

Learning from Environmental Regularities is Grounded in Specific Objects not Abstract Categories

Lauren L. Emberson (lle7@cornell.edu)

Sackler Institute for Developmental Psychobiology; Weill-Cornell Medical School
Box 140, 1300 York Ave., NY, NY 10065 USA

Dani Rubinstein (rubinstein@nih.gov)

MEG Core Facility, National Institute of Mental Health
Building 10 - Magnuson Clinical Center, B1D65B, 10 Center Drive, Bethesda, MD 20892 USA

Abstract

This paper examines statistical learning in the presence of predictive regularities at multiple levels of abstraction. Participants were presented with streams of pictures where picture order was predicted by both object identity and the categories these objects belong to. In Experiment 1, we establish that participants do learn based on the specific objects and not solely at the abstract, categorical level. In Experiment 2, we discount the possibility that participants gain abstract knowledge in addition to more concrete, object-based knowledge. Moreover, we consistently find equal learning in those who viewed the atypical exemplars and those who viewed the typical exemplars of the categories. Overall, our results suggest that when learning from environmental regularities, object-specific information takes precedence over more abstract, category level information when both are predictive.

Keywords: Statistical learning; environmental learning; visual development; perceptual learning; object perception; categorization.

Introduction

Throughout our lifetimes, it is clear that experience shapes our mental model of the world. Focusing on adulthood, adults learn to recognize new objects and categories as well as new properties of familiar objects; they learn new words and adapt to changing patterns in the ambient language, all by adapting future behavior based on experience. Despite the clear importance of learning from information in the environment, the nature of the mechanisms that support learning from real-world experience is largely unknown. A central problem in this literature is how learning mechanisms operate given the richness of the information we get from the world. Are learning mechanisms *a priori* constrained to learn particular patterns? Can learning proceed along many types of perceptual information and/or at different levels of abstraction¹ simultaneously?

In this paper, we focus on a type of learning called “statistical learning” where participants passively learn from

stimuli embedded with probabilistic information². Previous research has supported the view that these experiential learning mechanisms are unconstrained: statistical learning has been demonstrated in multiple sensory modalities (Conway & Christiansen, 2005), across a wide range of perceptual input. For example, in the visual modality, learning can occur from sequences of gestures (Baldwin, Anderrson, Saffran, & Myers, 2007) as well as abstract shapes (Fiser & Aslin, 2001). While the majority of these studies have focused on learning probabilistic relations of individual items or objects, there is evidence that learning can occur at higher levels of informational abstraction including over new categories of nonsense words (Saffran, 2002) and based on familiar semantic categories (Brady & Oliva, 2008).

Overall, these studies support the view that environmental learning is unconstrained. That is, if there is any reliability probabilistic information in the environment, humans can learn from it regardless of level of abstraction or perceptual properties. If learning is entirely unconstrained, it is unclear how learning mechanisms operate in complex environments where information from multiple sources and at many levels of abstraction abounds.

However, these behavioral demonstrations of an entirely unconstrained learning mechanism arise from paradigms in which information is only predictive at a single perceptual and/or informational dimension. For example, while Brady and Oliva (2008) demonstrate learning of categories of scenes, participants were presented with a new scene from the category during each successive presentation. In this paradigm, individual scenes (e.g. beach₁ and beach₂) are not predictive of picture order, only the category of pictures are (e.g. a beach predict a kitchen but beach₁ does not predict kitchen₁), thus it would be impossible for participants to learn based on individual scenes. Thus, these results provide an existence proof of an unconstrained learning mechanism but they arise under specific, restricted conditions.

In actuality, environmental stimuli exhibit statistical regularities at many levels of abstraction, simultaneously.

¹ By “levels of abstraction” we are broadly referring to the multiplicity of ways in which a cognitive system can represent a given object or experience: e.g. your pet could be “Rex”, a beagle, a dog, an animate being, a brown object etc.

