

UC San Diego

UC San Diego Previously Published Works

Title

Tracking the emergence of memories: A category-learning paradigm to explore schema-driven recognition

Permalink

<https://escholarship.org/uc/item/2ps14151>

Journal

Memory & Cognition, 45(1)

ISSN

0090-502X

Authors

De Brigard, Felipe
Brady, Timothy F
Ruzic, Luka
[et al.](#)

Publication Date

2017

DOI

10.3758/s13421-016-0643-6

Peer reviewed

Tracking the emergence of memories: A category-learning paradigm to explore schema-driven recognition

Felipe De Brigard^{1,2,3} · Timothy F. Brady⁴ · Luka Ruzic^{2,5} · Daniel L. Schacter^{6,7}

© Psychonomic Society, Inc. 2016

Abstract Previous research has shown that prior knowledge structures or *schemas* affect recognition memory. However, since the acquisition of schemas occurs over prolonged periods of time, few paradigms allow the direct manipulation of schema acquisition to study their effect on memory performance. Recently, a number of parallelisms in recognition memory between studies involving schemas and studies involving category learning have been identified. The current paper capitalizes on these findings and offers a novel experimental paradigm that allows manipulation of category learning between individuals to study the effects of schema acquisition on recognition. First, participants learn to categorize computer-generated items whose category-inclusion criteria differ between participants. Next, participants study items that belong to either the learned category, the non-learned

category, both, or neither. Finally, participants receive a recognition test that includes old and new items, either from the learned, the non-learned, or neither category. Using variations on this paradigm, four experiments were conducted. The results from the first three studies suggest that learning a category increases hit rates for old category-consistent items and false alarm rates for new category-consistent lures. Absent the category learning, no such effects are evident, even when participants are exposed to the same learning trials as those who learned the categories. The results from the fourth experiment suggest that, at least for false alarm rates, the effects of category learning are not solely attributable to frequency of occurrence of category-consistent items during learning. Implications for recognition memory as well as advantages of the proposed paradigm are discussed.

Electronic supplementary material The online version of this article (doi:10.3758/s13421-016-0643-6) contains supplementary material, which is available to authorized users.

Keywords Schema · Categorization · Recognition · Memory · False alarms

✉ Felipe De Brigard
felipe.debrigard@duke.edu

Introduction

Since Bartlett's pioneering work, it has been known that episodic memory is influenced by previously acquired knowledge. Bartlett's influential insight was that memory involves "an active organization of past reactions" (Bartlett, 1932, p. 201), so that every new experience is encoded and retrieved not only as an individual event, but also as related to a knowledge structure of previously encoded similar events. Bartlett called this knowledge structure a *schema*, but this general idea later reemerged in the literature under different names, such as *script* (Schank & Abelson, 1977), *frame* (Minsky, 1975), and *gist* (Brainerd & Reyna, 1990). So enduring is the notion of schema that it is still widely discussed in memory research, as evidenced by recent studies in cognitive neuroscience

- ¹ Department of Philosophy, Duke University, 203A West Duke Building, Durham, NC 27708-0743, USA
- ² Center for Cognitive Neuroscience, Duke University, Durham, NC, USA
- ³ Duke Institute for Brain Sciences, Durham, NC, USA
- ⁴ Department of Psychology, University of California, San Diego, La Jolla, CA, USA
- ⁵ Department of Psychology, Duke University, Durham, NC, USA
- ⁶ Center for Brain Science, Harvard University, Cambridge, MA, USA
- ⁷ Department of Psychology, Harvard University, Cambridge, MA, USA

exploring its neural underpinnings (e.g., van Kesteren, Ruiters, Fernández, & Henson, 2012).

Traditionally, studies on the effects of schemas on memory have produced two kinds of results (Lampinen, Copeland, & Neuschatz, 2001). First, schematic knowledge increases recognition of schema-inconsistent information relative to schema-consistent information (Bower, Black, & Turner, 1979; Rojahn & Pettigrew, 1992). Second, it increases false alarms to schema-consistent lures relative to schema-inconsistent lures (Brewer & Treyners, 1981; Lampinen et al., 2001). In addition, recent results show that schemas also influence short-term memory (Hollingworth & Henderson, 2003; Oliva, 2005), suggesting they can affect retention as early as in working memory (Brady, Konkle, & Alvarez, 2009) and as late as in long-term memory retrieval (Hemmer & Steyvers, 2009; Steyvers & Hemmer, 2012).

Unfortunately, because the acquisition of schemas is thought to occur over prolonged periods of time, it is not easy to track the process of acquisition beginning with learning the background knowledge to the study material and memory test in the controlled environment of the laboratory. Thus, attempts to directly assess how memory performance is affected by differences in acquired knowledge typically involve comparing within-subject performance for different pre-acquired schemas (Graesser & Nakamura, 1982; Roediger & McDermott, 1995; Taylor & Crocker, 1981) or between-subject performance for individuals with different schematic expertise (e.g., Arkes & Freedman, 1984; Castel, McCabe, Roediger, & Heitman, 2007; Chase & Simon, 1973; de Groot, 1966). Therefore, these experimental paradigms do not allow the direct manipulation of schema acquisition to study their effect on memory performance.

Recently, however, researchers have started to remark on a number of similarities between extant results from memory studies on schema-driven recognition and studies on category learning that include a recognition memory component (Sakamoto & Love, 2004; Sakamoto, 2012). In one such study, for instance, Palmeri and Nosofsky (1995) asked participants to learn to classify 16 geometric stimuli according to a simple rule. Although most stimuli fit the rule, the learning list included exceptions. A subsequent recognition memory test showed that subjects recognized category-inconsistent items at a higher rate than category-consistent ones. This recognition advantage for rule-inconsistent relative to rule-consistent exemplars parallels the aforementioned recognition advantage for schema-inconsistent versus schema-consistent items. But the parallels do not end here. In a meta-analysis on schema-dependent recognition memory, Rojahn and Pettigrew (1992) report that the recognition advantage for schema-inconsistent information increases as the proportion of schema-inconsistent to schema-consistent items becomes smaller. To investigate whether this effect was evident in studies involving category learning, Sakamoto and Love (2004)

manipulated the strength of the category rule by varying the frequency of the category-consistent items during learning. They demonstrated that when the rule was stronger (e.g., included more rule-consistent items during learning), exceptional items were remembered at a higher rate than when the rule was weaker.

These results have led a number of researchers to postulate strong connections between the cognitive processes underlying schematic and categorical learning (e.g., Love, 2013; Davis, Xue, Love, Preston, & Poldrack, 2014; Sakamoto & Love, 2004; Sakamoto, 2012). Sakamoto (2012) has recently argued that these results on recognition memory and category learning challenge existing models of category learning—such as exemplar and prototype-based models—and also lend credence to cluster-based models (e.g., SUSTAIN; Love, Medin, & Gureckis, 2004) that take category learning as tantamount to schema acquisition. Specifically, following the notion that a schema is a knowledge structure that provides a set of expectations based on past experiences, Sakamoto (2012) argues that category learning should be understood as the process of building “schema-like representations in which rule-following items are encoded as a set of expectations, and rule-violating items are stored separately” (p. 2961). Unfortunately, as Sakamoto also remarks, very little work has been done linking category learning and schema acquisition in memory research, despite asking similar questions and finding parallel results (but see Clapper, 2008). The relatively few and recent attempts to bridge these two disciplines have come from the category learning side, and oftentimes measures that are of interest to memory researchers (e.g., separate analyses on hits and false alarms) have not been directly addressed by the experimental paradigms employed.

The current paper seeks to contribute to the integration of these two disciplines from the perspective of memory research by offering a novel experimental paradigm that can help to extend the use of category-learning manipulations to study the effect of structured knowledge acquisition on hit and false alarm rates during recognition memory tests. The paradigm allows researchers to manipulate different features of category learning and encoding between individuals to study the effects of schema acquisition on hit and false-alarm rates during subsequent memory performance.

The paradigm consists of three stages. In the first stage (*Learning*), participants learn to categorize computer-generated items—in this case, flowers—either as belonging or as not belonging to a category. The current variation employs a rule-based strategy, but other strategies can be employed as well (Ashby & O’Brien, 2005). In addition, although the structure of the paradigm allows the inclusion of exceptions to explore schema inconsistent effects, the current variation reported here does not include category-exceptions in the learning stage—i.e., every item that belongs to the category follows the rule—since our interest is to explore

schema-consistent effects on recognition memory. Here, the feature that determines category membership—i.e., the rule—which is counterbalanced across participants, occurs in half of the trials. The flowers that include this feature belong to the learned category. Another feature, also counterbalanced across participants, occurs as frequently as the learned category feature and coincides with flowers from the learned category half of the time. Flowers that include this second feature belong to what we refer to as the non-learned category. As such, the main difference between the learned and the non-learned category rests on the precise feedback participants receive, which in this case aims exclusively to help them categorize flowers from the learned category. Thus, despite the non-learned category sharing the same statistical features as the learned category, participants did not engage in any explicit task to learn it, or even acknowledge its presence. Finally, the remaining flowers (one-quarter) belong to neither category.

In the next stage (*Study*), participants study flowers that belong to either the learned category, the non-learned category, both, or neither. In the final stage (*Test*), participants receive a recognition test that includes old flowers as well as new flowers (lures) from the learned category, the non-learned category, or neither category. Of note, none of the flowers seen during the study and test stages overlapped with the flowers seen during the learning stage. We employed four variations of the paradigm. Experiment 1 implements the paradigm as explained above. In Experiment 2, the learning stage is eliminated to test whether the effects on recognition memory observed in Experiment 1 were attributable to or independent of prior category learning. Experiment 3 also tests whether the differential effects on recognition memory were attributable to category learning, by increasing the difficulty of the category learning task so that some participants learned the category and other participants did not. This manipulation allows subsequent memory performance to be compared with otherwise identical conditions between those who learned and did not learn the category. Finally, in Experiment 4 we test the hypothesis, derived from category-learning models, that the effects on recognition memory found in the first three experiments may be attributable to the relative frequency of presentation of the category-consistent items relative to category inconsistent-items.

Consistent with previous studies on recognition memory for schema-consistent items (Brainerd and Reyna, 2005; Lampinen et al., 2001), we expected that participants who learned to sort flowers into a particular category will have higher hit rates for old items of the learned category relative to old items of the non-learned category, despite the fact that items of the non-learned category were presented with the same frequency as items of the learned category during learning. Likewise, we expected to see higher false alarm rates for lures of the learned category relative to lures of the non-

learned category. In turn, we expected lower hit and false alarm rates for flowers that did not belong to either the learned or the non-learned categories. Additionally, we expected that in the absence of learning these effects would not be evident. Finally, we conjectured that even after equating for frequency of presentation during encoding, the effect of schema acquisition would still be evident in a recognition task.

Experiment 1

Methods


Participants One hundred individuals were recruited and run using Amazon Mechanical Turk (<https://www.mturk.com>). All participants were from the USA, gave informed consent, and received approximately US\$0.60 for participating.

Stimuli For this study, 1,024 different flowers were generated in MATLAB. Flowers varied across five possible dimensions, with each dimension taking up one of four possible values, as follows: (a) number of petals: 2, 4, 6, or 8; (b) color of petals: red, blue, green, or yellow; (c) shape of center: circle, triangle, square, or star; (d) color of the center: orange, pink, bright green, or turquoise; and (e) number of sepals: 0, 1, 2, or 3. Flowers appeared in the center of the screen against a white background (Fig. 1).


Procedure The experiment consisted of three stages: *learning*, *study*, and *test* (Fig. 1). For the *learning* stage, participants were told they will see a flower on the screen, and that they will be asked whether or not such a flower belonged to the species *avlonia*. Participants were told that *avlonias* differed from other flowers in one simple way (e.g., only *avlonias* had four petals, or pink sepals, or a heart in the center), and that their task was to find out what that simple way was. Participants were informed of the five dimensions across which all the flowers varied and shown two example flowers to highlight these dimensions. They were then told they would see a total of 60 flowers, one at a time. Whenever a flower appeared on the screen, the question “Is this an *avlonia*?” was also displayed, and participants were asked to answer “Yes” or “No.” Each trial was self-paced, and the flower and question were displayed through the trial. The word “Correct” or “Incorrect” was presented as feedback immediately following the participant’s response. Participants were told that at the beginning they would have to guess, but that eventually they would figure out what makes a flower an *avlonia*.

There were four kinds of trials during the *learning* phase. One-quarter of the trials showed flowers that included the learned category feature—i.e., the feature that made flowers *avlonias* (e.g., yellow center). One-quarter of the trials displayed flowers with a different feature, which determined

a. Learning phase



b. Study phase



c. Test phase




Fig. 1 Experimental paradigm. **a** Learning phase: 60 MATLAB generated flowers varying across five dimensions, each one taking up four possible values (number of petals [2, 4, 6, 8], color of petals [red, blue, green, yellow], shape of center [circle, triangle, square, star], color of center [yellow, pink, turquoise, bright green], number of sepals [0, 1, 2, 3]). For each flower, participants responded “Yes” or “No” to the criterion-inclusion question (“Is this an avlonia?”), and received immediate feedback. Trials were self-paced. The learning phase lasted at most 7 min. **b** Study phase: after a 2-min display of the instructions for this stage, 20 study flowers, not included in the learning phase, were presented randomly for 5 s each. Five such flowers from the Learned category (in this example, the category-inclusion feature is “yellow center”), five were

flowers from a Not Learned category (in this example, the category inclusion feature is “six petals”), five were flowers from Both categories, and five were flowers from Neither category. The study phase lasted about 1.5 min. **c** Test phase: after a 2-min display of the instructions for this last stage, 50 test flowers were presented randomly for a recognition test. Twenty such flowers were old (i.e., presented at test), and 30 were new. Ten of the new flowers were not previously presented lures from the Learned category (Learned-Lures), ten were not previously presented lures from the Not Learned category (Not Learned-Lures), and ten were not previously presented lures from Neither category (Neither-Lures). The test phase was self-paced and lasted at most 3.5 min

the category inclusion for a non-learned category (e.g., six petals). One-quarter of the trials showed flowers with both features (e.g., yellow center and six petals), so they belonged to both the learned and the non-learned category. Finally, one-quarter of the trials displayed flowers that belonged to neither category. Therefore, in effect, 50 % of the trials should produce “Yes” responses, but of those only half were unique to either the learned category or the non-learned category. The features that determined both category memberships were counterbalanced across participants, and the remaining features were randomly selected. (Table 1 displays the abstract category structure of the learning phase. Figure 2 displays a frequency distribution matrix of the average pairwise co-occurrence of each dimension’s values across trials during the learning phase.)

After the *learning* phase, participants saw a screen with the instructions for the second phase: *study*. This screen

was displayed for a minimum of 90 s. On this instructions screen, participants were told that in the next part of the experiment they would see 20 flowers one at a time, and that their task was to remember them as best as they could. They were told that they would receive an extra bonus for remembering them all. They were also told that some of the flowers they would have to remember may or may not be avlonias. During this *study* phase, participants saw five flowers from the learned category, five from the non-learned category, five from both categories, and five from neither, in random order. Each flower was displayed for 5 s, followed by a 1-s blank and then the next flower. Flowers were counterbalanced across participants, and none of these 20 study flowers were presented during the *learning* phase.

Finally, after the *study* phase, participants saw a screen with the instructions for the final phase: *test*. This screen

Table 1 Abstract category structure used in the learning stage of Experiment 1. Each row represents a unique stimulus (flower) from one of four possible groups of 15 stimuli each. Each number represents one of four possible values for each dimension. Here, the first value of the first dimension (1,1) corresponds to the category inclusion value for Learned-Flowers, while the first value of the second dimension (2,1) corresponds to the category-inclusion value for Not Learned-Flowers. All Both-Flowers include both values (1,1) and (2,1), whereas none of the Neither-Flowers does. Which value corresponded to the inclusion criteria for both Learned and No Learned was counterbalanced between subjects, and every value was used the same amount of times across subjects. The other values were chosen randomly

Category	Stimulus	Dimension value				
		1	2	3	4	5
Learned	L1	1	2	1	1	1
	L2	1	2	1	1	1
	L3	1	2	1	1	2
	L4	1	2	1	2	2
	L5	1	2	2	2	3
	L6	1	3	2	2	3
	L7	1	3	2	3	4
	L8	1	3	2	3	4
	L9	1	3	3	3	1
	L10	1	3	3	4	1
	L11	1	4	3	4	2
	L12	1	4	3	4	2
	L13	1	4	4	1	3
	L14	1	4	4	2	3
	L15	1	4	4	3	4
Not Learned	NL1	2	1	1	1	1
	NL2	2	1	1	1	1
	NL3	2	1	1	1	2
	NL4	2	1	1	2	2
	NL5	2	1	2	2	3
	NL6	3	1	2	2	3
	NL7	3	1	2	3	4
	NL8	3	1	2	3	4
	NL9	3	1	3	3	1
	NL10	3	1	3	4	1
	NL11	4	1	3	4	2
	NL12	4	1	3	4	2
	NL13	4	1	4	1	3
	NL14	4	1	4	2	3
	NL15	4	1	4	3	4
Both	B1	1	1	1	1	1
	B2	1	1	1	1	1
	B3	1	1	1	1	2
	B4	1	1	1	2	2
	B5	1	1	2	2	3
	B6	1	1	2	2	3
	B7	1	1	2	3	4
	B8	1	1	2	3	4

