



Published in final edited form as:

Cognition. 2016 August ; 153: 63–78. doi:10.1016/j.cognition.2016.04.010.

Statistical Learning Is Constrained to Less Abstract Patterns in Complex Sensory Input (but not the Least)

Lauren L. Emberson^{1,2} and Dani Rubinstein^{3,4}

¹Brain and Cognitive Sciences Department, University of Rochester

²Psychology Department, Princeton University

³Department of Neuroscience, Brown University

⁴Clinical and Translational Neuroscience Branch, National Institutes of Health

Abstract

The influence of statistical information on behavior (either through learning or adaptation) is quickly becoming foundational to many domains of cognitive psychology and cognitive neuroscience, from language comprehension to visual development. We investigate a central problem impacting these diverse fields: when encountering input with rich statistical information, are there any constraints on learning? This paper examines learning outcomes when adult learners are given statistical information across multiple levels of abstraction simultaneously: from abstract, semantic categories of everyday objects to individual viewpoints on these objects. After revealing statistical learning of abstract, semantic categories with scrambled individual exemplars (Exp. 1), participants viewed pictures where the categories *as well as* the individual objects predicted picture order (e.g., bird₁—dog₁, bird₂—dog₂). Our findings suggest that participants preferentially encode the relationships between the individual objects, even in the presence of statistical regularities linking semantic categories (Exps. 2 and 3). In a final experiment we investigate whether learners are biased towards learning object-level regularities or simply construct the most detailed model given the data (and therefore best able to predict the specifics of the upcoming stimulus) by investigating whether participants preferentially learn from the statistical regularities linking individual snapshots of objects or the relationship between the objects themselves (e.g., bird_picture₁—dog_picture₁, bird_picture₂—dog_picture₂). We find that participants fail to learn the relationships between individual snapshots, suggesting a bias towards object-level statistical regularities as opposed to merely constructing the most complete model of the input. This work moves beyond the previous existence proofs that statistical learning is possible at both very high and very low levels of abstraction (categories vs. individual objects) and suggests that, at least with the current categories and type of learner, there are biases to pick up on statistical regularities between individual objects even when robust statistical information is

This manuscript version is made available under the CC BY-NC-ND 4.0 license.

Correspondence to: Lauren L. Emberson, Peretsman-Scully Hall, Psychology Department, Princeton University, Princeton, NJ, 08540 USA, Phone: (607) 342-1690, Fax: (609) 258-1113, laurenle@princeton.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

present at other levels of abstraction. These findings speak directly to emerging theories about how systems supporting statistical learning and prediction operate in our structure-rich environments. Moreover, the theoretical implications of the current work across multiple domains of study is already clear: statistical learning cannot be assumed to be unconstrained even if statistical learning has previously been established at a given level of abstraction when that information is presented in isolation.

Keywords

Psychology; learning; statistical learning; implicit learning; cognitive development; pattern recognition; vision; perception; perceptual learning; object perception; categorization

Introduction

Throughout our lifetime, experience shapes our mental model of the world; learning from patterns, regularities, or statistics in the environment is a way that experience might powerfully and incidentally shape cognition. Rich mental models or representations of the environment in turn support the prediction of upcoming sensory input facilitating both rapid and accurate cognitive processing when these predictions are correct, or error-signals and further learning and/or adaptation when these predictions are incorrect.

Indeed, statistical learning¹, or the sensitivity to statistical information in sensory input, is increasingly considered to form the foundation of numerous and diverse perceptual and cognitive abilities. For example, not only is statistical learning believed to be an essential component of the development of language (Saffran, Aslin & Newport, 1996; Romberg & Saffran, 2003; Thiessen & Erickson, 2013), but researchers have increasingly turned to mechanisms of adaptation, or sensitivity to statistical or distributional patterns in the environment, to explain adult comprehension (e.g., Fine & Jaeger, 2013; Kleinschmidt & Jaeger, accepted). Likewise in vision, there have been many demonstrations that the visual system is sensitive to statistical information both immediately after exposure (e.g., neural demonstrations: Summerfield et al., 2008; Turk-Browne et al., 2009; 2010) and cumulatively over a lifetime of visual experience (Purves, Wojtach & Beau Lotto, 2011; Purves, Monson, Sundararajan & Wojtach, 2014). Indeed, even young infants have the ability to perform complex visual statistical learning tasks (as young as 2-months-old for sequential visual statistical learning tasks; Kirkham, Slemmer & Johnson, 2002; Fiser & Aslin, 2002 for spatial statistical learning tasks with 9-month-olds)². The more the field searches, the more it finds that the sensory environment is filled with statistical information. Paired with the

¹This paper focuses on “statistical learning,” which is typically studied with participants passively learning from stimuli endowed with probabilistic, distributional or statistical information (though also see Hunt & Aslin, 2001; Perruchet & Pacton, 2006 for related behavioral paradigms). However, many, and perhaps all, types of learning can be considered *statistical* (i.e., diverse learning tasks can be commonly characterized by the patterns or statistical information presented). In addition, the ability to incidentally learn from statistical regularities in the environment has been found to recruit multiple classically-defined learning and memory systems (see Emberson & