus remained on the computer screen until the participant categorized it, followed by a 1000 ms blank screen inter-trial interval. No feedback was given throughout the experiment.

The experiment consisted of a 20-trial training phase and a 160-trial testing phase. The selection and sequencing of the training trials was manipulated across conditions while the testing phase was identical across conditions and consisted of 80 A stimuli and 80 B stimuli randomly intermixed. Participants' categorization strategy was evaluated both early (i.e. during first 20 testing phase trials, testing length comparable to several previous studies) as well as across the entire testing phase (i.e. the total of 160 testing phase trials, to evaluate the effect of more extensive exposure).

Three exemplar sequencing conditions were examined. In the intermixed condition, the training phase consisted of 10 A and 10 B stimuli that were randomly selected and presented in an intermixed order. Following Clapper & Bower (2002), the presentation sequence of the stimuli was controlled in the blocked condition, so that a long streak of stimuli from only one category was presented before presenting the first stimulus from the other category. Clapper and Bower

found a similar advantage for the blocked condition both when they presented a block of A stimuli followed by a block of B stimuli as when they presented a block of A stimuli only. As it was not obvious whether to expect the same finding here, we decided to further subdivide the blocked condition into a blocked (AB) group and a blocked (A only) group. In the blocked (AB) group, the training sequence consisted of 10 A stimuli followed by 10 B stimuli. In the blocked (A only) group, the training sequence consisted of 20 A stimuli only. Stimuli within each category were selected and ordered randomly. Of the 39 participants in the blocked condition, 20 participants were assigned to the blocked (AB) group and 19 were assigned to the blocked (A only) group. To anticipate, both blocked groups – the blocked (AB) group and the blocked (A only) group – performed about equally on all measures.

Results and Discussion

Categorization Strategy Analysis Method—We hypothesized that a participant may use one of three categorization strategies that we call bimodal, unimodal and flat. The bimodal categorization strategy assumes that a participant consistently uses a rule along the bimodal, category-relevant stimulus dimension – e.g. sorting short lines into one category and long lines into another category when length is indeed the bimodal dimension. The unimodal categorization strategy assumes that a participant consistently uses a rule along the category-irrelevant, unimodal stimulus dimension. The flat categorization strategy assumes that a participant uses a category label with a fixed probability, irrespective of which stimulus is presented. Participants who were best fit by the flat model could be responded randomly or could be adopting a number of different strategies across the experiment.²

We used a decision bound model to instantiate each strategy. A detailed formalization of the decision bound model is provided elsewhere (see Ashby, 1992), but here we briefly describe the basic procedure. Each unidimensional (bimodal and unimodal) model assumes that an observer sets a criterion on a single perceptual dimension and then makes a decision about the stimulus value on that dimension. When the stimulus value falls above the criterion, it is classified as A, and when the stimulus value falls below the criterion, it is classified as B. For example, in the present experiment, observers might use the rule: respond A if the line length is short, respond B if it is long. The unidimensional (bimodal and unimodal) models have two free parameters: a decision criterion on the selected perceptual dimension and the variance of internal (perceptual and criterial) noise. The bimodal model assumes that the criterion was created along the bimodal dimension (e.g. Figure 2, left panel); the unimodal model assumes that the criterion was created along the unimodal dimension (e.g. Figure 2, middle panel). Unlike the two unidimensional models, the flat model assumes that the participant responds with a fixed probability of responding A to all stimuli, irrespective of their stimulus values along either dimension. Generally, the flat model has one free parameter, p (the probability of responding A). We also fit a zero-parameter version of the flat model, assuming a fixed probability of responding A at $p = .5$.

²It is possible that some participants consistently used one strategy, but occasionally switched response buttons because they forget the category assignment that they were using. Such a participant would be classified as a flat responder. It is also possible that participants base their categorization decision on both dimensions. To capture either such behavior, we developed an alternative dependent measure that determined whether changes in the stimulus dimensional values from trial to trial were predictive of a switch from one category response to the other. This analysis provided us with model-free weighting of each dimension on individual participant's responses. For instance, a participant would score high weight on length if he responded with the same button on consecutive trials when the current and previous stimulus had similar lengths, but responded with different buttons on consecutive trials when the current and previous stimulus had dissimilar lengths. In short, the results of this alternative analysis closely match those of the response strategy analysis, with vast majority of participants identified as using bimodal or unimodal strategy indeed showing high weight for that dimension and near-zero weight for the other. These results suggest that the assumption of a unidimensional strategy is a valid one. Additionally, the vast majority of participants classified as flat showed near-zero weight on both dimension, suggesting that these participants were mostly responding randomly and did not appear to be simply forgetting the button assignments.

Each of the models was fit to the test phase categorization responses of each participant, once considering only the first 20 testing phase trials (for comparison with existing literature; “early”) and once to the all 160 testing trials (“across”). Model parameters were estimated using maximum likelihood procedures and the Bayesian Information Criterion (BIC) (Schwarz, 1978). The BIC value (Equation 1) is computed from the likelihood of the responses (L), but takes into account the number of free parameters (k) and the number of observations (N) that were used to compute the likelihood (penalizing for extra free parameters).

$$\text{BIC} = -2 \ln L + k * \ln N \quad (1)$$

Model fit evaluation—By selecting the model associated with the smallest BIC value, we identified the categorization strategy (bimodal, unimodal or flat) that best characterized each participant's responses. The majority of participants' data was best fit by the bimodal or the unimodal model (see Figure 3 below). Visual inspection of each participant's response pattern closely matched those of the formal model fits. A representative response pattern for three participants, each using a different categorization strategy, is presented in Figure 2.

The mean goodness-of-fit value (BIC) for all models and the proportion of responses accounted for by the models, separately for participants that were best fit by the bimodal, unimodal and flat strategies, are presented in Table 2. The proportion of responses accounted for by a model can be computed by binarizing the likelihood of a response to each stimulus, i.e. predicting response A when likelihood of A according to the model is higher than likelihood of B (and vice versa), and then computing the proportion of participant's responses that match the predicted response. Table 2 shows that for participants identified as using a bimodal or a unimodal strategy, the corresponding proportion of correctly predicted responses were very high, around 90%, confirming that these models did an excellent job of characterizing participants' responses.³

Categorization strategy results—The proportions of participants within each group using the bimodal, the unimodal and the flat strategy early (first 20 testing phase trials; left panel) and across the full 160 trial testing phase (right panel) are depicted in Figure 3. To assess strategy preferences in each condition, we tested whether (among the participants exhibiting a consistent -- bimodal or unimodal -- strategy) one of the strategies was selected over the other strategy significantly more frequently than expected by chance, as evaluated using a binomial test with probability of $p = .5$ for each of the two strategies. The results are presented in Table 3. The results (Figure 3 and Table 3) show that blocked sequences did not lead to better learning in terms of a match between participants' categorization strategies and the underlying category structure. On the contrary, the bimodal categorization strategy was significantly preferred over the unimodal strategy in the intermixed condition, suggesting that participants were sensitive to the underlying category structure when presentation order was not manipulated. In contrast, there was no significant preference for either strategy in either the blocked (AB) condition or the blocked (A only) condition. In fact, the unimodal categorization strategy was selected numerically more frequently than the bimodal strategy in both blocked conditions.

³For brevity, we report the goodness-of-fit and proportion of responses accounted for by the models only for Experiment 1A. It is important to note that similar excellent fits were obtained in all following experiments reported in this paper, suggesting that the unidimension (bimodal and unimodal) models provided good account for participants' responses in all conditions that were investigated.

Experiment 1B

The advantage of blocked over intermixed exemplar sequences in unsupervised incidental learning has been shown repeatedly (Clapper, 2006; Clapper & Bower, 1994, 2002). This makes the reversed pattern observed in Experiment 1A for intentional learning quite interesting. To ensure that the reversal of the effect observed here was not due to stimulus, category structure or other methodological differences, but indeed due to a different learning mode, we replicated Experiment 1A but added learning mode (intentional vs. incidental) as a second factor in Experiment 1B. As we observed no significant differences between the blocked (AB) group and blocked (A only) group, we include only the blocked (AB) version as it equates the number of training stimuli from both categories. For consistency across the testing phase of Experiments 1A and 1B, we manipulated the participants' learning mode only during the presentation of the training trials, keeping the testing phase identical to Experiment 1A in all conditions.