Table 1 (continued)

Category	Stimulus	Dimension value				
		1	2	3	4	5
Neither	B9	1	1	2	3	1
	B10	1	1	3	4	1
	B11	1	1	3	4	2
	B12	1	1	3	4	2
	B13	1	1	3	1	3
	B14	1	1	4	2	3
	B15	1	1	4	3	4
	N1	2	2	1	1	1
	N2	2	2	1	1	1
	N3	2	2	1	1	2
	N4	2	2	2	2	2
	N5	2	3	2	2	3
	N6	3	3	2	2	3
	N7	3	3	3	3	4
	N8	3	3	3	3	4
N9	3	4	3	3	1	
N10	3	4	3	4	1	
N11	4	4	4	4	2	
N12	4	4	4	4	2	
N13	4	2	4	1	3	
N14	4	3	4	2	3	
N15	4	4	1	3	4	

was displayed for a minimum of 90 s. Participants were told that in this final stage 50 flowers would be shown, one at a time, and they would be asked to answer, for each flower, whether or not it was in the study list. Of the 50 test flowers, 20 were old, ten were new flowers from the learned category (Learned-lures), ten were new flowers from the non-learned category (Not learned-lures), and ten were new flowers from neither category (Neither-lures). None of these 50 flowers were presented during the *learning* phase. Flowers were displayed one by one in random order, and each trial was self-paced.

Results

Performance data during the learning phase are depicted on Fig. 3a, and recognition data during the test phase are depicted in Fig. 3b (Hits) and c (FA). Hit and FA rates were analyzed separately. An ANOVA comparing hit rates for old items from all four categories (Learned, Not Learned, Both, Neither) revealed a main effect of Category, $F(3, 97) = 5.359, p = .002, \eta^2 = .14$. Paired direct contrasts showed that the hit rate for items from the

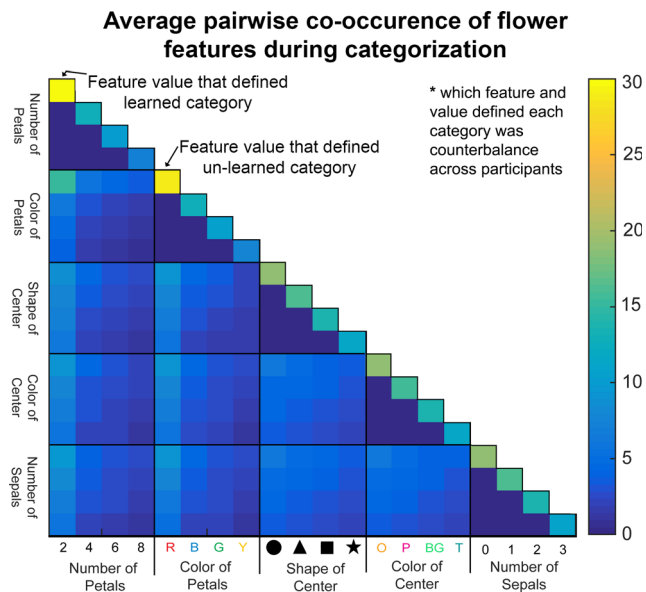


Fig. 2 Frequency distribution matrix depicting average pairwise co-occurrence of flower's features during the learning phase in Experiment 1. The groups of columns correspond to the five dimensions across which flowers varied, and the columns within them correspond to the particular values along those dimensions. Colors indicate the frequency of presentation of that value for that dimension across trials. As such, the first dimension (thick column) corresponds to the learned rule dimension, and the leftmost value (the first thin column within the dimension) corresponds to the learned rule value. In this illustration, the category inclusion value for the learned category corresponds to two petals. The second dimension corresponds to the not learned rule dimension, and the leftmost value to the not learned rule. In this illustration, the value for the not learned category corresponds to red petals

Learned category ($M = .71$, $SD = .17$) was significantly higher than for items from the Not Learned category ($M = .67$, $SD = .20$), $t(99) = 2.58$, $p = .011$, $d = .22$, as well as for items from Neither category ($M = .61$, $SD = .21$), $t(99) = 3.88$, $p < .001$, $d = .52$. Likewise, the hit rate for items from the Both categories was significantly higher ($M = .70$, $SD = .25$) than for items from the Not Learned, $t(99) = 2.02$, $p = .046$, $d = .13$, as well as the Neither category, $t(99) = 3.10$, $p = .002$, $d = .39$. Interestingly, the hit rate for items from the Not Learned category was higher than for Neither, $t(99) = 2.59$, $p = .011$, $d = .29$. There was no difference in the hit rate for items from the Learned and the Both Categories, $p = .699$.

An ANOVA comparing FA rates for Lures of the Learned, Not Learned, and Neither categories revealed a main effect of Lure, $F(2, 98) = 17.45$, $p < .001$, $\eta^2 = .26$. Paired direct contrasts revealed that FA rates for lures from the Learned category ($M = .54$, $SD = .21$) were higher than for lures of the Not Learned ($M = .44$, $SD = .22$), $t(100) = 4.20$, $p < .001$, $d = .47$, as well as the Neither categories ($M = .40$, $SD = .19$), $t(100) = 5.75$, $p < .001$, $d = .70$. The difference between the FA rates for lures from the Not Learned and the Neither categories,

although close, did not reach significance, $t(100) = 1.89$, $p = .061$, $d = .20$.¹

To assess differences in discriminability in recognition memory for items from the Learned, Not Learned, and Neither categories, two further analyses were conducted. First, we calculated differences in Hit minus FA rates for items from all three conditions (Learned: $M_{[\text{Hit-FA}]} = .17$, $SD = .29$; Not Learned: $M_{[\text{Hit-FA}]} = .23$, $SD = .30$; Neither: $M_{[\text{Hit-FA}]} = .21$, $SD = .30$), and performed an ANOVA across the three measures. No effects were found ($p > .05$). Second, sensitivity and bias measures were estimated in a signal detection theory (SDT) framework by calculating measures of d' and C for all three conditions (Learned: $d' = .56$, $C = -0.37$; Not Learned: $d' = .72$, $C = -0.18$; Neither: $d' = .65$, $C = -0.01$). A one-way ANOVA comparing measures of d' revealed no effects of sensitivity ($p > .05$). However, a second one-way ANOVA comparing measures of C revealed a main effect of condition, $F(2, 98) = 19.12$, $p < .001$, $\eta^2 = .28$. Follow-up paired t -tests indicated higher bias in recognition responses for items from the Learned relative to the Not Learned, $t(99) = 4.12$, $p < .001$, $d = .41$, and the Neither category, $t(99) = 6.07$, $p < .001$, $d = .82$, as well as for the Not Learned relative to Neither category, $t(99) = 3.19$, $p = .002$, $d = .37$.

Discussion

The results of Experiment 1 suggest that participants who learned to categorize flowers according to a certain criterion had higher hit rates for category-consistent items as well as higher FA rates for category-consistent lures even if, during study, they saw an equal number of flowers that belonged to another, non-learned category. These results show how acquiring categorical knowledge can reliably increase the probability of both correctly and falsely recognizing items that instantiate the learned category in a subsequent memory test. Moreover, the current paradigm demonstrates that such

¹ After conducting this study, we noticed a minor bug in the code during the learning phase that made the flowers from the Not Learned and the Both categories appear between 14 and 16 times, rather than a fixed 15 times each, as with the flowers from the Learned and the Neither categories. Although the difference is negligible, we conducted a replication of this study ($N = 100$) in which each flower trial was presented exactly 15 times. The effect of Lure replicated, $F(2, 98) = 24.47$, $p < .001$, $\eta^2 = .33$, and the directed contrasts were virtually identical, with FA rates for lures of the Learned category ($M = .54$, $SD = .21$) being higher than for Not Learned ($M = .47$, $SD = .21$), $t(100) = 2.91$, $p = .004$, and for Neither ($M = .39$, $SD = .21$), $t(100) = 7.03$, $p < .001$. The difference between FA rates for lures of the Not Learned and the Neither categories did reach significance in the replication, $t(100) = 3.72$, $p < .001$. On the other hand, the Hit rates effect did not replicate, $F(3, 97) = 1.84$, $p = .14$, although the hit rate for items from Neither category was still numerically lower ($M = .60$, $SD = .21$) than for Learned ($M = .65$, $SD = .20$), Not Learned ($M = .64$, $SD = .21$), and Both ($M = .66$, $SD = .28$). This suggests that perhaps the effect on FA alarms is more robust than the effect on Hit rates (see results in Fig. 8, Supplementary Materials).