² In the current paradigm, a stream of pictures is embedded with regularities that predict picture order—predictive regularities. If participants learn from this probabilistic environmental information, they should be able to distinguish picture orders that they observed from scrambled or foil orders of pictures.

For example, the predictive relationship between dogs and leashes exists based on abstract categories as well as in the actual objects or exemplars seen in the world (e.g. dogs have their specific leashes). The learning paradigms reviewed above do not reflect this important aspect of information that we receive from the world: information is often redundant across multiple levels of abstraction.

The current paper systematically investigates learning where participants are exposed to environmental regularities at multiple levels of abstraction. Do participants learn from the multiple levels of predictive dependencies simultaneously or are they biased to information at a certain level of abstraction? To address this question, we devised a novel statistical learning task where predictive regularities are learnable and redundant at multiple levels of abstraction. Specifically, participants were presented with sequences of new exemplars from known categories. Both the categories (e.g. dogs-fish, flowers-birds) and the individual exemplars of these categories (e.g. dog₁-fish₁, dog₂-fish₂) were predictive of picture order (see Figure 1). In two studies, we examined whether participants learn simultaneously based on both types of information or whether participants learn preferentially based on categorical or object-based regularities.

We believe that the current experimental design provides ample opportunity for learning at the abstract, categorical level. First, previous research has established that the categories used in the current experiment are initially processed at the basic-level (dog as opposed to the subordinate level of beagle; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) which is the level of categorical regularities of the picture stream. Second, we employed the same stimulus timing (short durations and long inter-stimulus-intervals) as employed by Brady and Oliva (2008) which will likely tap into the fast, gist-based recognition of the pictures. Finally, the stream has fewer pairs of categories than objects (see Figure 1). Thus, category level learning is, in some sense, easier than object-based learning.

Using the same methodology as Brady and Oliva (2008) as described above, pilot testing confirmed that when categorical regularities are predictive of picture order but individual objects or exemplars are not participants can learn based on categorical regularities: mean = 62.6%, $t(13) = 2.80, p < 0.05$. These results confirm that if object-based regularities are not present, category-level statistical learning is possible using the current stimuli and categories.

Finally, in order to more closely examine how learning proceeds at the categorical level of information, we manipulated the typicality of the exemplars that participants viewed: roughly half the participants were familiarized with typical exemplars of the categories and the rest were familiarized with atypical exemplars (see Appendix 1 for the atypical exemplars). Research has consistently shown that atypical exemplars are processed differently from typical exemplars (Dale, Kehoe, & Spivey, 2007) and tend to be more quickly processed below the basic-level categories (e.g. penguin as opposed to bird; Jolicoeur,

Gluck, & Kosslyn, 1984). Thus, we expect the participants familiarized with atypical exemplars to have weaker learning at the category-level but equivalent learning at the object or exemplar specific level. This typicality manipulation provides another way to examine performance for evidence of learning across different levels of abstraction.

Experiment 1: Testing for Object-Level Learning

The first experiment examines learning based on regularities of individual objects where both objects and object categories are predictive of picture order. Figure 1 illustrates a sample familiarization stream. We employed a testing procedure that is well-established in the statistical learning literature (e.g. Brady & Oliva, 2008; Fiser & Aslin, 2001): participants were asked to distinguish pairs of pictures from familiarization (e.g. bird₁-dog₁) from a foil pair created from the same pool of pictures but which violates contingency pattern of the familiarization stream. To isolate knowledge at the object-specific level, the foils were designed to violate object-based regularities while maintaining categorical regularities (e.g. bird₁-dog₂, see top panel of Figure 2). Thus, participants *require* object-level knowledge of the familiarization stream in order to distinguish the foils from the pairs. Given this experimental design, if participants are able to consistently distinguish pairs from foils, this is evidence for learning based on the objects and not the categories presented during familiarization.




Figure 1: A sample familiarization stream. Pictures were organized into pairs of categories (e.g. birds > dogs) as well as specific objects within these categories (e.g. robin > beagle). Thus, predictive regularities were redundant across multiple levels of abstraction resulting in two pairs of categories and eight pairs of objects or exemplars of these categories.