Method

Participants—Eighty University of Texas at Austin students participated in the experiment as partial fulfillment of a course credit requirement or for pay. All participants had 20/20 or corrected to 20/20 vision. Participants were randomly divided into four groups of 20 participants each: Intentional intermixed, intentional blocked, incidental intermixed, incidental blocked.

Stimuli and apparatus—The stimuli and presentation apparatus were identical to those of Experiment 1A.

Procedure—The procedure was similar to Experiment 1A except that 240 testing trials were included instead of 160. As in Experiment 1A, we analyzed categorization strategy both early (during the first 20 testing trials) and across testing phase. For the across testing phase analysis, we focused on both 160 test trials (as in Experiment 1A) and on all 240 trials. For brevity and consistency across studies, we report only the results based on the first 160 trials. Including all 240 trials yielded minimal differences (maximum of one participant changing strategy in any condition).

The exemplar sequence and learning mode were manipulated during the 20-trial training phase. A 2 (sequencing order: intermixed vs. blocked) $\times$ 2 (learning mode: intentional vs. incidental) factorial design was used. The selection and sequencing of the initial 20 trials in the intermixed and the blocked sequencing order condition was the same as for the intermixed group and the blocked (AB) group from Experiment 1A. The instructions for participants in the intentional learning mode condition were identical to those from Experiment 1A. Participants in the incidental learning mode condition were kept naïve before and during the training phase regarding the nature of the experimental task. To facilitate possible incidental category learning during the training phase, we adopted a pleasantness rating task and cover story used successfully in previous unsupervised category learning research (Love, 2003). Participants were asked to view the training stimuli and rate how much they liked each line on the scale from 1 to 7. Participants were told that the ratings would be used to select stimuli for a future experiment. After viewing and rating the training stimuli (sequenced either randomly or blocked by category, just as in the intentional condition), the participants were told that the stimuli they just rated, as well as the stimuli they were about to see, were members of two naturally distinct categories. They might have noticed the distinction already. If not, they would discover the method of separating the stimuli into the two naturally distinct categories during the rest of the experiment. The testing phase was identical for all participants and involved a categorization judgment without feedback on each trial.

Results and Discussion

The proportion of participants using the bimodal, unimodal, and flat response strategies within each group is depicted in Figure 4. The categorization strategy preferences in each condition across participants are presented in Table 4. In the intentional learning mode conditions, the bimodal categorization strategy was significantly preferred over the unimodal strategy in the intermixed condition, but not the blocked condition, replicating the results of Experiment 1A. The opposite was found for the incidental learning mode: while participants in the intermixed condition did not show any strategy preference, the participants in the blocked condition were significantly more likely to use the bimodal strategy. This suggests that participants in the incidental intermixed condition did not endorse the underlying category structure in the training stimuli during the pleasantness rating task and were equally likely to pick the bimodal or the unimodal dimension when surprised with the categorization task. On the other hand, when the presentation of the training exemplars was blocked, the incidental participants were more likely to match their preference to the underlying category structure. The interaction of factors (learning mode $\times$ sequencing) on the probability of using the bimodal strategy was significant for early testing phase (early: $B = 0.934$, $S.E. = 0.464$, $p = .044$; across: $B = 0.449$, $S.E. = 0.470$, $p = .340$; logistic regression for probability of using bimodal strategy).

The results from Experiment 1 suggest that learning mode modulates the effect of exemplar sequencing in unsupervised categorization. In incidental learning, blocking training exemplars by category increased the likelihood of participants using strategy matching the underlying category structure, consistent with previous findings (Clapper & Bower, 1994, 2002). However, this pattern reverses in intentional learning.

Experiment 2

Experiment 2 further explored sequencing effects in intentional unsupervised learning. Research on recency effects demonstrated how similarity between successive stimuli affects categorization responses. For example, people tend to assign a new stimulus into a different category than the previous stimulus when the new stimulus differs saliently from the previous stimulus (Jones, Love, & Maddox, 2006; Jones & Sieck, 2003; Stewart, Brown, & Chater, 2002). We hypothesized that in unsupervised learning, participants may be more likely to separate stimuli along the dimension that changes its value more dramatically. In the blocked condition, stimulus-to-stimulus changes were large along the unimodal dimension, but minimal along the bimodal dimension, working against the effect of the underlying category structure. In the intermixed condition, stimulus-to-stimulus changes could be large along both dimensions, thus leading the participant to rely mainly on the underlying category structure. Thus, we hypothesized that stimulus-to-stimulus changes along the two dimensions are the critical factor in categorization strategy selection when the learning mode is intentional.

To test this hypothesis, we constructed three training sequences that accentuated or attenuated the stimulus-to-stimulus contrast on the two dimensions. First, we constructed a training sequence with maximal stimulus-to-stimulus changes along the bimodal dimension and minimal changes along the unimodal dimension. A schematic of this sequence is presented in Figure 5, left-most panel (“zipper” condition). If stimulus contrast affects strategy selection in intentional learning as we hypothesized, this manipulation should lead to an increase in the bimodal strategy preference. Second, we constructed a training sequence that was a variation of the blocked sequence, but avoided large stimulus changes along the unimodal dimension (“blocked-gradual” condition; Figure 5, middle panel). We expected higher bimodal strategy use in this blocked-gradual condition than the blocked (and perhaps even intermixed) condition from Experiment 1. Finally, we constructed a second variation of the blocked condition that minimized the stimulus contrast along the bimodal dimension and maximized the stimulus contrast along the unimodal dimension (“blocked-sudden” condition; Figure 5, right-most

panel). This sequencing should further bias the participants toward the unimodal dimension beyond what was observed in the blocked condition of Experiment 1. Experiment 2 compares performance across these three training sequences under intentional learning conditions.

Method

Participants—Eighty-one students from the University of Texas at Austin participated in the study, either as a partial fulfillment of a course requirement or for pay. All participants were tested for 20/20 vision.

Stimuli and apparatus—The stimuli and presentation apparatus were identical to Experiment 1 and 2.

Procedure—There were three exemplar sequencing conditions as described above: zipper, blocked-gradual and blocked-sudden (see Figure 5). The actual training stimuli were selected randomly from the same distributions as in Experiment 1. For each participant, the first training stimulus was randomly selected from one of the four extreme stimuli.

Results and Discussion

The percentage of participants using the bimodal, unimodal and flat categorization strategies is presented in Figure 6. The strategy preference in each condition across participants is presented in Table 5. To summarize, participants in the zipper condition overwhelmingly selected the bimodal categorization strategy. Participants in the blocked-gradual condition were equally likely to select the unimodal strategy as the bimodal strategy, and participants in the blocked-sudden condition were more likely to select the unimodal dimension. Surprisingly, the blocked-gradual presentation did not prevent participants from focusing on the unimodal dimension and did not lead to better learning than observed in the blocked conditions from Experiment 1. One hypothesis motivating the blocked-gradual condition was that large changes along the unimodal dimension in the blocked condition in Experiment 1 led participants to use the unimodal strategy before the first stimulus from the second category was presented because of large stimulus changes along the unimodal dimension. The blocked-gradual finding in Experiment 2 suggests that participants in the blocked condition tend to form two categories even in the absence of large changes along the unimodal dimension. A more detailed examination of the results is reserved for the General Discussion.

General Discussion

As teachers, parents or students, we are often faced with situations in which we present, or are presented with, a body of information that must be grouped or organized. In many cases, there is some underlying structure to the information, but learning is unsupervised (or at the very least, trial by trial feedback is absent). This article investigates the effects of different training sequences on subsequent category formation when an underlying category structure exists. Sequencing (order) effects have been shown to affect both intentional supervised category learning (Goldstone, 1996; Jones, Love, & Maddox, 2006; Jones & Sieck, 2003; Stewart, Brown, & Chater, 2002) and unsupervised incidental category learning (Clapper, 2006; Clapper & Bower, 1994, 2002). However, very little is known about their effects on unsupervised intentional category learning, an everyday form of categorization.

In this article, we focused on this understudied form of categorization and investigated how exemplar sequencing affects peoples' ability to match their categorization strategy to an underlying category structure during intentional learning. We demonstrated that the order of initial stimuli and primarily a contrast between successive stimuli has a major effect on categorization strategy selection, and that this effect is different for intentional learning than

that previously found for incidental learning (Clapper, 2006; Clapper & Bower, 1994, 2002). Below, we summarize these results and discuss how stimulus contrast and other factors affect unsupervised category learning.