Fig. 3 Learning and recognition performance from Experiment 1. **a** Categorization performance during the learning phase. **b** Hit rates during the test phase. H-Learned corresponds to the hit rates for items from the Learned category; H-Not Learned corresponds to the hit rates for items from the Not Learned category; H-Both corresponds to the hit rates for items of Both categories; and H-Neither corresponds to the hit rates for

items from Neither category. Error bars indicate SEM. **c** False alarm rates during the test phase. FA-Learned corresponds to the false alarm rates for lures of the Learned category; FA-Not Learned corresponds to the false alarm rate for lures of the Not Learned category; and FA-Neither corresponds to the false alarm rate for lures of Neither category

learning can occur quickly and in conditions under which low-level perceptual features have been controlled for. In addition, the differences in measures of bias found in Experiment 1 suggest that learning a category increases the participant's susceptibility toward answering "old" for items of the learned category relative to items from the non-learned or from neither category.

However, although the counterbalanced design in Experiment 1 allowed us to find differences in Hit and FA rates for items from the Learned and the Not Learned categories, it may still be possible that the effects found for items from Neither category were at least partly due to features of the study and testing materials. Likewise, the fact that we find essentially no bias for items of the Neither category relative to both the Not Learned and the Learned categories, suggests that at least some of the effects on discriminability may be due to features of the study and testing phases rather than learning alone. Thus, to assure that the effects in hit and FA rates were only attributable to learning differences rather than study or testing differences, we conducted a second experiment that did not include a learning stage. We hypothesized that if the differences in hit rates for items of the Learned versus Not Learned and Neither categories, as well as the between-group differences in the Learned versus the Not Learned and Neither categories in FA rates observed in Experiment 1 were due to differences in learning alone, then these effects would disappear in Experiment 2.

Experiment 2

Methods

Participants One hundred individuals were recruited and run with Amazon Mechanical Turk (<https://www.mturk.com>). All

participants were from the USA, gave informed consent, and received approximately US\$0.60 for participating.

Stimuli and procedure Stimuli and procedure were identical to Experiment 1 except Experiment 2 did not include the learning phase, only the study and test phases.

Results

Recognition data during the test phase are depicted in Fig. 4a (Hits) and b (FA). An ANOVA was employed to compare hit rates for old items from all four categories (Learned, Not Learned, Both, Neither). No effect was found, $p = .095$. An ANOVA comparing FA rates for Lures of the Learned, Not Learned and Neither categories revealed a main effect of Lure, $F(2, 98) = 7.89$, $p = .001$, $\eta^2 = .14$. Paired direct contrasts revealed no difference between FA rates for lures of the Learned ($M = .48$, $SD = .22$) and the Not Learned categories ($M = .48$, $SD = .22$), $t(100) = .13$, $p = .90$. However, FA rates for lures of the Learned and Not Learned categories were higher than for lures of Neither category ($M = .40$, $SD = .21$), smallest $t(100) = 3.44$, $p = .001$, $d = .37$.

As in Experiment 1, discriminability in recognition memory was assessed, first by calculating differences in Hit minus FA rates for items from all three conditions (Learned: $M_{[\text{Hit-FA}]} = .21$, $SD = .31$; Not Learned $M_{[\text{Hit-FA}]} = .20$, $SD = .30$; Neither: $M_{[\text{Hit-FA}]} = .23$, $SD = .29$), and performing an ANOVA across the three measures. No effects were found ($p > .05$). Second, measures of d' and C were also calculated (Learned: $d' = .64$, $C = -.26$; Not Learned: $d' = .61$, $C = -.24$; Neither: $d' = .72$, $C = -.02$), and two one-way ANOVAs comparing them across conditions were conducted. While no effects for d' were revealed ($p > .05$), for C there was a main effect of condition, $F(2, 98) = 7.73$, $p = .001$, $\eta^2 = .14$. Follow-up paired t -tests indicated lower bias in recognition responses

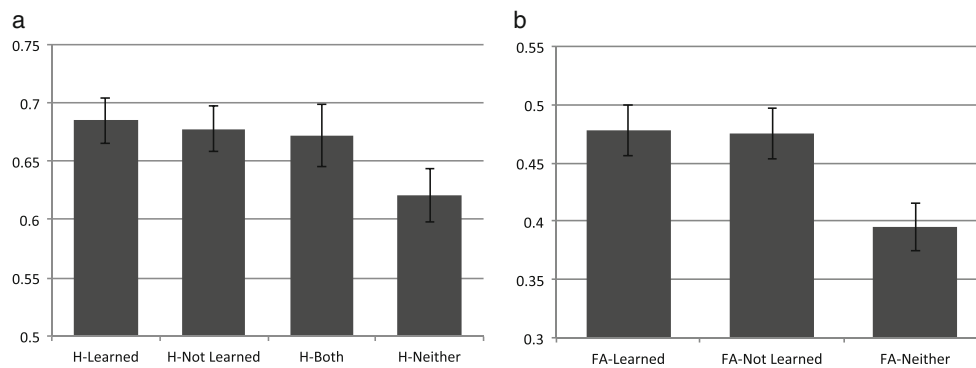


Fig. 4 Recognition performance from Experiment 2. **a** Hit rates during the test phase. H-Learned corresponds to the hit rates for items from the Learned category; H-Not Learned corresponds to the hit rates for items from the Not Learned category; H-Both corresponds to the hit rates for items of Both categories; and H-Neither corresponds to the hit rates for

items from Neither category. Error bars indicate SEM. **b** False alarm rates during the test phase. FA-Learned corresponds to the false alarm rates for lures of the Learned category; FA-Not Learned corresponds to the false alarm rate for lures of the Not Learned category; and FA-Neither corresponds to the false alarm rate for lures of Neither category.

for items from Neither category relative to both Not Learned, $t(99) = 3.34$, $p = .001$, $d = .42$, and Learned, $t(99) = 3.89$, $p < .001$, $d = .46$. However, no difference in C was found for items of the Learned versus Not Learned categories ($p > .05$).

Discussion

To explore whether the effects in FA rates from Experiment 1 were attributable to learning differences alone, rather than study or testing differences, Experiment 2 did not include a learning stage. As expected, there were no differences in FA for either list. However, FA rates for lures of both the Learned and Not Learned categories were higher than for lures of Neither category, suggesting that the higher frequency of some flowers' features during the study phase may still have an effect on FA during recognition even in the absence of explicit category learning. Similarly, the fact that we still find some bias for items from the Learned and the Not Learned lists relative to items from Neither category suggests that at least some of the participants' tendencies to respond "old" to items during the test phase may be due to features of the study and testing phases rather than learning alone. This observation may also explain why the hit rates of the learned categories in Experiment 1 were also higher than the hit rates for items of neither category—a pattern of results that was also replicated in Experiment 2, and in the replication of Experiment 1 (see footnote 1), albeit in neither case reaching statistical significance. We explore this issue in Experiment 4. But before that, we report the results from Experiment 3, in which we further test whether the effects found in Experiment 1 were indeed attributable to learning differences. To that end, Experiment 3 involves a more complex category learning phase, where flowers varied across seven rather than five dimensions. This increased level of difficulty resulted in only about half of the participants successfully learning the category, thus allowing us to compare recognition performance between

participants that completed all the learning trials and successfully learned the category versus participants that completed all the learning trials but failed to learn the category.

Experiment 3

Methods

Participants Two hundred individuals were recruited and run with Amazon Mechanical Turk (<https://www.mturk.com>). All participants were from the USA, gave informed consent, and received approximately US\$0.60 for participating.

