



Learning concepts when instances never repeat

Donald Homa¹ · Mark Blair² · Samuel M. McClure¹ · John Medema¹ · Gregory Stone¹

Published online: 12 November 2018
© Psychonomic Society, Inc. 2018

Abstract

Three experiments explored the learning of categories where the training instances either repeated in each training block or appeared only once during the entire learning phase, followed by a classification transfer (Experiment 1) or a recognition transfer test (Experiments 2 and 3). Subjects received training instances from either two (Experiment 2) or three categories (Experiments 1–3) for either 15 or 20 training blocks. The results showed substantial learning in each experiment, with the notable result that learning was not slowed in the non-repeating condition in any of the three experiments. Furthermore, subsequent transfer was marginally better in the non-repeating condition. The recognition results showed that subjects in the repeat condition had substantial memory for the training instances, whereas subjects in the non-repeat condition had no measurable memory for the training instances, as measured either by hit and false-alarm rates or by signal detectability measures. These outcomes are consistent with prototype models of category learning, at least when patterns never repeat in learning, and place severe constraints on exemplar views that posit transfer mechanisms to stored individual traces. A formal model, which incorporates changing similarity relationships during learning, was shown to explain the major results.

Introduction

The empirical study of concepts has a long history, beginning with the introspective analyses of Moore (1910) and Fisher (1916), and soon followed by the analytical methodology of Hull (1920). Hull's paradigm, in which variables were introduced in a learning phase followed by a transfer test, has since become the paradigm of choice. For the purposes of the present study, the interesting aspect of Hull's paradigm, and virtually all studies since, is that the learning set is typically recycled until some learning criterion is met, such as errorless performance or a predetermined number of learning blocks has occurred.

This paradigm has provided important insights into the formation and representation of concepts, including the identification of learning variables that shape concepts (e.g., Homa, 1984; Wills & Pothos, 2012) as well as the development of formal, quantitative models of classification and category learning (e.g., Busemeyer & Diederich, 2009; Hintzman, 1986; Nosofsky, 1988; Minda and Smith, 2001). Nonetheless, when

learning many real-world categories, for example birds or faces or cars, only a tiny portion of the instances are ever exactly repeated. Even individual exemplars vary subtly across time (e.g., faces appear different depending on view angle or expression) or different instances of a category have individuating characteristics (e.g., every red Toyota Prius differs in license plate, cleanliness, and decoration). The exact repetition of instances is more likely to fall within the domain of formalized training or instruction. The category learning paradigm, in which exact instances are repeated throughout the learning phase, gives little insight into what kind of performance to expect when learning non-repeating stimulus sets.

A handful of studies have explored category learning when instances do not repeat in the training phase. Medin, Dewey, and Murphy (1983) compared the acquisition of two face categories where the faces were either repeated on each trial block or not. Only the terminal levels of learning were reported, and subjects in the non-repeat condition found the task quite difficult (only 28% of the subjects reached errorless criterion). Learning was probably slowed because the dimensions relevant for classification – hair and shirt color, hair length, and whether a smile was present or not – were unrelated to natural grouping based on facial properties, and classification was solely determined by their value on facially-irrelevant dimensions. Ashby and his colleagues (e.g., Ashby & Gott, 1988; Ashby & Maddox, 1990, 1992; Casale, Roeder, & Ashby, 2012) have explored categories in which patterns are randomly sampled from a bivariate normal distribution, thereby generating a

✉ Donald Homa
donhoma@asu.edu

¹ Department of Psychology, Arizona State University, Tempe, AZ 85287, USA

² Department of Psychology, Simon Fraser University, Burnaby, BC, Canada

virtual limitless number of exemplars in learning. However, their task differs from the experiments reported here in a number of ways: (1) subjects typically learned two categories; (2) stimuli were composed of stimuli that varied along two readily-identifiable dimensions; and (3) learning and transfer were not contrasted for categories composed of repeated patterns versus the learning of categories whose patterns never repeat¹.

A different approach was taken by Knowlton and Squire (1993) and Reed, Squire, Patalano, Smith, and Jonides (1999), who presented 40 patterns one time each from a single category to normal and amnesic patients. Their major focus was whether results based on subsequent classification and recognition could be explained by single- or multiple-memory systems. Disputes over the interpretation of transfer results have subsequently been provided by Nosofsky and Zaki (1998) and Zaki and Nosofsky (2001). However, the utility of the single category paradigm was questioned by Palmeri and Flanery (1999), who demonstrated that subsequent classification on the transfer test with this paradigm could be explained by factors unrelated to learning. More recently, Homa, Hout, Milikan and Milikan (2011) found that minimal category knowledge was acquired following the observational phase of a single category. Because degree of category knowledge is enhanced by the training on more, rather than fewer, categories (Homa & Chambliss, 1975), presumably because the additional categories in the learning set provide the subject with information about distinctive category cues, it is doubtful that subjects learn much about a category from exposure to the members of a single category.

In the present study, we directly contrasted the learning of multiple categories where the learning patterns were either repeated or not on each trial block. For subjects in the non-repeating condition, the learning patterns were replaced on each trial block with novel patterns of the same level of distortion from the same prototypes. Subsequent transfer included either recognition or classification tests. Our initial expectation was that categories could be learned when training patterns never repeated, but that this procedure would likely slow the rate and possibly the terminal level of learning, at least compared to the repetition condition.

Indeed, exemplar-based models of classification (e.g., Nosofsky, 1988; Nosofsky & Johansen, 2000; Shin & Nosofsky, 1992) clearly predict a faster rate of learning when

patterns repeat versus a condition where the patterns never repeat (or repeat less frequently). This prediction occurs because classification, either in learning or transfer, is based on the summed evidence favoring a category, and similarity of a pattern to itself is greater than the similarity of a pattern to any other pattern of that category. This self-similarity ensures that the summed similarity to a category always favors more rapid learning when patterns repeat in learning.

To illustrate this, assume that the subject assigns a pattern into a particular learning category based on its similarity to the members of that category, relative to members of the contrasting categories. As is typically done, assume that similarity is a monotonic function of its distance to other category instances in psychological space, where the psychological space is derived via multidimensional scaling (e.g., Homa, Dunbar, & Nohre, 1991; Nosofsky & Zaki, 1998; Shin & Nosofsky, 1992). Let the similarity between patterns *i* and *j* be defined as:

$$s(i, j) = \exp(-cd_{ij}) \quad (1)$$

where d_{ij} is the distance between items *i* and *j*. The parameter *c* functions as a scaling (sensitivity) parameter that reflects how well the category members are discriminated from each other. To formalize the learning algorithm used in the present study, first consider how learning was predicted by Nosofsky and Zaki (1998). The probability that pattern *i* is classified into category A rather than category B is given by:

$$P(A|i) = \frac{[\sum s(i, a) + \beta]^\gamma}{[\sum s(i, a) + \beta]^\gamma + [\sum s(i, b) + \beta]^\gamma} \quad (2)$$

The summed similarities of instance *i* to patterns in category A and B are represented by $\sum s(i, a)$ and $\sum s(i, b)$, respectively. The parameter β is background noise, and γ is a response-scaling parameter (e.g., the subject evaluates the evidence for each category and assigns the stimulus to the category based on probability matching when $\gamma = 1$ and to the category more deterministically when γ exceeds 1). Because summed similarities cumulate across learning blocks, and because the background noise is constant, learning improves across trials. Logically, *c* should be greater when learning patterns repeat during learning, since repetition across learning blocks should make these patterns more discriminable relative to patterns that never repeat.

This learning algorithm was modified in the present study to reflect the predictions of an exemplar model involving multiple categories (e.g., Homa, Powell, & Ferguson, 2014; Nosofsky & Johansen, 2000), where learning instances either repeated (REP) on each block or not (NREP). For the REP condition, each category was represented by five different patterns, and the subject was required to learn three categories (A, B, C). On each trial block, these 15 patterns were randomly presented, with subjects receiving either 20 (Experiment 1) or 15 (Experiments 2 and 3) trial blocks. When a pattern from set A is presented on a learning trial, this item has maximal similarity to itself, a within-category

¹ At least in the biological domain and except for close family members, the family pet(s), etc., the vast bulk of encountered category instances are unique, e.g., forest creatures, birds of prey, lake fish, etc., rarely make return appearances. The novelty of experiences is further magnified by temporal changes that occur moment by moment (e.g., viewing angle and distance, facial expressions, bodily positions) and by maturation across longer time periods (Do, Homa, & Koehler, 2014), what Bruner, Goodnow, and Austin (1956) called identity categories. This is not to dismiss the salience of a critical instance that is repeatedly encountered or images in books that may be re-visited, although the importance of instance frequency, at least in experimental studies, has been shown to lose its impact on transfer with increases in category size (Homa et al., 1991). Regardless, the present study focuses on a category learning paradigm that uses reasonably complex stimuli that never repeat during the learning phase, a paradigm largely ignored in the literature.

similarity to the remaining four patterns of that category, and a between-category similarity to the five patterns in each of the other two categories. A simplifying assumption that does not alter the predicted learning rate for the REP and NREP condition is to assume that the subject performs randomly on the initial learning block. Beginning with the second trial block and thereafter, classification is based on stored exemplar knowledge. Therefore, when a pattern is presented for learning on trial block N , all learning patterns have been encountered $(N-1)$ previous times. To generate specific learning functions, we substituted multidimensional distances (Kruskal, 1964; Shepard, 1962) from a study that approximated the categories and constraints used in the present study (Home, Proulx, Blair, 2008). The particular distances are less critical for the purposes of this illustration than the constraint that similarity decreases with distance in the multidimensional space. The MDS distances between two medium patterns from the same category were 0.72, with this distance equal to 1.60 for patterns belonging to different categories. Equation 2 can be rewritten, beginning with trial block $B = 2$, as:

$$P(A|i, B) = \frac{[(B-1)(e^{-0c} + 4e^{-.72c}) + \beta]^\gamma}{[(B-1)(e^{-0c} + 4e^{-.72c}) + \beta]^\gamma + [(B-1)10e^{-1.6c} + \beta]^\gamma}$$

$$= \frac{[(B-1)(1 + 4e^{-.72c}) + \beta]^\gamma}{[(B-1)(1 + 4e^{-.72c}) + \beta]^\gamma + [(B-1)10e^{-1.6c} + \beta]^\gamma}$$
(3)

For the NREP condition, the number of different patterns supplants pattern repetition across learning blocks. The corresponding equation for the NREP condition is, therefore:

$$P(A|i, B) = \frac{[(B-1)5e^{-.72c} + \beta]^\gamma}{[(B-1)5e^{-.72c} + \beta]^\gamma + [(B-1)10e^{-1.6c} + \beta]^\gamma}$$
(4)

The only difference between equations (3) and (4) is that exemplar similarity to itself occurs in the REP equation and not in the NREP equation, and the frequency of a pattern in the REP condition is swapped with an equivalent number of different patterns in the NREP condition, although the number of category exposures in the REP and NREP conditions is the same.

Figure 1 shows predictions for these two equations, based on variations of c and γ . Learning improves across trial blocks in each case, with the REP condition always exceeding the performance of the NREP version with comparable parameter values. The predicted difference between REP and NREP conditions is probably underestimated, because c should be greater in the REP condition because pattern discriminability should be higher whenever patterns are repeated in learning, as is the case in the REP condition.² Since learning rate is

² In general, the parameter c is assumed to reflect pattern discriminability, with higher values associated with enhanced discriminability. The prediction that the parameter c should be larger with increasing frequency in learning has been reported previously, e.g., Shin and Nosofsky (1992). Furthermore, c is decreased with delay of test (Zaki & Nosofsky, 2001), indicating that pattern discriminability is reduced with the passage of time.

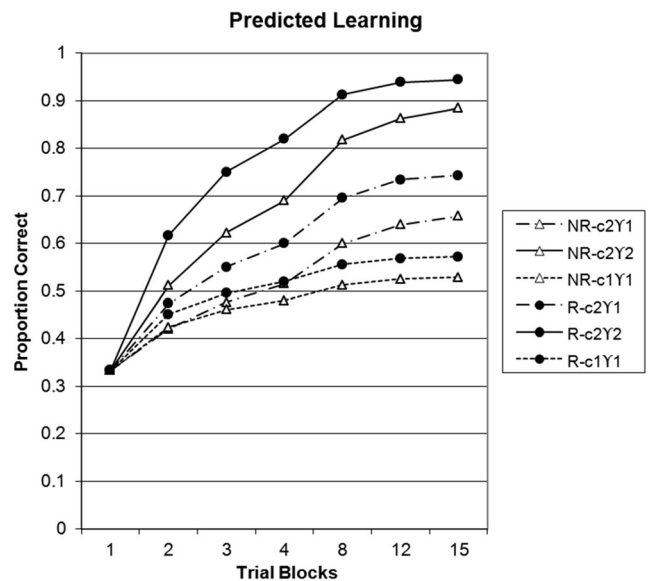


Fig. 1 Predicted learning rates by exemplar model for repeating (R) and non-repeating (NR) conditions, based on variations of c and γ

increased with increasing values of c (e.g., the learning rate in the REP condition with $c = 2$, $\gamma = 1$ is greater than the learning rate with $c = 1$, $\gamma = 1$), differences between learning rates in the REP and NREP condition should be greater as well.

An alternative view is that category learning and subsequent classification of novel instances on a transfer test is based on similarity to a category prototype (Homa, Proulx, & Blair, 2008; Smith & Minda, 1998, 2002). The non-repeating condition has a far larger category size compared to the repeating condition, and increases in category size are known to enhance subsequent generalization (e.g., Homa, Goldman, Cornell, & Cross, 1973; Homa & Vosburgh, 1976). However, the benefits of category size have typically been demonstrated for subsequent transfer following learning, rather than learning itself. As a consequence, the learning rate predicted for the REP and NREP conditions, based on a prototype influence, is unclear. This is further complicated by evidence that a prototype influence emerges only later in learning and only following exposure to categories of large size (Homa, Dunbar, & Nohre, 1991). Since increasing category size produces substantial benefits on later transfer, subsequent transfer might be as good, if not better, in the non-repeating condition.

Recognition performance following learning in the REP and NREP conditions was also of interest. Because subsequent recognition of the training patterns should be poor in the non-repeating condition, classification of novel category members on a transfer test might also be poorer for the non-repeating condition as well.

By having subjects learn multiple categories, the interpretative problems attendant with single-category learning (Knowlton & Squire, 1993; Nosofsky & Zaki, 1998; Palmeri

& Flanery, 1999) was avoided. Furthermore, this procedure allows us to objectively track learning across trial blocks, something impossible in the single-category paradigm. Recognition and classification were assessed following a common learning procedure. In the experiments by Knowlton and Squire (1993), Reed et al. (1999), Nosofsky and Zaki (1998), and Zaki and Nosofsky (2001), different learning tasks preceded classification and recognition transfer and thus the data outcome – an apparent dissociation between recognition and classification – is compromised, i.e., what was demonstrated was that it is possible to find a task that produced equivalent classification but different recognition, not that classification transfer occurred in the absence of memorial traces (recognition).

Finally, the recognition transfer test for the non-repeating condition contained training patterns that had occurred on each of the previous training blocks. This has the advantage of determining the level of memory for training patterns that had been presented early, mid-way, or late in learning. A reasonable expectation is that whatever recognition memory exists for these patterns should be modulated by their placement in time during original learning.

We were especially interested in three performance issues and one major theoretical issue: (1) Can subjects readily learn categories with instances that never repeat in the training phase? (2) Do subjects exhibit non-zero recognition of these training patterns? (3) Does training that likely degrades memory for particular training instances also degrade subsequent transfer to novel instances from these categories? And (4) If the exemplar-based model of classification (Nosofsky, 1988) fails to predict the rate of learning when patterns repeat or not on each trial block, can an alternative model be shown to capture the results?

All experiments in the present study used multiple learning blocks, usually with three categories (although Experiment 2 also included a two-category paradigm). In Experiment 1, the transfer test involved classification of novel instances belonging to the learning categories. In Experiments 2 and 3, the transfer test involved a recognition test containing either old, new, and foil patterns (Experiment 2) or old, new, and prototype patterns (Experiment 3). Foil patterns were stimuli generated from prototypes that were different from those used in learning. The modification of the recognition test in Experiment 3 placed a severe constraint on exemplar knowledge, since all patterns were members of the learned categories but only a portion of these were old. The memorial integrity of training patterns as revealed in recognition is relevant to the evaluation of current models of categorization. Also, signal detection measures (d' and β) can be calculated for the REP and NREP conditions in Experiments 2 and 3.

Experiment 1

Method

Subjects The subjects were 58 undergraduates from an Introductory Psychology course at Arizona State University, 26 subjects in the Repetition condition (REP) and 32 in the Non-repetition condition (NREP).³ Subjects were randomly assigned to the REP or NREP condition.

Materials and apparatus The subject sat in a sound-dampened chamber about 20 in. from a 17-in. computer monitor. The patterns used in learning and transfer were the distorted forms used previously (e.g., Homa, 1978). In brief, a form category is created by first generating a random nine-dot configuration within a 50×50 grid and then connecting the dots with lines. This pattern is arbitrarily designated as the category prototype; different members of this category are then generated by statistically moving each of the dots of the prototype. In the present study, six different prototypes (A–F) were generated, with about half the subjects receiving one set (A, B, C) and half receiving the other set (D, E, F).

The amount of dot displacement determines the distortion level of a pattern. Low-, medium-, and high-level pattern distortions have vertices that are displaced, on average, about 1.20, 2.80, and 4.60 units, respectively, from each corresponding dot of the prototype. Patterns belonging to different prototype categories have their dots displaced, on average, by 10–15 units (Homa, 1978). All learning patterns were medium-level distortions; transfer patterns included low-, medium-, and high-level distortions, as well as the category prototype for each category. All patterns were randomly generated from a statistical algorithm, with the only restriction being that the generated pattern fell within a pre-specified distortion range. During the learning phase, patterns were randomly sampled for each trial block and subject to the restriction that the patterns in the REP condition also appear in the NREP condition.

Each pattern appeared in white against a black background. The subject responded by depressing the A, B, and C keys on a standard keyboard

Procedure All subjects received 20 trial blocks in the learning phase. In the repeating condition (REP), each trial block contained 15 patterns, five in each of three categories. The order of the patterns within a trial block was randomized. The same patterns were presented in each of the 20 trial

³ Based on previous studies with these stimuli and similar procedure, 25–35 subjects per condition provide an expected power of .80 and greater for relevant interactions involving known variables such as pattern distortion level. We targeted this number of subjects for all experiments and conditions herein.

blocks. For subjects in the non-repeating condition (NREP), the training patterns were always different.

The procedure on each learning block was the same – a pattern was presented in the center of the screen (about 7.5 cm along the horizontal and vertical dimensions) and remained visible until the subject responded with a key press indicating their category judgment. Immediately following the subject’s response, the correct category name appeared for 1 s below the pattern, followed by the next pattern. The presentation of the learning patterns was seamless across trials blocks, with no temporal break between blocks. For subjects in the REP conditions, each learning pattern appeared 20 times; for subjects in the NREP conditions, each pattern appeared once. All subjects saw 300 patterns.

Transfer Following learning, a 5-min, self-paced distracter task was used (rating CVCs for their pronounceability on a 7-point scale). The transfer task immediately followed the distracter task. On the transfer test, each subject classified, without feedback, 48 patterns, 16 from each of the three categories. The 16 patterns were composed of the prototype and five instances each that were low-, medium-, and high-distortions of the category prototype.

Results

Learning The mean proportion of correct classifications across the 20 trial blocks for the REP and NREP conditions is shown in Fig. 2. The effect of trial blocks was significant, $F(19, 1,064) = 108.19$, $MSe = .0093$, $\eta^2 = .659$, $p < .001$, but the effect of conditions (REP vs. NREP) was not, $F(1, 56) = 0.06$, $MSe = .146$, $\eta^2 = .001$, $p = .809$. The block \times condition interaction was significant, $F(19, 1,064) = 2.15$, $MSe = .0093$, $\eta^2 = .037$, $p = .003$, and reflected the slight crossover, with NREP exceeding REP performance for the initial 6 blocks, and the REP generally exceeding the NREP for the final 12.

Transfer Proportion correct performance on the transfer test is shown in Fig. 3 as a function of training condition.⁴ The main effect of item type (low, medium, high, prototype) was significant, $F(3, 168) = 31.58$, $MSe = .0052$, $\eta^2 = .361$, $p < .001$, but the overall difference of condition was not, $F(1, 56) = 1.63$, $\eta^2 = .028$, $MSe = .0012$, $p = .207$. The item \times condition interaction was significant, $F(3, 168) = 2.86$, $MSe = .0052$, $\eta^2 = .049$, $p = .038$. The interaction reflected the greater decline in accuracy with increasing distortion for the REP condition compared to the NREP condition. Subsequent tests revealed that the classification rate on medium distortions was significantly higher in the NREP versus REP condition ($t(56) = 2.02$, $p =$

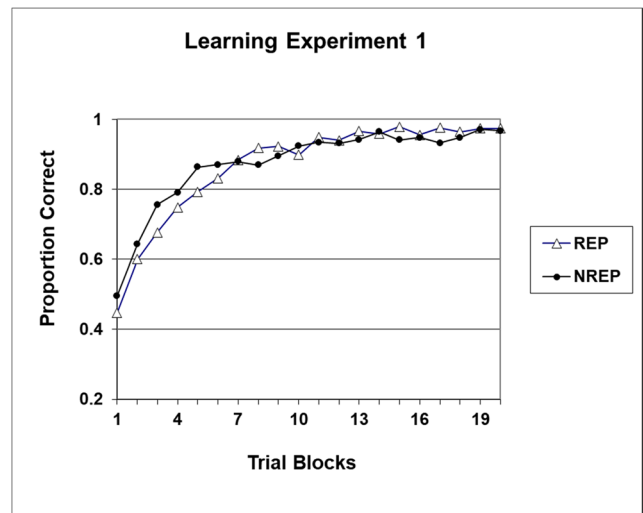


Fig. 2 Learning performance across the 20 trial blocks for the REP and NREP in Experiment 1 (standard errors at each block for each condition averaged between .01 and .03)

.049, two-tailed); the differences on the remaining items failed to reach significance ($p > .15$ in each case).

Discussion

Two major results emerged from Experiment 1. First, the rate of learning was not affected by having novel instances occur within each trial block. In fact, performance in the REP condition exceeded that for the NREP condition only late in learning. This latter outcome is hardly surprising, since subjects in the REP condition should eventually memorize the small set of training patterns, an outcome not possible in the NREP condition. Second, transfer performance was at least as good, if not slightly better, in the NREP condition. Taken together, having new patterns appear in each trial block did not affect the rate of learning while slightly enhancing subsequent transfer.

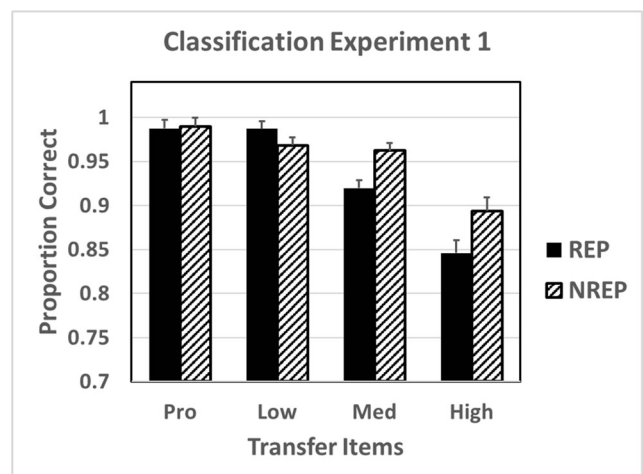


Fig. 3 Mean proportion correct performance (with standard error bars) on the transfer test

⁴ The standard errors shown on each figure and table reflect the recommendation by Cousineau (2005) that removes between-subject variability inherent in mixed designs containing repeated-measure factors.

Experiment 2

In Experiment 2 recognition rather than classification was assessed following REP and NREP learning. The number of learning blocks was reduced from 20 to 15 because learning appeared to asymptote after 15 learning blocks. Following learning, all subjects received a mixture of old (training), new, and foil patterns, where the foil patterns were novel patterns from different prototypes. We anticipated that subjects receiving REP training would clearly discriminate training patterns from new and foil patterns, but were less certain whether subjects receiving NREP training would show any ability to discriminate old from new patterns. Also, subjects learned either two or three categories prior to transfer, primarily to assess the generality of the acquisition results of Experiment 1.

Method

Subjects The subjects were 127 Arizona State University undergraduates, randomly assigned to the four between-subject conditions of number of categories learned \times mode of presentation (REP, NREP). Of the 58 subjects in the two-category condition (2C), 33 were randomly assigned to the repeating condition (REP) and 25 to the non-repeating (NREP) condition. Of the 69 subjects in the three-category condition (3C), 34 were in the REP condition and 35 in the NREP condition. None of the subjects had served in Experiment 1.

Procedure All subjects received 15 trial blocks in the learning phase. For the two-category condition, each trial block contained ten patterns, five in each of two categories; for the three-category condition, each trial block contained 15 patterns, five in each of three categories. Otherwise, the learning procedure was identical to that of Experiment 1.

Transfer On the transfer test, the subject briefly inspected the pattern and indicated whether it occurred on the learning set, typing “O” for old or “N” for new on the keyboard. In the 2C REP condition, a total of 29 patterns were presented, ten old (five from each category), ten new (five per category), and nine foils. All new patterns were medium-level distortions (as were the old); the foil patterns were medium-level distortions from three different prototypes (three each from three different prototypes), generated from prototypes not used in learning. In the 3C REP condition, a total of 39 patterns were presented, 15 old (five per category), 15 new (five per category), and nine foils (again, three each from each of three different prototype categories). In the 2C NREP condition, a total of 69 different patterns were used, 30 old (15 old from each of two categories), 30 new (15 each from each of two categories), plus nine foils (as before). In the 3C NREP, there were again 69 patterns, 30 old (ten old from each of three categories), 30 new (ten per category), and nine foils. Although the number of

transfer patterns differed in the various conditions, each condition had similar proportions of old and new patterns.

In the NREP conditions there were two old patterns on the transfer test from each training block, randomly selected from the available training patterns. The rationale for this manipulation is that memory strength for patterns that appeared only once in learning might be manifested primarily in the later blocks.

Results

Learning Learning performance in the REP and NREP conditions is shown in Fig. 4, with each condition, as before, showing substantial learning. The main effect of Blocks was significant, $F(14, 938) = 88.72$, $MSe = .014$, $\eta^2 = .570$, $p < .001$. Neither the main effect of Learning condition (REP vs. NREP), $F(1, 67) = 2.38$, $MSe = .185$, $\eta^2 = .034$, $p = .13$ nor the Block \times Condition interaction, $F < 1$, were significant.

Transfer – Recognition Figure 5 shows the likelihood that old, new, and foil patterns were called old on the transfer test. In the REP conditions, there was a clear discrimination between old and new patterns (Old = .896, New = .618), $t(66) = 12.77$, and old and foil patterns (Foils = .154), $t(66) = 27.18$, both $p < .001$, an outcome unchanged by the number of training categories. For the NREP condition, no discrimination was found between old and new patterns (Old = .771, New = .780), $t(59) = 0.72$, $p = .474$, with old exceeding foils (Foils = .188), $t(59) = 19.89$, $p < .001$. The only effect of number of training categories for the NREP condition was a higher rate of false alarming to the foil patterns in the three-category versus the two-category condition (.244 vs. .133), $t(58) = 2.14$, $p = .037$.

The initial analysis included the variables of condition (REP, NREP), number of training categories (two, three), and item type (old, new, foil). The main effect of item type was significant, $F(2, 246) = 776.51$, $\eta^2 = .863$, as was the Item \times Condition interaction, $F(2, 246) = 31.22$, $MSe = .0197$, $\eta^2 = .202$, both $p < .001$. Neither the main effect of Condition, $F(1, 123) = 1.15$, $p = .286$, nor Number of Training categories, $F(1, 123) = 2.93$, $p = .089$ was significant. The Number \times Items interaction was also significant, $F(2, 246) = 4.35$, $MSe = .0187$, $\eta^2 = .034$, $p = .014$.

Separate analyses for number of training categories did not alter the results. With either two or three training categories, the main effect of Items and the Item \times Condition interaction were each significant, $p < .01$ in each case; the main effect of Condition was not significant.

Recognition of old patterns in the NREP condition was also analyzed across blocks. Figure 6 shows the proportion of “old” responses to training patterns when they occurred in trial blocks 1–15. Analyses revealed no obvious trend across training blocks, with linear fits across trial blocks resulting in a non-significant, slightly negative slope for both the two- and the

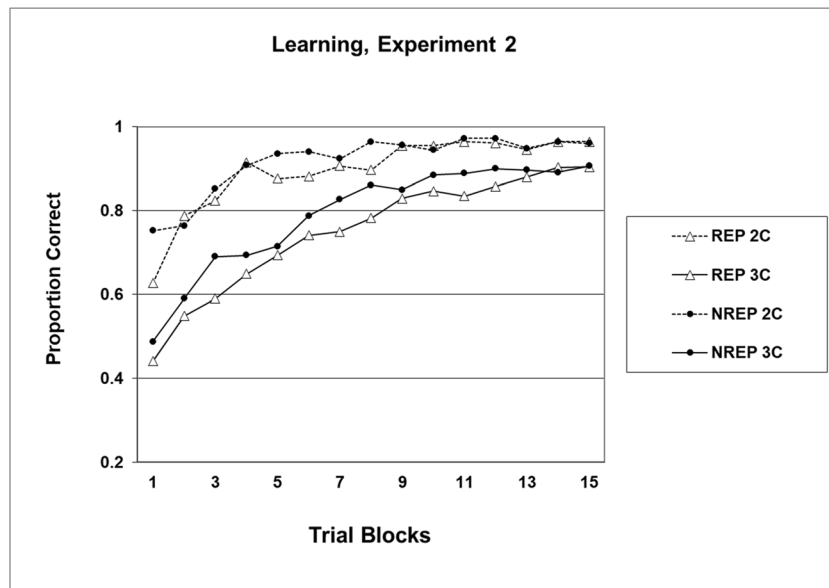


Fig. 4 Mean proportion correct classification across trial blocks for the REP and NREP conditions as a function of number of training categories, Experiment 2 (standard errors at each block and condition averaged between .01 and .03)

three-category learning conditions. In particular, calling an old pattern “old” did not increase with increasing block number.⁵

Signal detection analysis Each subject had their hit and false-alarm rate for old, new, and foil patterns converted to the signal detection measures of d' and a criterion. Hit and false-alarm rates of 1.00 and .000 were converted to .98 and .02 for computational purposes. For REP, the d' for the old versus foil patterns was 2.85 and 2.45 for subjects who learned two and three categories, respectively. For new items, these values were 1.81 and 1.31. The corresponding d' values for the NREP condition for old items were 2.32 and 1.66 for two and three categories, respectively. For the new items, these values were 2.17 and 1.83. The d' for old versus new patterns was, for REP subjects, 1.04 and 1.14 for the two- and three-category conditions, respectively. For the NREP condition, the d' for old versus new was 0.15 and -0.17, respectively.⁶

An analysis revealed that the main effect of pattern type on the transfer test was highly significant, $F(1, 123) = 116.79$, $MSe = .156$, $\eta^2 = .487$, $p < .001$, as was the number of training categories, $F(1, 123) = 10.40$, $MSe = 1.352$, $\eta^2 = .078$, $p = .002$. The significant type \times condition interaction was caused by the substantial differences in d' for the old and new patterns in the REP condition and the absence of a discrimination in the NREP condition, $F(1, 123) = 121.08$, $MSe = .156$, $\eta^2 =$

.496, $p < .001$. The only other significant source was the Type \times Condition \times Number of Categories, $F(1, 123) = 4.29$, $\eta^2 = .034$, $p = .04$, $MSe = .156$. The REP \times Number of Categories was not significant, $F = 0.04$, $p > .20$.

Figure 7 shows a signal detection representation at transfer for the Old, New, and Foil patterns for the three-category condition, separately for the REP (top panel) and NREP (bottom panel) training conditions. A notable outcome is the clear separation between old, new, and foil patterns in the REP condition, and the minimal separation between old and new patterns in the NREP condition. Placement of the criterion for calling an item old was similar for both conditions. A similar outcome (not shown) was obtained for the two-category training condition.

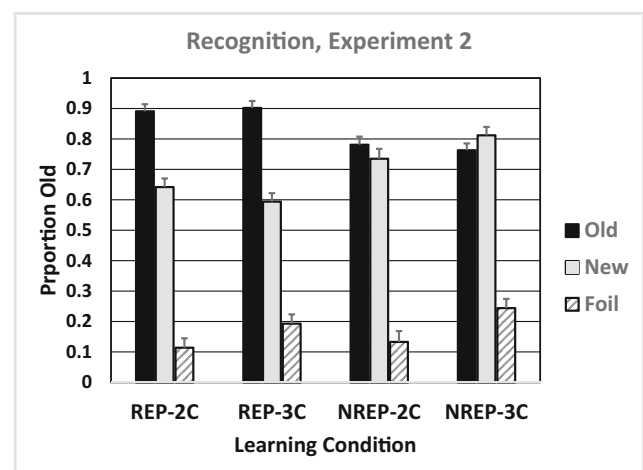


Fig. 5 Mean proportion of old responses (with standard error bars) to old, new, and foil patterns on the transfer test, Experiment 2

⁵ In the non-repeating condition, accuracy on the initial, middle, and final 23 recognition trials was .51, .52, and .49, respectively ($p > .20$). As a result, the failure to discriminate old from new stimuli in the non-repeating condition was not due to the longer test series, compared to the repeating condition.

⁶ For new versus foil patterns, d' is a measure of discriminability between these patterns.

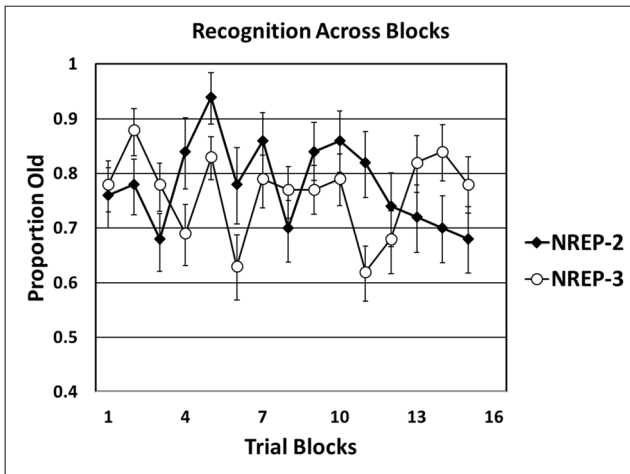


Fig. 6 The proportion of old responses (with standard error bars) on the transfer test to training patterns that occurred in trial blocks 1–15, Experiment 2

Discussion

Five major results were found in Experiment 2: (1) As was the case in Experiment 1, learning was unaffected by having different patterns represent each category on every block; (2) Subjects in the REP condition readily discriminated old, new, and foil patterns. In contrast, subjects in the NREP condition discriminated old and new patterns from foil patterns but could not discrimination old from new patterns; (3) Having subjects learn two or three categories altered none of the major results, and thus, the conclusions here are likely robust across number of training categories; (4) An analysis of recognition in the NREP condition across training blocks revealed no apparent trend, i.e., training patterns were called old at a rate independent of where in learning these patterns occurred; and (5) A signal

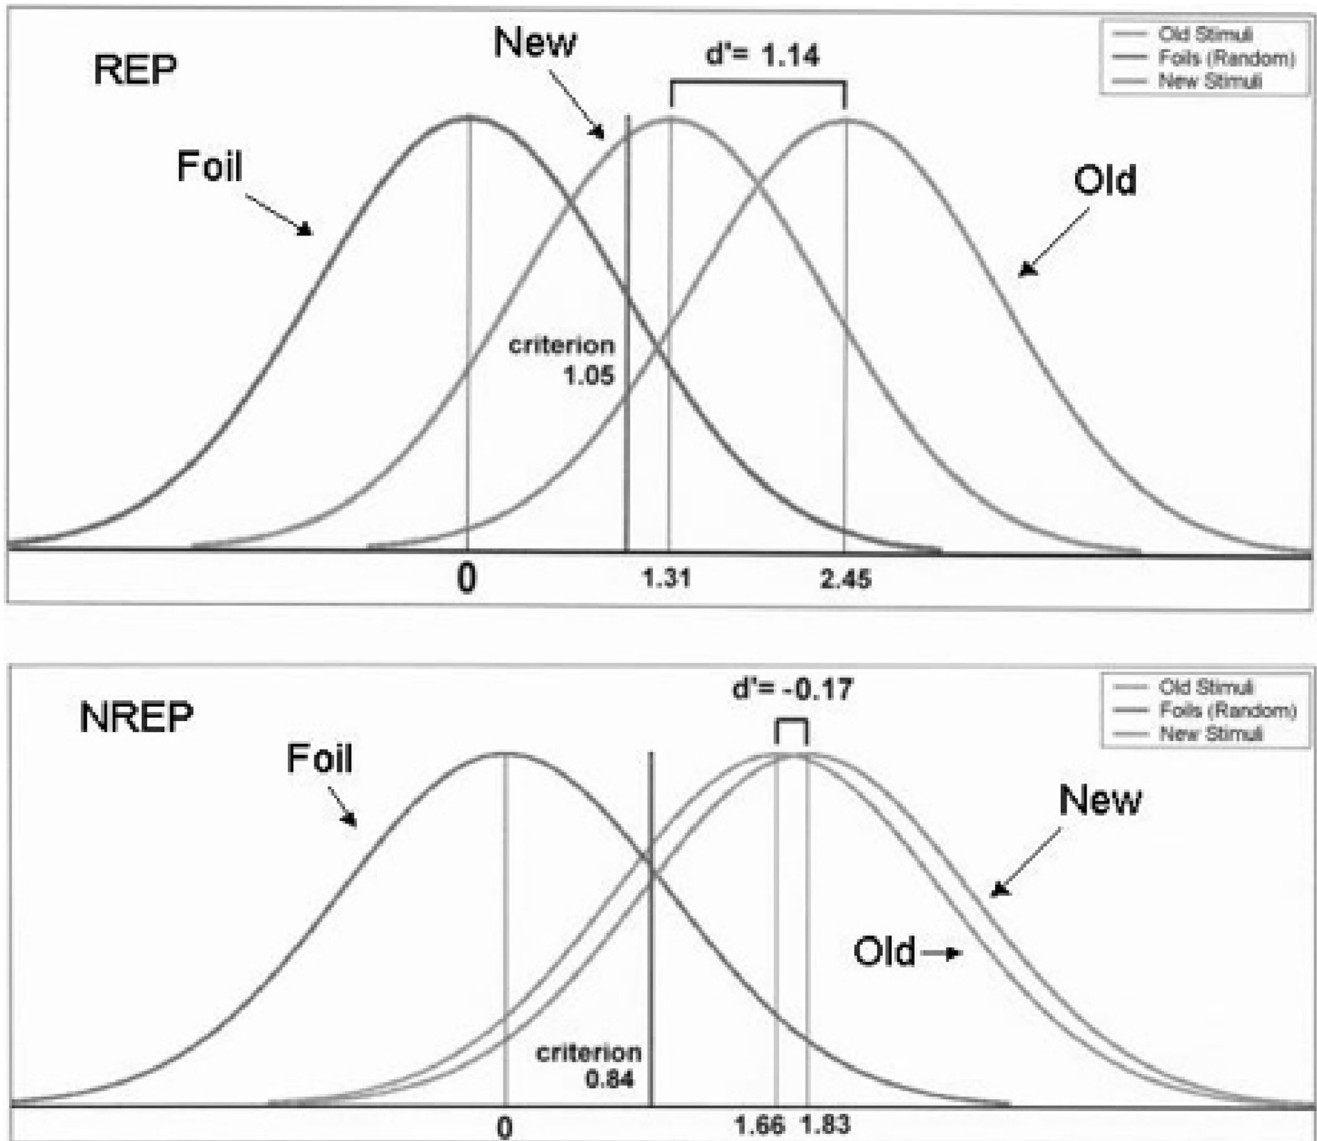


Fig. 7 A signal detection representation for the transfer items, three-category condition, Experiment 2

detection analysis suggested that subjects established a similar criterion for calling patterns old, with the old, new, and foil patterns clearly separated in the REP condition. However, no such discrimination between old and new patterns was revealed in the NREP conditions, although these patterns were clearly separated from the foils. The implication of the results in Experiment 2 was that subjects had little difficulty in learning the categories in the NREP condition, and performed at a rate equal to the REP subjects, even though no discernible memory for the training patterns was evident.

An alternative interpretation for the recognition results is that subjects in the NREP condition had some memory for the training instances but that they established a criterion to discriminate category patterns from non-category patterns, not old from new. To address this possibility, a final experiment was done in which no foils from other categories were used. Thus, all transfer patterns were patterns from the learning categories. If subjects can discriminate old from new patterns in the NREP condition, it should be apparent in Experiment 3. If, in contrast, subjects had no discernible memory for the training patterns, then no discrimination between old and new patterns should be evident. The category prototypes for the learning categories were also included in the recognition set.

Experiment 3

Experiment 3 was identical to Experiment 2 with two exceptions: (1) The transfer set contained old, new, and prototype patterns, and no foils from categories outside the learning categories were used; and (2) All subjects learned three categories.

Method

Subjects The subjects were 55 undergraduates at Arizona State University, 29 in the REP condition, and 25 in the NREP condition. None of the subjects had served in Experiments 1 or 2.

Procedure The procedure was identical to Experiment 2. Following learning, the transfer test contained only patterns from the categories represented in the learning phase. For the REP condition, the transfer test contained 33 patterns, 15 old (five from each category), 15 new (five from each category), and the three-category prototypes. In the NREP condition, there were 63 different patterns, 30 old (10×3 categories), 30 new (10×3 categories), plus the three-category prototypes. As was the case in Experiment 2, the transfer test for the NREP subjects contained two old patterns from each of the 15 trial blocks.

Results

Learning Figure 8 shows the learning performance across the trial blocks for the REP and NREP conditions. The substantial learning across blocks was significant, $F(14, 742) = 55.00$, $MSe = .014$, $p < .001$, but neither the main effect of condition, $F < 1$, nor the Condition \times Blocks interaction, $F < 1$, was significant, both $ps > .20$.

Transfer

Figure 9 shows the likelihood that old, new, and prototype patterns were called old on the transfer test. An analysis revealed that the main effect of condition was not significant, $F(1, 53) = 1.67$, $MSe = .056$, $\eta^2 = .031$, $p = .201$. However, the main effect of pattern type was significant, $F(2, 106) = 45.60$, $MSe = .016$, $\eta^2 = .462$, as was as the Pattern Type \times Condition interaction, $F(2, 106) = 18.15$, $MSe = .016$, $\eta^2 = .255$, both $ps < .001$. As was the case in Experiment 2, subjects discriminated between old and new patterns in the REP condition but failed to discriminate between these patterns in the NREP condition. The higher rate of calling the prototype old in the NREP versus the REP condition (.923, .863) was not significant, $t(53) = 1.32$, $p = .187$.

The likelihood that the old patterns in the NREP were called old as a function of where they appeared in acquisition is shown in Fig. 10. Once again, performance appeared to randomly fluctuate across training blocks. A linear fit across blocks revealed no evidence that the training patterns were more likely to be called old if they appeared later in training.

The signal detection analysis of the transfer results revealed that, once again, subjects discriminated old from new in the REP condition ($d' = 1.09$) but not in the NREP condition ($d' = -.04$). The d' values for the prototype-new discrimination was somewhat higher in the NREP than in the REP condition ($d' = 1.18$ vs. 0.81).

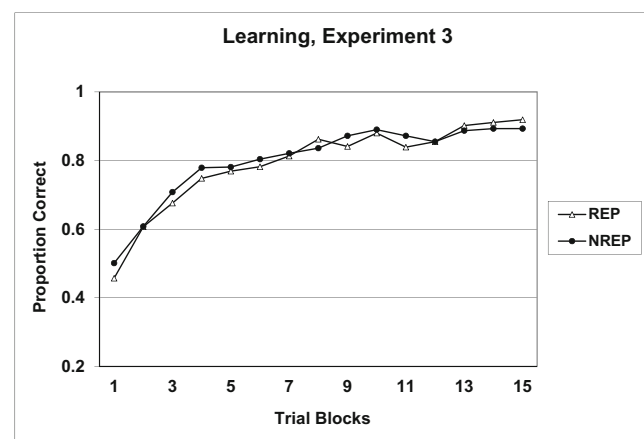


Fig. 8 The mean proportion correct classification across trial blocks for the REP and NREP conditions, Experiment 3 (standard errors at each block and condition averaged between .01 and .03)

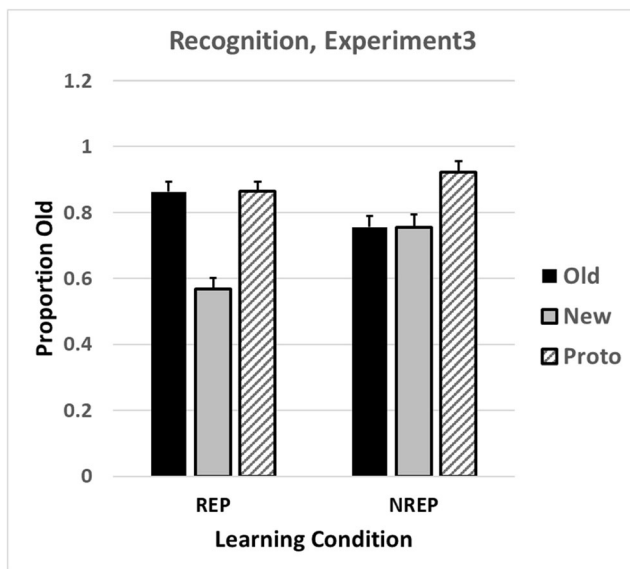


Fig. 9 Mean proportion of old responses (with standard error bars) to old, new, and prototype patterns on the transfer test, Experiment 3

General discussion

The main focus of the present study was to assess whether categories could be readily learned when novel patterns appeared on each trial block, rather than repeating them, as is typically done in categorization research. The results were clear-cut. In each experiment, generally when three categories were learned but also when two were used, the degree of learning was substantial and no worse than when the category patterns were repeated on each trial block. Interestingly, exemplar-based models of classification (Medin & Schafer, 1978; Nosofsky, 1988) predict a faster rate of learning in the repetition condition, an outcome forced by the greater

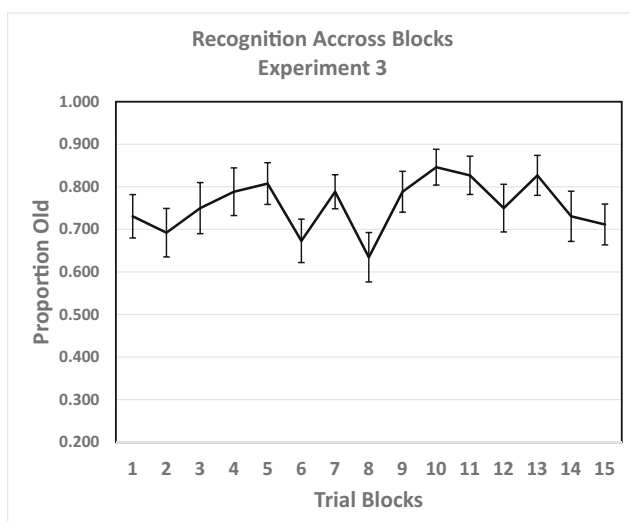


Fig. 10 The proportion of old responses (with standard error bars) on the transfer test to training patterns that occurred in trial blocks 1–15, Experiment 3

similarity to memorial traces when patterns are repeated in learning. Since this outcome was not obtained, alternative learning mechanisms must be considered.

As revealed in Experiment 1, subsequent generalization to novel patterns was not impacted by learning from non-repeating stimuli. Classification of novel patterns occurred at a rate equal to, or slightly better than, the classification rates obtained following repetition training. Notably, rapid learning and excellent classification occurred even though subjects exhibited no significant memory for the individual training patterns. Specifically, in the non-repeating condition, recognition of the old patterns was no better than chance, as indicated by the equivalent hit and false-alarm rates of old and new patterns. In contrast, substantial recognition of the old patterns occurred in the traditional repeating condition. Subsequent signal detection analysis revealed that subjects establish a comparable criterion for oldness judgments following repetition and non-repetition learning, with d' indistinguishable from zero for the non-repeat condition. Converging support arose from analysis of the hit rates for the patterns that occurred on each training block. One might have anticipated some memory for patterns, especially those that occurred late in the learning phase. Nonetheless, the recognition rate for all patterns was flat across blocks and at a level that did not differ from the false-alarm rate for new patterns. Finally, the subjects in the non-repeat condition, who failed to recognize any of the old patterns, nonetheless had a false-alarm rate for the category prototypes that was higher than for any other pattern and higher than the false-alarm rates for the prototypes in the repeat condition.

Taken together, these results pose substantial problems for exemplar-based models of classification. Proponents of exemplar theory might maintain that minimal memory of the training instances still existed following acquisition, which was responsible for the classification and recognition performance. Indeed, this is the basis of the argument by Nosofsky and Zaki (1998), who reassessed and re-interpreted the performance of amnesic patients in the study by Knowlton and Squire (1993). The amnesic subjects in their study had a lower recognition rate of the training patterns compared to the normal controls, but recognition accuracy by amnesics was about 65% (where chance was 50%). This reduced, but not absent, memory was sufficient to generate classification and recognition results similar to that obtained by Knowlton and Squire. Nosofsky and Zaki also addressed the performance of one subject, E. P., whose recognition performance *was* at chance yet whose classification was near normal, concluding that: “The baseline version of the exemplar model presented in this article is unable to fit this pattern of data.” (p. 254). In effect, the subjects in the present study are like the patient E.P., who had normal transfer for novel patterns even though his recognition memory for the training patterns was non-existent.

The argument might be made that our subjects extracted pattern fragments in the learning phase, rather than intact

patterns, sufficient to promote rapid learning and later generalization even though recognition for individual patterns might be virtually zero. This interpretation seems unlikely because the storage of pattern fragments might result in reduced memory for intact patterns, but not zero. That is, the existence of these pattern fragments would be matched with the entire pattern at the time of recognition, something not possible with the new patterns. As we noted previously (Homa et al., 1991), novel pattern fragments are unlikely to exactly repeat in another pattern when ill-defined, infinitely variable patterns are used, and thus, stored fragments could be exactly matched only with the training pattern.

What we believe is responsible for the learning, classification, and recognition performance in the non-repeating conditions is the abstraction of a category prototype early in learning, which continues to evolve as more and more patterns are encountered in training. Further, this abstraction process is more likely to occur when the patterns never repeat, since an initial strategy of attempting to memorize individual patterns is doomed to failure as more and more unique patterns are encountered. This fosters an abstraction process early on, with abstracted prototypes formed and continually modified as different learning patterns are presented. There is, in principle, no necessity to store particular learning patterns in order to learn the category. The subsequent classification is easy to explain – research has shown for many years that increasing category size aids later classification, an outcome that has been repeatedly obtained (e.g., Homa et al., 1973; Homa et al., 1991; Homa, Proulx & Blair, 2008). Subjects in the non-repeating condition are exposed to precisely that – a category of substantial size, relative to the category size of the repeating conditions. The recognition results would similarly follow if, indeed, the subject stores the abstracted prototype and little else.

In the following section, we introduce a model of categorization for the non-repeating condition that assumes that continual learning increasingly alters the similarity relationship to an abstracted prototype. The current preliminary model can capture the learning, transfer classification, and transfer recognition obtained in the present study.

Model of categorization The proposed model uses a common set of parameters to capture three critical sets of results: (1) Learning is as rapid when patterns never repeat as when they do; (2) Following learning in the non-repeating condition (NREP), subjects demonstrate little or no memory for the training patterns; and (3) Classification of novel patterns is no worse and perhaps slightly superior following training on categories whose patterns never repeat. The model has one major modification compared to previous models – we assume that similarity among the transfer patterns is modified and dependent upon the level of learning. That is, across training, categories begin to emerge, represented by increasingly distinctive clusters in the multidimensional space. Otherwise,

the proposed model shares similarity to other models that represent similarity by the empirically determined multidimensional distances separating the patterns (e.g., Homa et al., 2008; Nosofsky & Johansen, 2000; Minda and Smith, 2001).

The current model makes one other critical assumption: when a category is represented by a small number of patterns that are repeatedly presented in learning, the subject stores these exemplars and uses these stored traces for both recognition and classification. However, when the patterns representing a category are sufficiently numerous (especially when never repeated, as in the present study), the subject abstracts the central tendency, which then forms the basis for all subsequent decisions. There is support for the view that small-sized categories, e.g., three to five patterns per category, are represented as singular traces whereas larger-sized categories, e.g., ten or more, are more likely to rely upon summary representations, such as an abstracted prototype (e.g., Homa, Sterling, & Trepel, 1981; Homa et al., 2008). There is also support for the assertion that similarity relationships are increasingly modified by degree of prior learning (Homa, Rhodes, & Chambliss, 1979). In Homa et al. (1979), the multidimensional space of three categories was increasingly modified via different levels of learning. In particular, additional learning increasingly reduced within-category distances (and, therefore, dissimilarity) as reflected by the subsequent multidimensional configuration, especially for categories represented by numerous instances.

Our initial model adopted the simplifying assumption that categories are largely unstructured at the outset of learning and that the distances among the patterns common to a category are increasingly reduced with additional learning blocks. Specifically, we assumed in the version presented here that the multidimensional distance among, for example, the distance between any two medium level patterns terminated close to the value of 0.72. This latter distance is the mean value obtained from the multidimensional scaling of patterns following extensive learning (e.g., Homa et al., 2008). On the initial block, we assumed that this value was 1.60 (the mean distance between any two patterns prior to learning). The convergence of distance between any two medium-level distortions, common to a category, to its terminal learning value obeyed the formula:

$$d(B) = \frac{d_b - d_{mm}}{B^\lambda} + d_{mm}$$

Here B is the block number, d_b is the MDS distance between any two patterns from different categories, d_{mm} is the MDS distance between any two patterns from the same category, and λ is a decay parameter that regulates how quickly similarity is modified by learning.

Since $d_b = 1.60$ and $d_{mm} = .72$, the equation reduces to:

$$d(B) = \frac{1.60 - 0.72}{B^\lambda} + 0.72$$

Therefore, following an extremely large number of learning trials, the MDS distance between any two medium-level

Table 1 Fixed parameter values for all conditions, with values for present Experiment

C	Number of categories	3
B _i	Current training block	varies
B _T	Transfer testing block after training	21
N	Number of exemplars used in REP training (per block, per cat)	5
T _{ol}	Number of old items tested during transfer for low distortion	0
T _{om}	Number of old items tested during transfer for medium distortion	5
T _{oh}	Number of old items tested during transfer for high distortion	0
T _{nl}	Number of new items tested during transfer for low distortion	5
T _{nm}	Number of new items tested during transfer for medium distortion	5
T _{nh}	Number of new items tested during transfer for high distortion	5

distortions was 0.72. The decay parameter (λ) allowed pattern distance to change at different rates across learning blocks for the REP and NREP conditions.

In our model, performance in the REP conditions is exemplar based, whereas performance in the NREP condition is prototype based. In the following equations, parameters in upper case are fixed parameters based on the design of the experiment (e.g., C, the number of categories is a fixed parameter; compared to c, the sensitivity-free parameter). The fixed parameters for all conditions are defined in Table 1, along with their values for this study. Empirically determined parameters, the MDS distance measures, are presented in Table 2, along with their definitions and values in this study. Lower-case parameters are free parameters and were adjusted in the process of finding values that minimized the RMSE error to mean data across subjects (using the `fminsearch` function in Matlab; e.g., β is processing noise that is hypothetical). The free parameters along with their definitions and best fit values are given in Table 3. Least square best fits between observed and predicted values were determined simultaneously for all conditions and three data sets – learning, transfer-classification, and transfer-recognition. In addition, best fits were determined by assigning equal weights across these three data sets. The adequacy of the data fits is addressed later.

Similarity measures are defined by MDS distances and scaling parameter c during training by the equation (REP and NREP, respectively):

$$S_{mm} = \exp \left[-c_r \left(D_{mm} + \frac{D_b}{B_i} \right) \right] \quad (5a)$$

$$S_{pm} = \exp \left[-c_n \left(D_{pm} + \frac{D_b}{B_i} \right) \right] \quad (5b)$$

Here, S_{mm} is the similarity between two medium-level distortions in the REP condition; S_{pm} is the similarity between a medium-level distortion and prototype pattern in the NREP condition. The added term D_b / B_i reflects the assumption that similarities increase with more training, i.e., block number (B_i).

The similarity measures for transfer are defined by the equation (REP and NREP, respectively):

$$S_{my} = \exp(-c_r D_{my}) \quad (5c)$$

$$S_{py} = \exp(-c_n D_{py}) \quad (5d)$$

The y subscript denotes whether we are calculating similarity scores for low, medium, or high-distortion exemplars.

Training equations Performance is determined by Luce's choice axiom (Luce, 1959). For REP training, the probability of correctly categorizing a stimulus in training block B_i is given by:

Table 2 Fixed multidimensional distance values for all conditions

D _x	Distances for exact matches (identity match) (used in all REP portions of all experiments)	0.0
D _{mb}	Between category distances, equal to D _{pb} (used in all REP portions of all experiments)	1.6
D _{pb}	Between category distances, equal to D _{mb} (used in all NREP portions of all experiments)	1.6
D _{ml}	Distance between medium distortion and low distortion exemplars (used in REP/transfer)	0.5
D _{mm}	Distance between medium and medium distortion exemplars (used in REP/transfer and REP/recognition)	0.72
D _{mh}	Distance between medium and high distortion exemplars (used in REP/transfer)	0.86
D _{pl}	Distance between prototype and low distortion exemplars (used in NREP/transfer)	0.24
D _{pm}	Distance between the prototype and medium distortion exemplars (used in NREP/transfer and NREP/recognition)	0.45
D _{ph}	Distance between the prototype and high distortion exemplars (used in NREP/transfer)	0.82

Table 3 Parameter values and root mean square error (RMSE) for the mixed and exemplar model

	Mixed	Exemplar
c_r Sensitivity parameter for REP	2.953	2.933
c_n Sensitivity parameter for NREP	1.532	1.944
γ Response scaling parameter	2.336	2.272
β_r Processing noise parameter for REP	3.009	2.866
β_n Processing noise parameter for NREP	0.000	0.000
θ_r Recognition threshold for REP	545.520	454.453
θ_n Recognition threshold for NREP	0.127	917.197
λ Similarity modification across blocks	0.552	0.624
RMSE	0.028	0.048

$$P_{correct} = \frac{[\beta_r + (B_i - 1)(1 + (N - 1)S_{mm})]^\gamma}{[\beta_r + (B_i - 1)(1 + (N - 1)S_{mm})]^\gamma + (C - 1)[\beta_r + (B_i - 1)\lambda S_{mb}]^\gamma} \quad (6)$$

This equation mirrors equations (1–4), with the major change that similarity among the training patterns is continuously modified as learning progresses.