Exemplar sequencing and stimulus contrast effects in intentional unsupervised learning

Across two experiments, we examined unsupervised categorization strategy preferences in five conditions that varied the sequence of the training exemplars. The contrast between the current and the previous stimulus has been shown to strongly influence categorization responses in supervised intentional categorization tasks (Jones, Love, & Maddox, 2006; Jones & Sieck, 2003). Inspired by this research, we hypothesized that participants under intentional learning conditions may use the contrast between successive stimuli to guide categorization strategy selection. To quantify the stimulus-to-stimulus changes associated with each sequencing condition, we computed the cumulative stimulus value change along each dimension observed by each participant in each condition. Graphical depiction of the actual cumulative changes along each dimension observed in the five sequencing conditions investigated under intentional learning mode is presented in Figure 7. The associated categorization strategy preferences during early testing trials are presented in Figure 8.

Figures 7 and 8 show how the cumulative changes along the two dimensions affect the observed strategy preferences. First, participants were more likely to select the bimodal categorization strategy when the bimodal dimension showed larger cumulative changes. Second, conditions associated with larger changes along the unimodal than the bimodal dimension were associated with a lack of preference for the bimodal strategy, suggesting that emphasizing the unimodal dimension can occlude the underlying category structure. Interestingly, eliminating the large stimulus-to-stimulus changes along the unimodal dimension in the blocked-gradual condition was not sufficient to induce a preference for the bimodal categorization strategy. Rather, large changes along the bimodal dimension (as present in the intermixed condition and to the extreme in the zipper condition) were necessary to induce a preference for the bimodal categorization strategy.

To quantify the effect of the cumulative change along each dimension on strategy preferences, we ran a logistic regression analysis with the probability of selecting the bimodal vs. unimodal strategy as a dependent variable and the cumulative changes along each dimension as independent variables, including all participants across the five sequencing conditions. The odds of selecting a bimodal strategy (i.e., $p(\text{bimodal})/p(\text{unimodal})$) increase with increased cumulative change along the bimodal dimension ($b = 0.084$, $SE = 0.020$, $p < .001$, $\exp(b) = 1.087$), and decrease with increased cumulative change along the unimodal dimension ($b = -0.031$, $SE = 0.021$, $p = .139$, $\exp(b) = 0.967$). Adding dummy-coded variables for the five sequencing conditions into the logistic model did not significantly improve the fit of the model ($\chi^2(4) = 3.78$, $p = .626$), suggesting that the model provides a reasonable fit for all five sequencing conditions.

Although the explanatory effect of the unimodal dimension cumulative change beyond what was already explained by the bimodal dimension did not reach significance, it is important to note that this lack of significance may reflect the design of our experiments rather than signifying no effect (which would be surprising given the results of the blocked-sudden condition). As we aimed to explore extreme cases of sequencing, the resulting cumulative changes along the unimodal dimension were, across participants in all conditions, negatively correlated with the changes along the bimodal dimension ($r = -.360$, $p < .001$). Therefore, we cannot unambiguously estimate the independent effects of two dimensions. When testing for the effect of the unimodal dimension alone, the effect is significant ($b = -0.053$, $SE = .019$, $p = .006$, $\exp(b) = 0.948$).

The effect of the initial sequencing (Figure 8) stand in contrast to the findings of Colreavy and Lewandowsky (2008) who found no effect of the initial stimulus order on participants' strategy selection using a similar category structure. Importantly, though, they did not introduce explicit sequence manipulations like those included in the current report. To replicate their condition, we ran the logistic regression analysis including only data from our analogous, intermixed condition. In agreement with their findings, the effect of the cumulative change along either dimension was nonsignificant (all $p > .2$). Thus, as the variability of cumulative change along each dimension was fairly limited within each condition compared to the variability between conditions (see Figure 7), the between conditions comparison was needed to reveal the relationship.

Intentionality of learning

The results observed here under intentional learning conditions are in direct contrast to those observed previously under incidental learning conditions (Clapper, 2006; Clapper & Bower, 1994, 2002). For example, Clapper and Bower (1994) found minimal evidence for unsupervised category learning in their intermixed training condition but found significant evidence for category learning in their blocked condition. In a slightly different incidental paradigm, Medin and Bettger (1994) found superior category learning (as measured by recognition memory) when they minimized stimulus-to-stimulus changes compared to maximized stimulus-to-stimulus changes, just the opposite than observed here for intentional category learning. We hypothesized and tested in Experiment 1B that intentionality of learning is the main factor responsible for the different pattern of results. In contrast to intentional learning, we found that the participants in the incidental learning mode did show significant preference for the bimodal strategy in the blocked, but not the intermixed condition. Therefore, intentionality of learning needs to be taken into account when studying the mechanism of unsupervised category learning.

Other Factors

It is important to note that other factors, beyond the scope of this article, have also been shown to affect unsupervised category learning. First, the nature of the underlying category structure plays a major role in unsupervised learning (Anderson, 1991; Ashby, Queller, & Berretty, 1999; Billman & Knutson, 1996; Pothos & Chater, 2002; Stewart & Chater, 2002). For example, Pothos and Chater (2002) showed that highly separable categories (categories that maximize within category similarity and minimize between category similarity) are most intuitive. Second, prior knowledge aids unsupervised learning (Clapper, 2007; Kaplan & Murphy, 1999). For example, Clapper (2007) showed that prior knowledge makes different features less confusable and less interchangeable, increasing the probability that participants discover feature intercorrelations in complex category structures. Third, the learner's current task and goal also determines learned category representations (Love, 2003; Markman & Ross, 2003; Yamauchi, Love, & Markman, 2002). For example, Love (2003) showed that a particular task that a participant pursues during incidental category learning determines what category structures are likely to be acquired. How category structure, prior knowledge and participant's goals interact with exemplar sequencing and learning mode is a question for future research.

References

- Ahn, W.-k.; Medin, DL. A two-stage model of category construction. *Cognitive Science: A Multidisciplinary Journal* 1992;16(1):81–121.
- Anderson JR. The adaptive nature of human categorization. *Psychological Review* 1991;98(3):409–429.
- Ashby, FG. *Multidimensional models of perception and cognition*. Lawrence Erlbaum Associates, Inc.; Hillsdale, NJ, England: 1992. *Multidimensional models of categorization*; p. 449–483.

- Ashby FG, Queller S, Berretty PM. On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics* 1999;61:1178–1199. [PubMed: 10497436]
- Ashby FG, Townsend JT. Varieties of perceptual independence. *Psychological Review* 1986;93(2):154–179. [PubMed: 3714926]
- Billman D, Heit E. Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science* 1988;12(4):587–625.
- Billman D, Knutson J. Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 1996;22(2):458–475.
- Brainard DH. Psychophysics software for use with MATLAB. *Spatial Vision* 1997;10:433–436. [PubMed: 9176952]
- Clapper JP. When more is less: Negative exposure effects in unsupervised learning. *Memory & Cognition* 2006;34(4):890–902.
- Clapper JP. Prior knowledge and correlational structure in unsupervised learning. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale* 2007;61(2):109–127.
- Clapper JP, Bower GH. Category invention in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 1994;20(2):443–460.
- Clapper JP, Bower GH. Adaptive categorization in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 2002;28(5):908–923.
- Colreavy E, Lewandowsky S. Strategy development and learning differences in supervised and unsupervised categorization. *Memory & Cognition* 2008;36:762–775.
- Fried LS, Holyoak KJ. Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 1984;10(2):234–257.
- Goldstone RL. Isolated and interrelated concepts. *Memory & Cognition* 1996;24(5):608–628.
- Jones M, Love BC, Maddox WT. Recency Effects as a Window to Generalization: Separating Decisional and Perceptual Sequential Effects in Category Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2006;32(2):316–332.
- Jones M, Sieck WR. Learning myopia: An adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2003;29(4):626–640.
- Kaplan AS, Murphy GL. The acquisition of category structure in unsupervised learning. *Memory & Cognition* 1999;27(4):699–712.
- Love BC. The multifaceted nature of unsupervised category learning. *Psychonomic Bulletin & Review* 2003;10(1):190–197. [PubMed: 12747507]
- Markman AB, Ross BH. Category use and category learning. *Psychological Bulletin* 2003;129(4):592–613. [PubMed: 12848222]
- Medin DL, Bettger JG. Presentation order and recognition of categorically related examples. *Psychonomic Bulletin & Review* 1994;1(2):250–254.
- Medin DL, Wattenmaker WD, Hampson SE. Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology* 1987;19(2):242–279. [PubMed: 3581758]
- Nosofsky RM. Overall similarity and the identification of separable-dimension stimuli: A choice model analysis. *Perception & Psychophysics* 1985;38(5):415–432. [PubMed: 3831920]
- Nosofsky RM. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* 1986;115(1):39–57. [PubMed: 2937873]
- Pelli DG. The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision* 1997;10:437–442. [PubMed: 9176953]
- Pothos EM, Chater N. A simplicity principle in unsupervised human categorization. *Cognitive Science* 2002;26(3):303–343.
- Regehr G, Brooks LR. Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 1995;21(2):347–363.
- Schwarz G. Estimating the dimension of a model. *Annals of Statistics* 1978;6:164–464.
- Shepard RN. Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology* 1964;1:54–87.