Stimuli The stimuli were very similar to those used in Experiment 1 (Fig. 1), except that, for Experiment 3, 2,187 different flowers were generated in MATLAB. Flowers varied across seven possible dimensions, with each dimension taking up one of three possible values, as follows: (a) number of petals: 4, 6, or 8; (b) color of petals: red, blue, or green; (c) shape of petals: pointy, round, or concave; (d) shape of center: circle, triangle, or square; (e) shape of hole in the center: star, blob, or heart; (f) color of the center: yellow, pink, or turquoise; and (g) number of sepals: 0, 1, or 2. Flowers appeared in the center of the screen against a white background.

Procedure The procedure was the same as Experiment 1 except for the following differences: (a) The learning phase involved 100 as opposed to 60 trials. (b) These 100 trials were divided as follows: 25 displayed flowers from the Learned category, 25 displayed flowers from the Not Learned category, 25 displayed flowers from Both categories, and 25 displayed flowers from Neither category. Therefore, as in Experiment 1, 50% of the trials were hits, but of those only half were unique to either the Learned or the Not Learned category. (c) We considered participants to have successfully learned if they

classified at least 85 % of flowers correctly from trials 70 to 100. The rest of the procedure was identical to Experiment 1.

Results

Participants' performance in the learning task is depicted in Fig. 5a, and recognition data are depicted in Fig. 5b (Hits) and c (FA). Only 56.5 % of participants reached a successful learning criterion during the learning stage. Thus, recognition data were split between participants who successfully learned the category [Learning: $n = 113$] and those who did not [No Learning: $n = 87$]. To analyze Hit rates, a 2 (Group: Learning, No Learning) \times 4 (Category: Learned, Not Learned, Both, Neither) ANOVA was conducted, revealing an effect of Category, $F(3, 196) = 3.59, p = .015, \eta^2 = .05$, with no interaction, $p = .85$. Paired direct contrasts indicated that, regardless of group, hit rates for the Neither category were significantly lower than for Learned, $t(199) = 2.73, p = .007$, Not Learned, $t(199) = 2.00, p = .047$, and Both, $t(199) = 2.94, p = .004$. The Hit rate for the Not Learned category was also lower than for Both, $t(199) = 2.66, p = .009$, but there was no difference between Learned and Not Learned, $p = .325$, or between Learned and Both, $p = .158$.

To analyze FA rates, a 2 (Group: Learning, No Learning) \times 3 (Lure: FA-Learned, FA-Not Learned, FA-Neither) ANOVA revealed a main effect of Lure, $F(2, 198) = 10.98, p < .001, \eta^2 = .05$, modified by a Lure \times Group interaction, $F(2, 198) = 3.39, p = .035, \eta^2 = .02$. Direct contrasts revealed that only for participants in the Learning group, FA for lures of the Learned category ($M = .58, SD = .21$) were significantly higher than for lures of both Not Learned ($M = .50, SD = .22$), $t(224) = 3.0, p = .003, r = .19$, and Neither ($M = .47, SD = .18$), $t(224) = 4.52, p < .0001, r = .27$, categories. There were no significant differences in FA rates for participants in the No Learning group (all $p > .05$).

To assess differences in discriminability in recognition memory, two further analyses were conducted. First, differences in Hit minus FA rates for items from all three conditions were calculated for both groups (Learning group: Learned: $M_{[Hit-FA]} = .09, SD = .21$; Not Learned $M_{[Hit-FA]} = .16, SD = .24$; Neither: $M_{[Hit-FA]} = .15, SD = .27$; No Learning group: Learned: $M_{[Hit-FA]} = .15, SD = .23$; Not Learned $M_{[Hit-FA]} = .15, SD = .21$; Neither: $M_{[Hit-FA]} = .15, SD = .24$), and then a 2 (Group: Learning, No Learning) \times 3 (Condition: Learned, Not Learned, Neither) mixed-design ANOVA was performed. No effects were found ($p > .05$). Second, measures of d' and C were calculated for both groups (Learning group: Learned: $d' = .25, C = -.39$; Not Learned: $d' = .49, C = -.26$; Neither: $d' = .46, C = -.13$; No Learning group: Learned: $d' = .45, C = -.38$; Not Learned: $d' = .43, C = -.34$; Neither: $d' = .58, C = -.28$), and two 2 (Group) \times 3 (Condition) mixed-design ANOVAs comparing each measure across conditions were conducted. While no effects for d' were revealed ($p > .05$), for C there was a main effect of Condition, $F(2, 197) = 7.65, p = .001, \eta^2 = .07$, with no interaction. Follow-up paired t-tests indicated higher bias in recognition responses for items from the Learned relative to Not Learned, $t(199) = 2.17, p = .03, d = .16$, and Neither category, $t(199) = 4.06, p < .001, d = .46$. No other effects were revealed.

Discussion

The difficulty of the learning task naturally divided participants into two groups: participants who succeeded in the categorization task (Learning) and participants who did not (No Learning). Participants who succeeded in the categorization task, and successfully discriminated items during the learning phase on the basis of the correct category inclusion-feature, showed a higher FA rate to lures of the learned category




Fig. 5 Learning and recognition performance from Experiment 3. **a** Categorization performance during the learning phase for participants that successfully learned the category (darker line) versus those participants that did not (lighter line). **b** Hit rates during the test phase. The darker bars correspond to the hit rates of participants who successfully learned the category whereas the lighter bars correspond to

the hit rates of participants that did not learn the categories. **c** False alarm rates during the test phase. The darker bars correspond to the false alarm rates of the participants who studied items from category list A, while the lighter bars correspond to the recognition performance of the participants who studied items from category list B

relative to all other lures. This effect was not apparent for participants who did not succeed at the categorization task. These results replicate those of Experiment 1, and corroborate those of Experiment 2, insofar as the effects for FA rates were only evident if the learning phase was successful. The pattern of FA rates for participants who saw all the learning trials but failed to learn the category mirror the pattern of FA rates in Experiment 2, where there was no learning phase. Taken together, these results indicate that learning a category increases the chance of falsely recognizing a lure as old if it is an instance of the previously acquired knowledge structure.

For the hit rates, however, the results were somewhat different, as there was no interaction between group and category. However, the analysis did reveal that hit rates were lower for items of the Neither category, regardless of group. Likewise, the analysis conducted on measures of discriminability revealed higher bias to answering “old” in items from the Learned relative to Neither category in both groups, and a small but significant effect of bias for items of the Learned relative to the Not Learned category, likely driven by the group that successfully learned the category. These results agree with the pattern of results that was evident for hit rates in Experiment 1, where hit rates for items from the Neither category were also lower than from any of the other three categories. The effect does not appear, however, when there is no learning phase, as in Experiment 2. Interestingly, a similar pattern of results occurs for FA rates in Experiments 1 and 3 (but not 2): FA rates for Neither-lures are lower than for the lures from the Learned and the Not Learned categories.

What might account for this pattern of results? Part of the explanation may be that, by the time the participant faces the test phase after the learning and study phases, they would have seen the value that determines category inclusion more frequently than any other value. As such, if some features of the flowers are oversampled during the learning phase, then it may be possible that the lower rates of hits and FA for items from the Neither category are simply due to oversampling rather than category learning per se. Thus, in Experiment 4, we modified the current paradigm so that the frequency of values for each dimension is kept constant during learning and encoding to assess whether the same hit rate and false alarm rate effects are found after equating frequency.

Experiment 4

Methods

Participants One hundred and five individuals were recruited and run with Amazon Mechanical Turk (<https://www.mturk.com>). All participants were from the USA, gave informed

consent, and received approximately US\$1.80 for participating.

Stimuli MATLAB generated flowers from the previous experiments were employed. Flowers varied across five possible dimensions with each dimension taking up one of three possible values, as follows: (a) number of petals: 4, 6, or 8; (b) color of petals: yellow, blue, or green; (c) shape of center: circle, triangle, or square; (d) color of the center: purple, orange, or turquoise; and (e) number of sepals: 1, 2, or 3. Flowers appeared in the center of the screen against a white background.