Methods

18 undergraduate students participated in Experiment 1 (age: mean = 20.7, std = 1.75; 2 left handed; 10F) and randomly assigned to each condition: 10 participants viewed the typical pictures, and 8 viewed the atypical pictures. All participants were from Cornell University, participated in exchange for course credit, and provided informed consent.

Familiarization A statistically-structured familiarization sequence was presented, using PsyScope X B53 on a MacMini computer with a 17in CRT monitor. Each picture was displayed for 300ms with a 700ms inter-stimulus interval (Brady & Oliva, 2008)

There were 4 categories of pictures: birds, dogs, fish, and flowers. For each category 4 different exemplars were used (dog₁, dog₂, etc.). The pictures were grouped into 8 pairs such that both the categories and the specific exemplars were predictive of picture order. For example, bird₁-dog₁ would always occur as a pair, as would bird₂-dog₂, bird₃-dog₃, and bird₄-dog₄. Thus, the familiarization stream contains multiple, redundant levels of predictive information: both the exemplar level and the more abstract category level of information are predictive. See Figure 1 for an illustration of the familiarization sequence. To ameliorate any effect of specific pairings on learning, different categories and object pairings were employed across participants. Participants saw each pair 28 times presented in random order without pairs repeating each other and were simply instructed to look at the pictures.

Testing After familiarization, the participants performed a test in which two pairs of pictures were presented sequentially: 700ms between pictures in the same pair, and 1200ms separating the pairs. One pair was from the familiarization (e.g. bird₁-dog₁), and one was a foil pair (e.g. bird₁-dog₂; see Figure 2). The foils were designed to violate the structure *only* at the exemplar level, and not the category level. Thus, participants require exemplar level knowledge of the familiarization stream in order to distinguish the foils from the true pairs. This test determines whether participants learn the familiarization sequence at the level of exemplars or specific figures or at the level of abstract categories. The participants were instructed to choose which of the pairs seemed more familiar, based on the familiarization task. No time constraint was imposed for their responses. There were 64 test trials.

After the experiment, the participants completed a survey in which they rated the pictures they had seen on a scale of 1-5 for interestingness and typicality. They were also asked to repeat the instructions of each task, to ensure they understood them correctly. Finally, they were asked whether they noticed any patterns during the familiarization sequence, to check for explicit knowledge of the sequence structure.

Results and Discussion

The current experiment was designed such that only exemplar specific knowledge could distinguish pairs seen during familiarization and foils. Performance was evaluated against chance (50%) for evidence of learning. Overall, participants demonstrate evidence of significant learning (mean = 72.7%; std = 23.2; $t(17) = 4.15$, $p < 0.0001$) indicating that participants acquired object-specific knowledge. See the bottom panel of Figure 2 for a graphical presentation of the results of this experiment.

12 participants reported evidence of explicit knowledge via the post-test questionnaires. The majority of these reports involved category level knowledge, some with knowledge of specific pairings within these categories (e.g. “particular flower with certain fish” and “maybe bird w/dog, flower w/fish”). A very small number of reports were exclusively at an object level (“white bird with white flower combo” and “black lab, sunflower, etc.”).

Data were submitted to an ANOVA examining the effects of exemplar typicality (Atypical vs. Typical) and explicit knowledge on test performance. Consistent with the findings mentioned in the introduction, we hypothesized that any contribution of categorical knowledge would be modulated by the typicality of the exemplars. We report no main effect of exemplar typicality ($F(1,14) = 0.307$; $p > 0.5$) nor interaction of typicality and explicit knowledge. The uniform performance across atypical and typical groups, as indicated in Figure 2, suggests no contribution from category-level knowledge in the current experiment.