- Stewart N, Brown GDA, Chater N. Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2002;28(1):3–11.
- Stewart N, Chater N. The effect of category variability in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2002;28(5):893–907.
- Yamauchi T, Love BC, Markman AB. Learning nonlinearly separable categories by inference and classification. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 2002;28(3):585–593.
- Zeithamova D, Maddox WT. The role of visuospatial and verbal working memory in perceptual category learning. *Memory & Cognition* 2007;35(6):1380–1398.

Figure 1.

Scatterplot of the stimuli used in all experiments. Each point represents a line with a particular value on the bimodal (relevant) dimension (e.g. length) and on the unimodal (irrelevant) dimension (e.g. orientation). Black circles denote members of one category, gray squares denote members of another category. Lengths ranged from approximately 100 - 500 pixels, and orientations ranged from 0 - 90 degrees.

Figure 2.

Examples of actual categorization responses for three participants who each used a different response strategy. Each symbol denotes a stimulus (line) with a particular length (x-axis) and orientation (y-axis). Star symbols denote the stimuli that the participant labeled as category A stimuli, open diamonds denote the stimuli that the participant labeled as category B stimuli. Left panel: responses of a participant who based their categorization decision on the bimodal dimension; middle panel: responses of a participant that based their categorization decision on the unimodal dimension; right panel: responses of a participant who was best fit by a flat response model.

Figure 3.

Percentage of participants within each group using the bimodal (black bars), unimodal (gray bars), and flat (white bars) categorization strategies. Left panel: strategy use during first 20 testing trials. Right panel: strategy use across 160 testing trials.

Figure 4.

Percentage of participants within each group using the bimodal (black bars), unimodal (gray bars), and flat (white bars) categorization strategies in Experiment 1B. Left panel: strategy use during first 20 testing trials. Right panel: strategy use across 160 testing trials.

Figure 5.

A schematic of the stimulus sequence order (1-20) for the 20 training trials in each condition of Experiment 3. Left panel: Ordering of stimuli in the zipper condition. Middle panel: Ordering of stimuli in the blocked-gradual condition. Right panel: Ordering of stimuli in the block-sudden condition.

Figure 6.

Percentage of participants using the bimodal (black bars), the unimodal (gray bars) and flat (white bars) categorization strategies in Experiment 2. Left panel: Best fitting strategy during first 20 testing trials. Right panel: Best fitting strategy across 160 testing phase trials.

Figure 7.

Cumulative stimulus-to-stimulus changes along the bimodal (black bars) and unimodal (gray bars) dimension observed during training by participants in different sequencing conditions. The data are in normalized units corresponding to a standard deviation of stimulus values along the unimodal dimension (100 pixels for length and 15 degrees for orientation). The error bars represent standard deviation across participants.

Figure 8.

Percentage of participants using the bimodal (black bars), unimodal (gray bars) and flat (white bars) categorization strategies during the early testing trials from five different exemplar sequencing conditions. Data from the zipper, blocked-gradual, and blocked-sudden conditions are taken from Experiment 2. Data for the intermixed and blocked condition were pooled across Experiment 1A (intermixed and blocked (AB) condition) and Experiment 1B (intentional intermixed and intentional blocked condition). Asterix denotes significant preference for a given strategy in that condition.