Procedure The procedure was similar to the previous experiments except for the following differences. The learning phase involved 54 trials. Each participant was assigned a learning rule consisting of one possible value (e.g., yellow) from one of the five possible dimensions (e.g., petals). This assignment was counterbalanced across participants. Flowers that belonged to the Learned category appeared on one-third of the trials—i.e., 18 trials. The other two-thirds of the trials included flowers displaying the other two values of the same dimension as the category-inclusion value. Half of these trials (18 trials) featured the second value (e.g., blue petals) while the other half (18 trials) featured the third value (e.g., green petals). Thus, the three values of the category relevant dimension were equally sampled during the learning stage. A value from a different dimension (e.g., one sepal) was assigned as the Not Learned category. This assignment was also counterbalanced across participants. The value that determined the Not Learned category was also displayed in 18 trials, half of which co-occurred with the value that determined the Learned category, so there was a total of nine flowers belonging to Both categories. The other two values of the dimension that determined the Not Learned category were displayed in equal frequencies of 18 trials each. As such, the three values of the category inclusion dimensions of the Learned and Not Learned categories were presented with equal frequency, 18 times each, nine of which co-occurred. Thus, contrary to Experiments 1 and 3, where the category inclusion value from both the Learned and the Not Learned categories occurred more frequently than the other values from the same dimensions, in Experiment 4 the frequency of the three possible values from the two dimensions that determined Learned and Not Learned flowers occurred with equal frequency. The other values of the remaining three dimensions were chosen pseudo-randomly, so that each of them co-occurred with each of the three values of the first two dimensions six times, resulting in identical frequencies for all of them. Participants were considered to have successfully learned the category if they classified correctly at least 80 % of the last 20 trials. (Table 2 displays the abstract category structure of the learning phase. Figure 6a displays a frequency

Table 2 Abstract category structure used in the learning stage of Experiment 4. Each row represents a unique stimulus (Flower). Each number represents one of three possible values for each of the five dimensions. Here, the first value of the first dimension (1,1) corresponds to the category inclusion value for the Learned category. The first value of the second dimension (2,1) corresponds to the category-inclusion value for the Not Learned category. Half of the category inclusion features for Not Learned-Flowers co-occurred with half of the Learned-Flowers to generate Both-Flowers. The values of the remaining three dimensions were selected pseudo-randomly, so that they covaried with equal frequencies with the values of each of the first two dimensions

Category	Stimulus	Dimension value				
		1	2	3	4	5
Learned	B1	1	1	1	1	3
	B2	1	1	1	2	1
	B3	1	1	1	2	2
	B4	1	1	1	3	3
	B5	1	1	2	3	1
	B6	1	1	2	3	2
	B7	1	1	3	1	1
	B8	1	1	3	2	1
	B9	1	1	3	2	3
	L1	1	2	1	1	3
	L2	1	2	2	1	2
	L3	1	2	2	2	1
	L4	1	2	3	3	1
	L5	1	2	3	3	3
	L6	1	3	1	1	3
	L7	1	3	2	1	2
	L8	1	3	2	2	2
	L9	1	3	3	3	2
Not Learned	NL1	2	1	1	1	3
	NL2	2	1	2	2	2
	NL3	2	1	3	1	3
	NL4	2	1	3	3	1
	NL5	3	1	2	1	3
	NL6	3	1	2	3	1
	NL7	3	1	2	3	3
	NL8	3	1	3	1	2
	NL9	3	1	3	2	2
	N1	2	2	1	1	1
	N2	2	2	1	3	1
	N3	2	2	1	3	2
	N4	2	2	2	1	1
	N5	2	2	2	3	3
	N6	2	2	3	1	2
	N7	2	2	3	2	3
	N8	2	3	1	2	2
	N9	2	3	1	3	2
N10	2	3	2	1	2	
N11	2	3	2	2	1	

Table 2 (continued)

Category	Stimulus	Dimension value				
		1	2	3	4	5
	N12	2	3	2	2	3
	N13	2	3	3	2	1
	N14	2	3	3	3	3
	N15	3	2	1	1	3
	N16	3	2	1	2	3
	N17	3	2	1	3	1
	N18	3	2	1	3	3
	N19	3	2	2	2	2
	N20	3	2	3	1	2
	N21	3	2	3	2	3
	N22	3	3	1	2	1
	N23	3	3	1	3	2
	N24	3	3	2	1	1
	N25	3	3	2	1	2
	N26	3	3	3	2	1
	N27	3	3	3	3	1

distribution matrix of the average pairwise co-occurrence of each dimension’s values across trials during the learning phase.)

To further control for the frequency of presentations of the critical values for flowers of the Learned and Not Learned categories during the study phase, the encoding session included 24 rather than 20 items, so each value of each relevant category was featured eight times. In total, the study phase included four Learned-Flowers, four Not Learned-Flowers, four Both-Flowers, and eight Neither-Flowers. (Figure 6b displays a frequency distribution matrix of the average pairwise co-occurrence of each dimension’s values across trials during the study phase.) Finally, the retrieval phase was identical to the previous ones.

Results

Participants’ performance in the learning task is depicted in Fig. 7a, and recognition data are depicted in Fig. 7b (Hits) and c (FA). An ANOVA was employed to compare hit rates for old items from all four categories (Learned, Not Learned, Both, Neither). No effect was found, $p = .779$. An ANOVA comparing FA rates for Lures from the Learned, Not Learned and Neither categories revealed a main effect of Lure, $F(2, 103) = 5.45, p = .006, \eta^2 = .10$. Paired contrasts revealed that FA rates for lures from the Learned category ($M = .59, SD = .23$) were higher than for lures of the Not Learned ($M = .52, SD = .23$), $t(104) = 2.93, p = .004, d = .30$, as well as the Neither categories ($M = .53, SD = .22$), $t(104) = 2.68, p = .009, d = .27$.

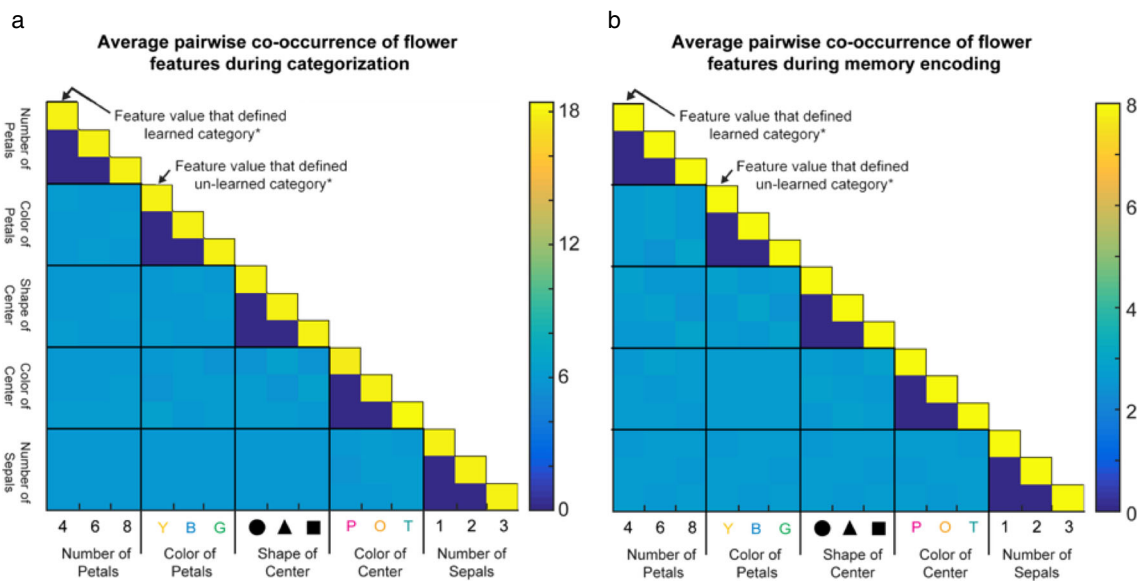


Fig. 6 Frequency distribution matrix depicting average pairwise co-occurrence of flower’s features during the learning phase (a) and the study phase (b) in Experiment 4. The groups of columns correspond to the five dimensions across which flowers varied, and the columns within them correspond to the particular values along those dimensions. Colors

indicate the frequency of presentation of that value for that dimension across trials. As in Fig. 2, the first dimension corresponds to the learned rule dimension, and the leftmost value corresponds to the learned rule value. * Which feature and value defined each category was counterbalanced across participants

However, there was no difference in FA rates for lures from the Not Learned and the Neither categories, $p = .75$.

As in Experiments 1 and 2, discriminability in recognition memory was assessed, first, by calculating differences in Hit minus FA rates for items from all three conditions (Learned: $M_{[Hit-FA]} = .03$, $SD = .34$; Not Learned $M_{[Hit-FA]} = .11$, $SD = .32$; Neither: $M_{[Hit-FA]} = .09$, $SD = .39$), and performing an ANOVA across the three measures. No effects were found ($p > .05$). Second, measures of d' and C were calculated (Learned: $d' = .14$, $C = -.37$; Not Learned: $d' = .40$, $C = -.25$; Neither: $d' = .38$, $C = -.28$), and ANOVAs comparing them across conditions revealed no effects of sensitivity or bias (both $p > .05$).