We do, however, report a marginal effect of explicit knowledge of sequence structure ($F(1,14) = 3.96$; $p < 0.07$). We will address this issue more deeply in the results section of Experiment 2. Given that the current experiment was designed such that categorical knowledge could not be used to distinguish foils from pairs, and most evidence for explicit knowledge came as a report of predictive dependencies involving category level knowledge, it is unclear how explicit knowledge boosts performance. One possibility is that participants who achieve a high level of knowledge also achieve lexical access to the categories. Possibly knowledge of many of the pairs of exemplars induces category-level explicit knowledge.

In sum, participants were exposed to a sequence of pictures containing predictive dependencies redundant at the level of individual object and at a more abstract level of the categories these objects belonged to. Test performance indicates that participants gained object-specific knowledge. In addition, results suggest that participants do not acquire additional knowledge from more abstract, categorical regularities. We hypothesized that if participants do acquire categorical level knowledge, it would be modulated by object typicality. Results indicate no difference in learning between participants who received exposure to typical or atypical exemplars. Failing to find any difference between these groups suggests that participants learned from object-level regularities exclusively.

However, some participants do report explicit knowledge of the sequences. The majority of these reports included and sometimes were exclusive to abstract, object categories. These results indicate some awareness of the abstract properties of the stream. Moreover, we find that explicit knowledge has a marginally significant effect on test performance. It is unclear how explicit knowledge of this kind could aid in performance given that the experiment was designed to tap into object-specific knowledge only. Thus, while Experiment 1 provides strong evidence for object level learning, it does not entirely exclude the possibility that participants acquire some more abstract knowledge. In the second experiment, we more directly examine the possibility that participants learn from both categorical and object level predictive dependencies.




Figure 2: Top Panel: The sole difference between experiments was the composition of foils used at test. In Experiment 1, foils were designed to assess learning at the object or exemplar-specific level, while Experiment 2 foils allow for knowledge at both levels of abstraction (object and category) to influence test performance.

Bottom Panel: Results across Experiments 1 and 2 indicate no effect of exemplar typicality or foils on test performance.

Experiment 2: Testing for Additional Category-Level Knowledge

The current experiment addresses whether participants learn from the predictive dependencies at multiple levels of abstraction simultaneously (e.g. objects: $bird_1$ - dog_1 ; abstract, categories: $birds$ - $dogs$). To this end, we modified the foils used in Experiment 1 while keeping all other aspects of the experiment the same (e.g. $bird_1$ - dog_2). The foils in Experiment 1 violated the statistical regularities at the level of individual objects but preserved categorical regularities. Thus, object-specific knowledge but not category level knowledge would be essential in order to distinguish the pair from the foil.

In Experiment 2, we changed the foils to violate both object-level and category-level statistical regularities (e.g. $bird_1$ - $flower_3$). Therefore, category knowledge as well as object-specific knowledge could be used at test. If it were the case that participants learn from predictive dependencies at *both* levels of informational abstraction, we hypothesize that it would be easier to distinguish foils in the current experiment, which violate both forms of statistical regularities, compared to the foils used in Experiment 1, where only object-level regularities were violated. However, if participants do not acquire abstract knowledge during familiarization, they will still be able to perform the test in the same manner as Experiment 1. Thus, if participants acquire abstract knowledge, we hypothesize a significant increase in performance in Experiment 2 from Experiment 1, and failure to observe a significant increase in test performance would indicate that learning does not occur at the abstract categorical level.

As in Experiment 1, participants viewed either typical or atypical exemplars. If participants acquire categorical knowledge during familiarization, this knowledge will likely be modulated by the typicality of exemplars. In Experiment 1, we did not observe any asymmetry of performance between these groups; however, categorical knowledge would interfere with test performance in this case. In the current experiment, categorical knowledge would be of benefit. Thus, we hypothesize that, if participants have access to category level knowledge after familiarization, participants who view typical exemplars will have a greater boost in test performance than those who view atypical exemplars.

Methods

Another 24 participants were recruited from the same subject pool and randomly assigned to each condition (16F, 1 left handed, age: mean = 19.6, std = 1.28); 12 viewed the typical pictures, and 12 viewed the atypical pictures. The procedure in this experiment differed from Experiment 1 in only one respect: the foil pairs during the test were designed to violate the statistical structure of the familiarization sequence at the exemplar *and* the category level.

Results and Discussion

We report significant learning overall in Experiment 2 (mean = 67.7%, std = 21.2; $t(23) = 4.10, p < 0.0001$). The data were submitted to a two-way ANOVA to evaluate the effects of typicality of exemplars and explicit knowledge. There is a main effect of explicit knowledge ($F(1, 20) = 110.2; p < 0.001$) and, as seen in Experiment 1, we report no main effect of typicality ($F(1, 20) = 0.537, p > 0.5$) or interaction between these factors. We hypothesized that if categorical knowledge was acquired during exposure that it would be modulated by the typicality of exemplars. The consistent null effect of exemplar typicality indicates that participants do not acquire abstract, category level knowledge during exposure to environmental regularities at multiple levels of abstraction.

We also hypothesized that if participants learned from statistical regularities about objects as well as categories, participants in Experiment 2 would perform better at test than participants in Experiment 1. Results from both experiments were analyzed in a 3-way ANOVA, to test for effects of experiment, typicality, and explicit knowledge on test performance. This analysis confirmed the pattern of results seen in the bottom panel of Figure 2: there is no main effect of Experiment ($F(1,35) = 0.440, p > 0.5$). Additionally, we confirm that across both experiments there is no main effect of typicality of exemplars ($F(1,35) = 0.081, p > 0.5$) and no interaction between these factors. Thus, test performance is equivalent across experiments indicating that participants likely did not acquire categorical knowledge during exposure to the familiarization stream.

Consistent with results found in both experiments separately, there is a main effect of explicit knowledge: $F(1,35) = 35.9, p < 0.001$. Pooling participants from both experiments, we find that participants with explicit knowledge performed better than those without (mean performance: 85.3% vs. 54.4%). However, both groups performed significantly better than chance (with knowledge: $t(20) = 8.07, p < 0.001$; without knowledge: $t(20) = 2.21, p < 0.02$). Thus, regardless of explicit knowledge there is evidence for learning in both groups.

To determine whether explicit knowledge is related to any of our experimental manipulations (e.g. typicality of exemplars), we examined whether number of participants who demonstrate explicit knowledge is biased towards either a particular experiment (Exp. 1 or 2) or typicality of the objects seen. Of the 42 subjects in both experiments, 21 reported knowledge of the structure, while 21 reported no such knowledge. Chi-square tests show that the proportion of participants who had explicit knowledge of the sequence structure was not significantly different between any of the experimental factors: Experiment 1 vs. Experiment 2: $c2(1, N = 42) = 3.5, p > 0.05$; typical vs. atypical: $c2(1, N = 42) = 1.09, p > 0.25$. These results indicate that explicit knowledge, while a significant factor affecting performance, is equally distributed across groups and thus should not disproportionately bias overall performance.

Finally, all participants rated both typical and atypical pictures on “interestingness” and typicality. T-tests comparing ratings within categories revealed that participants rate atypical and typical exemplars distinctly and also rate the atypical exemplars as more interesting ($t(134) > 3.5; p < 0.001$ within categories for both typicality and “interestingness”). These results validate the assumption that participants view atypical and typical exemplars differently.

Along with Experiment 1, these results support the view that participants learn from statistical regularities at the lowest level of representational abstraction even when more abstract statistical regularities are available to any learning mechanism. Specifically, the results from Experiment 2 cast doubt on the possibility that participants learn from predictive regularities at both levels abstraction.

General Discussion

Humans are able to learn from experience where complex regularities are present. We investigated behavior in a novel learning task designed to investigate a key aspect of the complexity of daily experience: participants viewed streams of pictures with predictive dependencies at multiple levels of abstraction. Specifically, both individual objects or exemplars and the semantic categories that these objects belonged to predicted picture order, thus both object and categorical information could be used to determine the structure of the familiarization stream. We consistently find evidence for learning at the lowest level of abstraction: participants respond at test according to the predictive dependencies of specific objects or category exemplars and do not show evidence of having learned at the more abstract level of categories even when abstract knowledge could aid test performance. Moreover, we find no modulation of learning by exemplar typicality. These findings suggest that while participants can learn from regularities of categories, they do not learn from more abstract regularities when less abstract, more grounded statistical information is present.