Table 1

Category structure parameters used in all experiments. Length given in pixels, orientation in degrees.

Length bimodal	μ_A	μ_B	σ
Length	187.5	412.5	12.5
Orientation	45	45	15
Orientation bimodal	μ_A	μ_B	σ
Length	300	300	100
Orientation	28.1	61.9	1.88

Note: Covariance is 0 in each category

Table 2

Model fit results from Experiment 1A. Mean (SEM) BIC goodness-of-fitting measure and proportion of participants' responses accounted for by bimodal model, unimodal model, and flat (one-parameter) model, separately for participants who were best fit by the bimodal, unimodal and flat response strategy. The fit of the zero-parameter version of the flat model depends solely on the number of trials (as there is always .5 prediction error for every response) and was thus the same for all participants: BIC = 28 for early 20 test trials, BIC = 222 across the testing phase, and proportion of responses accounted for = 0.5. *Early*: fit results for first 20 testing trials; *across*: fit results across 160 trials.

Strategy used	bimodal	BIC unimodal	flat	bimodal	accounted for unimodal	flat
early	bimodal	12 (1.3)	30 (0.4)	0.94 (0.01)	0.62 (0.02)	0.59 (0.01)
	unimodal	31 (0.7)	29 (0.6)	0.64 (0.02)	0.93 (0.01)	0.63 (0.02)
	flat	27 (2.4)	25 (2.4)	0.68 (0.04)	0.69 (0.04)	0.66 (0.04)
across	bimodal	63 (8.4)	226 (0.5)	0.94 (0.01)	0.54 (0.01)	0.52 (0.01)
	unimodal	217 (5.7)	207 (5.4)	0.63 (0.02)	0.87 (0.02)	0.64 (0.02)
	flat	224 (10.6)	213 (10.6)	0.60 (0.04)	0.60 (0.05)	0.59 (0.04)

Table 3

Categorization strategy preference across participants in each condition of Experiment 1A. Preference results and the p-values reflect a binomial test of a null hypothesis that each strategy was selected equally often against an alternative hypothesis that the proportion of participants using one strategy was greater than the other strategy. *bimodal*: significant preference for bimodal over unimodal strategy; *neither*: no significant preference for either strategy. Unimodal strategy was not preferred significantly in any condition of Experiment 1A.

condition	early	across
<i>intermixed</i>	bimodal (p = .006)	bimodal (p = .001)
<i>blocked (AB)</i>	neither (p = .593)	neither (p = .407)
<i>blocked (A only)</i>	neither (p = .500)	neither (p = .227)

Table 4

Categorization strategy preference across participants in each condition of Experiment 1B. Preference results and the p-values reflect a binomial test of a null hypothesis that each strategy was selected equally often against an alternative hypothesis that the proportion of participants using one strategy was greater than the other strategy. *bimodal*: significant preference for bimodal over unimodal strategy; *neither*: no significant preference for either strategy. Unimodal strategy was not preferred significantly in any condition of Experiment 1B.

condition	early	across
<i>incidental</i>		
<i>intermixed</i>	bimodal (p = .032)	bimodal (p = .010)
<i>blocked</i>	neither (p = .402)	neither (p = .314)
<i>intentional</i>		
<i>intermixed</i>	neither (p = .500)	neither (p = .180)
<i>blocked</i>	bimodal (p = .011)	bimodal (p = .018)

Table 5

Categorization strategy preference across participants in each condition of Experiment 2. Preference results and the p-values reflect a binomial test of a null hypothesis that each strategy was selected equally often against an alternative hypothesis that the proportion of participants using one strategy was greater than the other strategy. *bimodal*: significant preference for bimodal over unimodal strategy; *neither*: no significant preference for either strategy

condition	early	across
<i>zipper</i>	bimodal ($p < .001$)	bimodal ($p < .001$)
<i>blocked-gradual</i>	neither ($p > .5$)	neither ($p > .5$)
<i>blocked-sudden</i>	unimodal ($p = .010$)	* unimodal ($p = .061$)

* *unimodal*: significant (marginal) preference for unimodal over bimodal strategy.