Discussion

Given the structure of the category learning phases in Experiments 1 and 3, the values of the category relevant dimension were more frequently presented than the non-category relevant values along the same dimension. As a result, oversampling the value of a dimension during learning may have had the effect of increasing both true and false recognition rates. This may explain why both hit rates and FA rates for items from the Neither category were lower than for items of any other category in Experiments 1 and 3. Moreover, the oversampled features during encoding in the absence of a learning stage may have affected the FA rate of




Fig. 7 Learning and recognition performance from Experiment 4. **a** Categorization performance during the learning phase. **b** Hit rates during the test phase. H-Learned corresponds to the hit rates for items from the Learned category; H-Not Learned corresponds to the hit rates for items from the Not Learned category; H-Both corresponds to the hit rates for items of Both categories; and H-Neither corresponds to the hit rates for

items from Neither category. Error bars indicate SEM. **c** False alarm rates during the test phase. FA-Learned corresponds to the false alarm rates for lures of the Learned category; FA-Not Learned corresponds to the false alarm rate for lures of the Not Learned category; and FA-Neither corresponds to the false alarm rate for lures of Neither category

lures from the Neither category, whose features were relatively under-sampled, in Experiment 2. The results from Experiment 4, however, suggest that the category-learning effects found in Experiments 1 and 3 may not be merely due to oversampling, as the increased false alarm rate for lures of the learned category remained despite controlling for frequency of occurrence during learning and encoding. After equating the oversampling of category-relevant items during learning and encoding, the difference between hit rates across conditions, as well as the difference between FA rates for lures of the non-learned and the neither categories, is no longer present.

General discussion

The current paper reports the results of four experiments employing variations on a new category-learning paradigm that permits the manipulation of different features of the learning, encoding, and retrieval stages in order to study the effects of acquiring a novel knowledge structure or schema on recognition memory. The results from Experiment 1 indicate that learning to categorize items as belonging to a particular category reliably increases the hit rates for category-consistent “old” items as well as the false alarm rates for category-consistent “new” lures. To complement this observation, the results of the discriminability analysis also suggest that participants were biased toward answering “old” to items of the learned category more so than to items from the not learned and neither categories. Importantly, these results are independent of the particular perceptual feature used as criterion for category inclusion. The results from Experiment 2 further indicated that, although the lower hit and false alarm rates for items from Neither category might have been due to the features of the items in the study phase, the differences in hit and false alarm rates for items from the Learned versus non Learned categories found in Experiment 1 were actually dependent upon category learning alone. Indeed, the fact that participants were equally biased toward answering “old” to items from the lists of the Learned and the Not Learned categories relative to items from Neither category further suggests that oversampling features during the study phase is enough to affect recognition at test.

However, the observation that the differences in hit and FA rates for items from the Learned relative to Not Learned category is due to learning, is further strengthened by the results from Experiment 3, which did include a learning stage but a more complex one, so that only about half of the participants learned the category rule. The effects of category learning found in Experiment 1 were only evident in participants that successfully learned the complex category learning, but were absent in those participants that failed to learn the task, despite the fact that participants from both groups were exposed to the

same items during the learning phase. This pattern suggests that the mere exposure to a series of items is not sufficient to generate the hit and false alarm increases during recognition, because it is necessary to experience such items as forming a coherent knowledge structure or schema.

These results parallel previous findings on recognition memory from the schematic knowledge literature. Smith, Ward, Tindell, Sinfonis and Wilkenfeld (2000), for instance, reported higher rates of false recognition to category-consistent lures during a memory test whose study phase included words that were semantically associated by typicality. More recently, Castel et al. (2007) showed that individuals with expertise in football were more likely to falsely recognize football-consistent lures relative to non-experts, but also to correctly identify football-consistent names from a study list. However, all previously documented cases of increases in hit and false alarm rates to schema-consistent as opposed to schema-inconsistent items during a memory recognition test have included subjects whose schematic knowledge was already learned, making it difficult to track the development from schematic knowledge acquisition to encoding and subsequent recognition. The paradigm employed in the current study allows investigation of all three stages, verifying the effect of structured knowledge acquisition on subsequent memory performance. Moreover, given the parallelisms in the results obtained in the first three experiments, and those reported in the schematic knowledge tradition, we believe that this paradigm, as well as possible variations thereof, could contribute to further integrating these two lines of research.

Indeed, Experiment 4 aimed to do precisely that. A consistent result in Experiments 1, 2, and 3 is that items that did not belong to either the learned or the non-learned categories had lower hit and false alarm rates than items that either belonged to the learned or to the non-learned category. A possible explanation for this result stems from the fact that, by the time the participant faces the test phase, the flowers’ features that determine category membership either to the learned or the non-learned category have been seen more frequently than any other value. If one assumes that the probability of endorsing a category-consistent item as old during a recognition task is proportional to the probability function determined by the frequency of prior occurrences of category-consistent stimuli, then it should be expected that by equating the frequency of occurrence of items from the learned, non-learned and neither categories, the difference in hit and false alarm rates should go away (De Brigard, 2012).

The modified version of the paradigm employed in Experiment 4 tested this hypothesis, and found that although the differences between hit rates across conditions as well as false alarm rates between lures of the non-learned category and the neither category were eliminated, the main effect of category learning on false alarm rates for category-consistent lures remained. This finding suggests that the effect of

category learning on false alarms is not only due to oversampling, and that acquiring a knowledge structure can influence false alarm rates to lures that fall within the scope of the acquired schema even when frequency of exposure has been controlled for during learning and encoding. What could account for these results? One possibility is that attentional mechanisms may highlight the relevant value and/or dimension during learning and encoding, which in turn may increase the evidentiary weight given to those items relative to other items that have been presented with equal frequency (e.g., Love et al., 2004). Further research is needed to evaluate the precise mechanisms responsible for this effect above and beyond frequency sampling. Variations of the current paradigm may be well-suited for such a purpose.

Another advantage of the current paradigm is that it affords the possibility of testing the effect of different category learning strategies on subsequent recognition tests. Previous research has shown that there are critical differences in the way people learn perceptual categories (Ashby & Maddox, 2005; Richler & Palmeri, 2014). However, the influence of different category-learning strategies on memory recognition is an issue that remains understudied (but see Clapper, 2008; Sakamoto & Love, 2010). The paradigm employed in the current study offers the possibility of varying the category learning strategy employed in the learning phase. In the version of the paradigm employed here, we used a A/Not-A rule based task, but it could easily be modified to fit other rule-based tasks (e.g. A/B), as well as tasks that include more or less complex rules (Ashby & Maddox, 2005; Feldman, 2000), which could be either continuous or categorical (Hemmer and Steyvers, 2009; Huttenlocher, Hedges, & Duncan, 1991; Huttenlocher, Hedges, & Vevea, 2000; Steyvers and Hemmer, 2012). The learning phase could also be modified to include other kinds of category-learning strategies—some of which are more familiar to the schematic-knowledge literature in memory—such as prototype-distortion (Koutstaal, Verfaellie, & Schacter, 2001; Posner & Keele, 1968; Shin & Nosofsky, 1992; Cabeza, Bruce, Kato, & Oda, 1999; Schacter, Verfaellie, & Koutstaal, 2002; Slotnick & Schacter, 2004; Smith & Minda, 2002), supervised versus unsupervised learning (Clapper, 2008; Clapper & Bower, 1991, 1994), information-integration (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Smith & Minda, 1998) and weather-prediction tasks (Eldridge, Masterman, & Knowlton, 2002; Reber & Squire, 1999).

Further advantages of the paradigm employed here is that it allows us to modify the test phase to explore, for instance, recognition differences for whole items versus items' features as a function of category learning. Unlike the experiments reported here, earlier studies on schematic influences on recognition test participant's memory for individual features, parts or statements of a general schema (e.g., Bower, et al.,

1979). Here, however, participants were asked not to recognize a particular feature of an item—e.g., red petal—but rather a particular item as a whole. Future studies could easily explore these potential differences by incorporating changes in the testing phase. Similarly, future studies could also include lures from Both categories to more fully explore differences in sensitivity and bias with a signal detection modeling approach.²

Additionally, we believe that by allowing the category-learning strategies to vary while tightly controlling for the subsequent study and recognition stages, the current paradigm can potentially contribute to a more recent debate regarding the involvement of multiple memory systems in category learning (Ashby & O'Brien, 2005). Early views on category learning assumed that the acquisition of new categorical knowledge was mediated by a single memory system, and that regardless of the category-learning strategy the same kind of category representation was generated at the end. Recent behavioral, neuropsychological, and neuroimaging evidence, however, suggest that different category learning strategies are subserved by different memory systems. It has been shown, for instance, that individuals with amnesia due to temporal lobe damage are impaired in learning categories using an explicit rule-based strategy, but perform normally in weather prediction tasks (Gluck, Oliver, & Myers, 1996; Knowlton, Squire, & Gluck, 1994). In contrast, individuals with striatal damage due to Parkinson's or Huntington's disease show profound learning deficits in weather prediction tasks, but perform on par with healthy controls in explicit-rule based tasks (Knowlton, Squire, Paulsen, Swerdlow, et al., 1996). These findings suggest that while the mechanisms involved in declarative memory systems are essential for category learning via explicit rule-based strategies, they may not be critical for category learning via probabilistic strategies, such as in the weather prediction task, which tap instead into memory systems associated with procedural memory. The paradigm used in the current study, as well as the variations in the learning stage it affords, could potentially contribute to understanding differences in category-learning for hit and false alarm rates as a function of the different memory systems involved. Finally—and relatedly—we believe that the results reported in the present manuscript, as well as the variations afforded by the current paradigm, may contribute to understanding whether the processes of acquiring a schema and of learning a category may actually reflect the same cognitive mechanisms (e.g., Davis et al., 2014).

Acknowledgments The authors thank Roberto Cabeza for useful discussion. This research was supported by a grant from National Institute on Aging (NIA AG08441) awarded to D.L.S.

² We thank an anonymous reviewer for these useful suggestions.

References

- Arkes, H. R., & Freedman, M. R. (1984). A demonstration of the costs and benefits of expertise in recognition memory. *Memory & Cognition*, *12*, 84–89.
- 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.
- Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–178.
- Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Science*, *2*, 83–89.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press.
- Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, *11*, 177–220.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, *138*(4), 487–502.
- Brainerd, C. J., & Reyna, V. F. (1990). Gist is the grist: Fuzzy-trace theory and the new intuitionism. *Developmental Review*, *10*, 3–47.
- Brainerd, C. J., & Reyna, V. F. (2005). *The science of false memory*. New York: Oxford University Press.
- Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive Psychology*, *13*, 207–230.
- Cabeza, R., Bruce, V., Kato, T., & Oda, M. (1999). Prototype effect in face recognition: Extension and limits. *Memory and Cognition*, *27*, 139–151.
- Castel, A. D., McCabe, D. P., Roediger, H. L., III, & Heitman, J. L. (2007). The dark side of expertise: Domain specific memory errors. *Psychological Science*, *18*, 3–5.
- Chase, W. G., & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing* (pp. 215–281). New York: Academic Press.
- Clapper, J. P. (2008). Category learning as schema induction. In M. A. Gluck, J. R. Anderson, & S. M. Kosslyn (Eds.), *Memory and mind: A Festschrift for Gordon H. Bower*. New Jersey: Lawrence Erlbaum Associates.
- Clapper, J. P., & Bower, G. H. (1991). Learning and applying category knowledge in unsupervised domains. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 27, pp. 65–108). New York: Academic Press.
- Clapper, J. P., & Bower, G. H. (1994). Category invention in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 43–460.
- Davis, T., Xue, G., Love, B. C., Preston, A. R., & Poldrack, R. A. (2014). Global neural pattern similarity as a common basis for categorization and recognition memory. *Journal of Neuroscience*, *34*(22), 7472–7484.
- De Brigard, F. (2012). Predictive memory and the surprising gap. Commentary on Andy Clark's "Whatever Next? predictive brains, situated agents and the future of cognitive science". *Frontiers in Psychology*, *3*, 420.
- de Groot, A. D. (1966). Perception and memory versus thought: Some old ideas and recent findings. In B. Kleinmuntz (Ed.), *Problem solving* (pp. 19–50). New York: Wiley.
- Eldridge, L. L., Masterman, D., & Knowlton, B. J. (2002). Intact implicit habit learning in Alzheimer's disease. *Behavioral Neuroscience*, *116*, 722–726.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*(6804), 630–633.
- Gluck, M. A., Oliver, L. M., & Myers, C. E. (1996). Late-training amnesic deficits in probabilistic category learning: A neurocomputational analysis. *Learning and Memory*, *3*, 326–340.
- Graesser, A. C., & Nakamura, G. V. (1982). The impact of a schema on comprehension and memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 16, pp. 59–109). New York: Academic Press.
- Hemmer, P., & Steyvers, M. (2009). A Bayesian account of reconstructive memory. *Topics in Cognitive Science*, *1*, 189–202.
- Hollingworth, A., & Henderson, J. M. (2003). Testing a conceptual locus for the inconsistent object change detection advantage in real-world scenes. *Memory & Cognition*, *31*, 930–940.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, *129*, 220–241.
- Knowlton, B., Squire, L., Paulsen, J., Swedlow, N., Swenson, M., & Butters, N. (1996). Dissociations within nondeclarative memory in Huntington's disease. *Neuropsychology*, *10*, 538–548.
- Knowlton, B., Squire, L., & Gluck, M. (1994). Probabilistic classification learning in amnesia. *Learning and Memory*, *9*, 408–418.
- Koutstaal, W., Verfaellie, M., & Schacter, D. L. (2001). Recognizing identical vs. similar categorically related common objects: Further evidence for degraded gist-representations in amnesia. *Neuropsychology*, *15*, 268–289.
- Lampinen, J. M., Copeland, S. M., & Neuschatz, J. S. (2001). Recollections of things schematic: Room schemas revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1211–1222.
- Love, B. (2013). Categorization. In K. N. Ochsner & S. M. Kosslyn (Eds.), *Oxford handbook of cognitive neuroscience* (pp. 342–358). Oxford: Oxford Press.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*(2), 309–332.
- Minsky, M. (1975). A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 211–277). New York: McGraw-Hill.
- Oliva, A. (2005). Gist of the scene. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *The encyclopedia of neurobiology of attention* (pp. 251–256). San Diego: Elsevier.
- Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 548–568.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.
- Reber, P. J., & Squire, L. R. (1999). Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behavioral Neuroscience*, *113*, 235–242.
- Richler, J. J., & Palmeri, T. J. (2014). Visual category learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*, 75–94.
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 803–814.
- Rojahn, K., & Pettigrew, T. F. (1992). Memory for schema-relevant information: A meta-analytic resolution. *British Journal of Social Psychology*, *31*, 81–109.
- Sakamoto, Y. (2012). Schematic influences on category learning and recognition memory. N. M. Seel (Ed.). *Encyclopedia of the Sciences of Learning*. Springer.
- Sakamoto, Y., & Love, B. C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General*, *133*, 534–553.
- Sakamoto, Y., & Love, B. C. (2010). Learning and retention through predictive inference and classification. *Journal of Experimental Psychology: Applied*, *16*, 361–377.

- Schacter, D. L., Verfaellie, M., & Koutstaal, W. (2002). Memory illusions in amnesic patients: Findings and implications. In L. R. Squire & D. L. Schacter (Eds.), *Neuropsychology of memory* (3rd ed., pp. 114–129). New York: Guilford Press.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding*. Hillsdale: Erlbaum.
- Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General*, *121*(3), 278–304.
- Slotnick, S. D., & Schacter, D. L. (2004). A sensory signature that distinguishes true from false memories. *Nature Neuroscience*, *7*, 664–672.
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1411–1430.
- Smith, J. D., & Minda, J. P. (2002). Distinguishing prototype-based and exemplar-based processes in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 800–811.
- Smith, S. M., Ward, T. B., Tindell, D. R., Sifonis, C. M., & Wilkenfeld, M. J. (2000). Category structure and created memories. *Memory & Cognition*, *28*, 386–395.
- Steyvers, M., & Hemmer, P. (2012). Reconstruction from Memory in Naturalistic Environments. In B. H. Ross (Ed), *The Psychology of Learning and Motivation*, (126–144). Elsevier Publishing
- Taylor, S. E., & Crocker, J. (1981). Schematic bases of social information processing. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The ontario symposium* (Vol. 1, pp. 89–134). Hillsdale: Erlbaum.
- van Kesteren, M. T., Ruitter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences*, *35*, 211–219.