Interestingly, while we systematically find that categorical knowledge has no influence on test performance, some participants acquire explicit knowledge of the categorical knowledge of the sequence. This result is strikingly similar to Brady and Oliva (2008): in their Exp. 3, after viewing streams with regularities present solely at the categorical level, participants were able to perform consistently in a test where pictures were replaced with category labels. In Exp. 4, Brady & Oliva (2008) include regularities at the scene specific or object level in addition to categorical regularities and again find evidence for lexical access. We argue that lexical access results in Exp. 3 and 4 of Brady & Oliva (2008) are similar to the demonstration of abstract level explicit knowledge in the current experiment.

While demonstration of lexical access to categories is interesting and important, we repeatedly show that abstract knowledge does not have a clear effect in test performance, raising questions about the nature and function of this lexical knowledge. To date, there has been no demonstration


of generalization, an important hallmark of abstract categorical knowledge, when less abstract regularities are present. This is an important avenue for future study to clarify this lexical result. An alternative possibility is that the lexical access is a byproduct of the participant strategy of using mental labels for the familiar objects and scenes as they are being presented. Previous statistical learning studies have been careful to avoid recognizable visual objects for this reason (Conway & Christiansen, 2005; Fiser & Aslin, 2001).

Despite our findings that adults did not learn from abstract, category level regularities when object-based regularities were present, it is nevertheless clear that in a natural environment we do acquire knowledge of higher order regularities. Thus, our results may simply point to a direction of how this learning occurs: learning starts in the relation of specific objects when statistical regularities are comparable at multiple levels of abstraction. One possibility is that when once the least abstract regularities have been mastered, learning can proceed along more abstract dimensions. Nevertheless, this finding may have important implications for more efficient teaching methods and could inform computational modeling of learning and development of human cognitive processes where the abstraction of representation is often an assumption built into the model.

Overall, this study aims to uncover how simple learning mechanisms operate in complex, naturalistic environments. We increased the complexity of the learning task, relative to previous experiments, by having predictive dependencies at multiple levels of abstraction. Results indicate that participants learned based on the more concrete, less abstract predictive dependencies. Results also suggest that participants did not additionally learn the more abstract relationships as this knowledge consistently did not influence test performance. These results inform the on-going debate as to whether domain-general learning mechanisms are largely unconstrained, as previous behavioral studies would have suggested. We believe that these results show some level of constraint on learning where more grounded, less abstract statistical relationships are learned preferentially when categorical and object specific knowledge is redundant.

Acknowledgments

Dr. Dima Amso, Dr. Rick Dale and Jordan DeLong for helpful conversations, Claire Schmidt for data collection and in particular, we thank Dr. Michael Spivey for his support.



Appendix 1: All atypical exemplars used in the current paper, organized by category (from left: dog, flower, fish, bird).

References

Baldwin, D. Anderrson, A. Saffran, J. & Myers, M. (2007). Segmenting dynamic human action via statistical structure. *Cognition*, 106, 1382-1407.

Brady, T. F. & Oliva, A. (2008). Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent, *Psychological Science*, 19, 678-685.

Conway, C. & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences, *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 31, 24-39.

Dale, R., Kehoe, C.E. & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory and Cognition*, 35, 15-28.

Fiser, J. & Aslin, R. A. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12, 499-504.

Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection, *Cognitive Psychology*, 16, 243-275.

Marcus, G. F., Fernandes, K.J. & Johnson, S.P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, 18, 387-391.

Saffran, J. R. (2002). Constraints on statistical language learning. *Journal of Memory and Language*, 47, 172-196.

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.