The equation for training in the NREP condition is similar to that for the REP condition with two key differences. First, similarity values reflect the similarity of the exemplars to the prototype rather than to each other; second, when calculating similarity, the sensitivity parameter for REP (c_r) and NREP (c_n) (equation 5) were allowed to take on different values. The number of categories (C) remained the same between REP and NREP. For NREP training, the probability of correct categorization in training block B_i is

$$P_{correct} = \frac{[\beta_n + S_{pm}]^\gamma}{[\beta_n + S_{pm}]^\gamma + (C - 1)[\beta_n + S_{pb}]^\gamma} \quad (7)$$

Figure 11 shows the learning results obtained in Experiment 1 (top panel) and the model’s predictions (bottom panel), shown separately for the REP and NREP conditions. The critical result – that NREP fares no worse in learning than REP, except for the final learning blocks when subjects likely memorized the fewer patterns in the REP condition – is nicely captured by our model.

Transfer – classification Transfer patterns were composed of new patterns at one of three level of distortion from the prototype and the category prototype itself. For the REP condition, pattern similarities are computed based on similarity to stored patterns in memory; for the NREP condition, pattern similarities are based on similarity to the category prototype.

For a new medium-distortion exemplar in REP, the probability is:

$$P_{correct} = \frac{[\beta_r + (B_T - 1)NS_{mm}]^\gamma}{[\beta_r + (B_T - 1)NS_{mm}]^\gamma + (C - 1)[\beta_r + (B_T - 1)NS_{mb}]^\gamma} \quad (8)$$

Predicted classification of the low, high, and prototype stimuli requires only that the appropriate similarity be

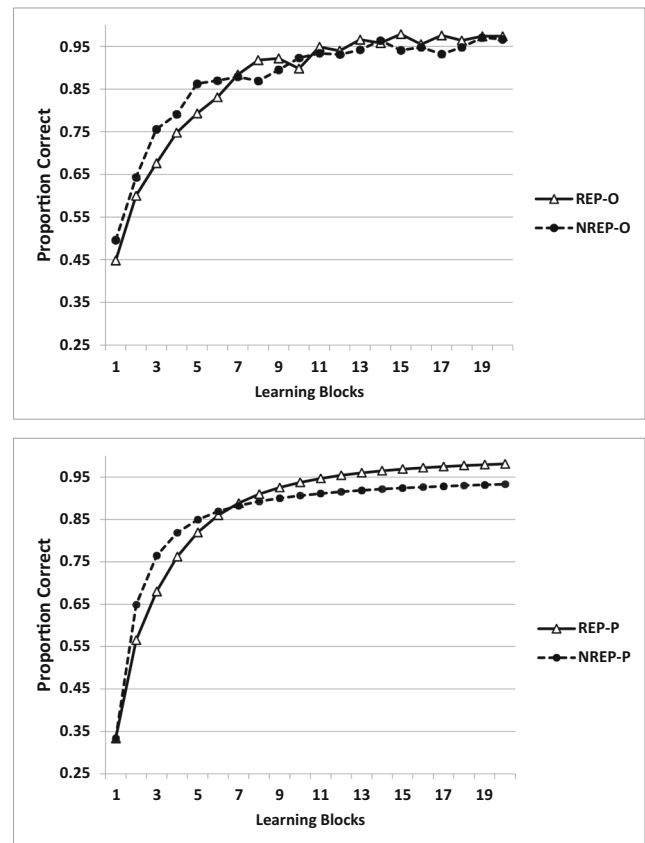


Fig. 11 Observed and predicted performance in learning for the REP and NREP conditions

substituted into equation (8). For example, low-distortion performance can be calculated as:

$$P_{correct} = \frac{[\beta_r + (B_T - 1)NS_{ml}]^\gamma}{[\beta_r + (B_T - 1)NS_{ml}]^\gamma + (C - 1)[\beta_r + (B_T - 1)NS_{mb}]^\gamma} \quad (9)$$

The NREP classification calculations, based on prototype models, are similar in structure. In the classification trial in which the prototype is presented, we must remember that the experimenter’s prototype is actually compared to the subject’s internal prototype representation, which for our purposes are assumed to be an exact match ($S_{pp} = 1$).

$$P_{correct} = \frac{[\beta_n + S_{pp}]^\gamma}{[\beta_n + S_{pp}]^\gamma + (C - 1)[\beta_n + S_{pb}]^\gamma} \quad (10)$$

All the other NREP classifications can be calculated by substituting S_{pp} with S_{py} , where y is the appropriate distortion level desired.

Figure 12 shows the obtained and predicted values for classification (Experiment 1) of the category prototype, low, medium, and high-level distortions. Overall, the model captures the gradient across distortion level, including the slight but

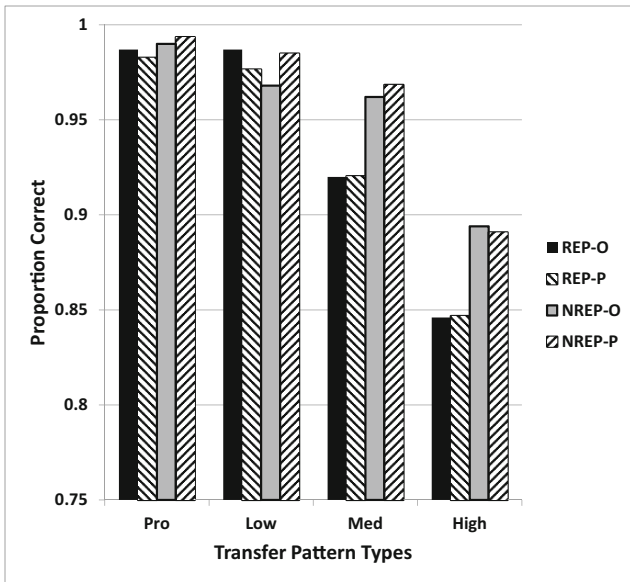


Fig. 12 Observed and predicted transfer performance on the classification of the different transfer patterns (prototype, low, medium, high), shown separately for the REP and NREP conditions

significant advantage of the NREP condition for the medium level distortions.

Transfer – recognition In recognition, we have different but conceptually similar equations for REP and NREP conditions for old, new, and prototype recognition test items. The difference between recognition judgments and categorization judgments is reflected in the use of a recognition threshold, θ_r for REP and θ_n for NREP.

To denote the change for REP, S_{rmm} is used (similarity between two medium patterns in the REP condition) instead of S_{mm} ; NREP similarity is denoted as S_{nmm} . The probability of calling “old” a previously-seen exemplar in the REP condition is modeled as:

$$P_{rep|old} = \frac{[\beta_r + (B_T - 1)(1 + (N - 1)S_{rmm} + (C - 1)NS_{rmb})^\gamma]}{[\beta_r + (B_T - 1)(1 + (N - 1)S_{rmm} + (C - 1)NS_{rmb})^\gamma] + \theta_r} \tag{11a}$$

The probability for a subject to call a new exemplar “old” is modeled as:

$$P_{rep|new} = 1 - P_{rep|old} \tag{11b}$$

NREP probability is similarly structured, except for the fact that we are using prototype-based models:

$$P_{nrep} = \frac{[\beta_n + S_{npm} + (C - 1)S_{npb}]^\gamma}{[\beta_n + S_{npm} + (C - 1)S_{npb}]^\gamma + \theta_n} \tag{12a}$$

The probability for a subject to call a new exemplar “old” is modeled as:

$$P_{nrep|new} = 1 - P_{nrep|old} \tag{12b}$$

The obtained and predicted performance on the recognition test is shown in Fig. 13. Critically, the model captures the result that subjects in the REP condition could discriminate the training from the new patterns whereas subjects in the NREP could not.

Fit to an exemplar model of categorization We also fit an exemplar-based model to the same set of results – learning, transfer-classification, and transfer-recognition for repeated versus non-repeated pattern learning – using the same parameters as used for our mixed model. The only change was that pattern similarity was based on the relationship of patterns to stored exemplars rather than the category prototype, regardless of whether learning involved the repetition of patterns across training blocks or the presentation of novel patterns within each block.

Overall, the mixed model outperformed the exemplar model (RMSE = .028 for the mixed model and .048 for the exemplar model). The advantage of the best-fitting mixed model, versus the exemplar model, arose primarily in the transfer data (e.g., subsequent classification [Experiment 1] and recognition [Experiment 3]). Figure 14 shows these disparities, with the mixed model in the upper panel and the exemplar model in the lower panel. Although the

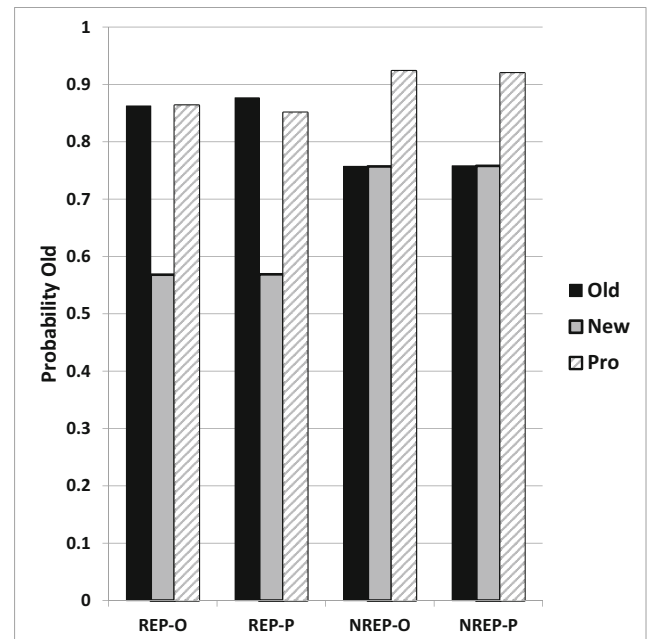


Fig. 13 Observed and predicted likelihood of calling a pattern “old” on the transfer test, as a function of pattern type (old, new), shown separately for the REP and NREP conditions

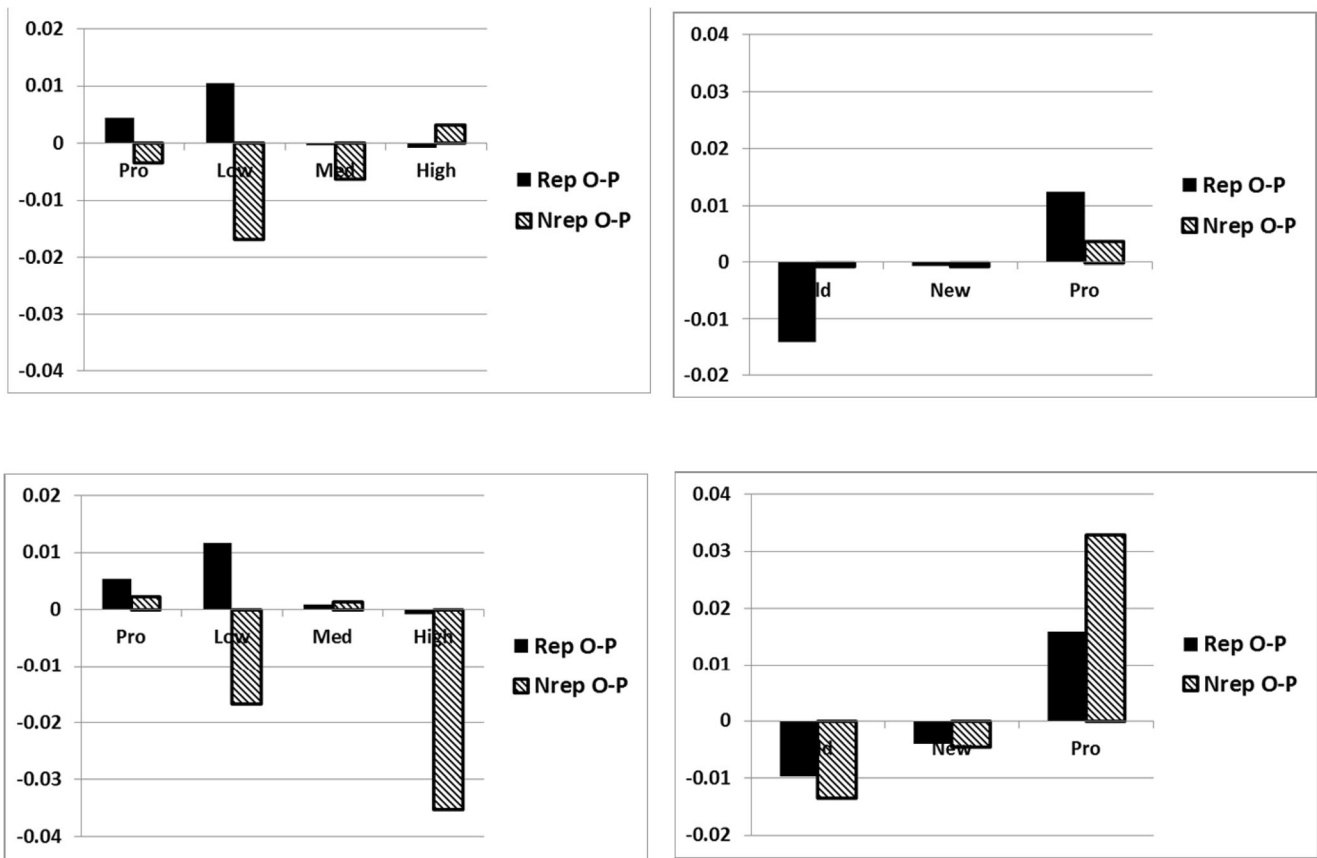


Fig. 14 Transfer predictions of Mixed (upper panel) and Exemplar (bottom panel) models for classification (left panels) and recognition (right panels)

exemplar model provides a reasonable fit to the data, including fit to the category prototype, critical contrasts generally favored the mixed model, especially for the NREP condition.⁷ In particular, the exemplar model over-predicted classification of the high distortions and under-predicted the false recognition rate of the prototype. The former outcome is consistent with previous concerns that the exemplar model generally failed to capture the magnitude of the gradient across decreasing similarity to the category prototype (e.g., Homa et al., 1981; Homa & Powell, & Ferguson, 2014; Smith & Minda, 2002). Two additional comments regarding the parameters are worth noting: (1) the best fitting threshold parameter is large for the mixed model for REP but close to zero for NREP. This implies that recognition performance is above chance for REP but is essentially at chance for NREP, as in the data; the exemplar model does not have this property; and (2) the noise parameter is zero for NREP but substantially above zero for REP. This suggests that the repetition of stimuli across learning blocks in the REP condition generates increasing noise whereas the virtual absence of memory for stimuli in the NREP condition does not. We should reiterate that the mixed model is an exemplar model for the repeat condition and a prototype model for the non-repeat condition. Not surprisingly, the mixed model

and the exemplar model predict with similar accuracy for the repeat condition.⁷

Conclusion

Our preliminary model does an adequate job in capturing the major results of our study. Learning is predicted to be as rapid when patterns repeat as when they do not, little or no memory for the training patterns exists when patterns are never repeated in learning, and subsequent classification of novel patterns, including those of increasing distortion level, was demonstrated.

Being able to capture results does not prove the legitimacy of our model (Wills & Pothos, 2012), but it demonstrates that a prototype model can capture critical results following non-repeating training better than the exemplar model explored in

⁷ One reviewer asked whether the high classification and recognition rates of the category prototype in the REP condition might reflect the joint contribution of memorized instances as well as the abstraction of the central tendency (prototype) of that category. We have elsewhere suggested that the joint influences of these factors might occur, modulated by factors like category size and level of learning (Homa et al. 2008). We cannot deny this alternative explanation for the present study, although exemplar influences alone were sufficient to capture results in the REP condition.

the present study. Some assumptions were adopted for simplicity's sake, such as the manner in which similarity changed across blocks and the assumption that learning on the first block would be at chance. We cannot claim that other models yet considered would fare less well than our current one, although the finding that learning can occur rapidly even when subjects cannot discriminate old from new patterns at better than chance accuracy poses a severe challenge for any exemplar or related connectionist models (e.g., ALCOVE, Kruschke, 1992). An interesting exception could be Minerva II (Hintzman, 1986), which stores particulars but can generate an echo mimicking the prototype. Given the severe memorial degradation of exemplars in the non-repeating condition, an issue is whether sufficient exemplar knowledge resides in memory to even generate a functional echo. A full examination of how other models may or may not capture our results is beyond the current scope of this paper. Regardless, the results in the present study are novel to the field and interesting in their own right since no classificatory cost was obtained when learning patterns never repeated, a characteristic likely to be common in natural learning situations but absent in current paradigms.

Acknowledgements We would like to thank Steve Goldinger and Fief Gunersoch for their comments on an earlier version of this manuscript, and LaMista Johnson, Tiffany Coor, and Rose Halterman for assisting with the data collection.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 33–53.
- Ashby, F. G., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 598–612.
- Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 50–71.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. New York: Wiley.
- Busemeyer, J. R., & Diederich, A. (2009). *Cognitive modeling*. Sage publications: New York.
- Casale, M.B., Roeder, J.L., & Ashby, F.G. (2012). Analogical transfer in perceptual categorization. *Memory & Cognition*, *40*, 434–449.
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, *1*, 42–45.
- Do, P., Homa, D., & Koehler, K. (2014). Identity categories and transformational paths for face changes across the age spectrum. *Memory & Cognition*, *42*, 340–353.
- Fisher, S. C. (1916). The process of generalizing abstraction; and its produce, the general concept. *Psychological Monographs*, *31* (whole No. 90).
- Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.
- Homa, D. (1978). Abstraction of ill-defined form. *Journal of Experimental Psychology: Human Learning and Memory*, *4*, 407–416.
- Homa, D. (1984). On the nature of categories. In G. H. Bower (Ed.), *The psychology of learning and motivation*, vol. 18. Academic Press: New York.
- Homa, D., & Chambliss, D. (1975). The relative contributions of common and distinctive information on the abstraction from ill-defined categories. *Journal of Experimental Psychology: Human Learning and Memory*, *1*, 351–359.
- Homa, D., Cross, J., Cornell, D., Goldman, D., & Shwartz, S. (1973). Prototype abstraction and classification of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology*, *101*, 116–122.
- Homa, D., Dunbar, S., & Nohre, L. (1991). Instance frequency, categorization, and the modulating effect of experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 444–458.
- Homa, D., Hout, M., Milliken, L., & Milliken, A. M. (2011). Bogus concerns about the false prototype enhancement effect. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *37*, 368–377.
- Homa, D., Powell, D., & Ferguson, R. W. (2014). Array training in a categorization task. *Quarterly Journal of Experimental Psychology*, *67*, 45–59.
- Homa, D., Proulx, M.J., & Blair, M. (2008). The modulating influence of category size on the classification of exception patterns. *The Quarterly Journal of Experimental Psychology*, *61*, 425–443.
- Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory*, *5*, 11–23.
- Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory*, *2*, 322–330.
- Hull, C. I. (1920). Quantitative aspects of the evolution of concepts. *Psychological Monographs*, *28* (whole No. 123).
- Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, *262*, 1747–1749.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.
- Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, *29*(1), 1–27.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 607–625.
- Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 775–799.
- Moore, T. V. (1910). The process of abstraction: An experimental study. *University of California Publications in Psychology*, *1*, 73–197.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of experimental Psychology: Learning, Memory, & Cognition*, *14*, 700–708.
- Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of multiple-system phenomena in perceptual categorization. *Psychonomic Bulletin & Review*, *7*(3), 375–402.
- Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, *9*, 247–255.

- Palmeri, T. J., & Flanery, M. A. (1999). Learning about categories in the absence of training: Profound amnesia and the relationship between perceptual categorization and recognition memory. *Psychological Science, 10*, 526-530.
- Reed, J. M., Squire, L. R., Patalano, A. L., Smith, E. E., & Jonides, J. (1999). Learning about categories that are defined by object-like stimuli despite impaired declarative memory. *Behavioral Neuroscience, 113*, 411-419.
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function: I. *Psychometrika, 27*(2), 125-140.
- Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121*, 278-304.
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the midst: 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 dot-pattern category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 800-811.
- Wills, A. J., & Pothos, E. M. (2012). On the adequacy of current empirical evaluations of formal models of categorization. *Psychological Review, 138*, 102-125.
- Zaki, S. R., & Nosofsky, R. M. (2001). A single-system interpretation of dissociations between recognition and categorization in a task involving object-like stimuli. *Cognitive, Affective, & Behavioral Neuroscience, 1*, 344-359.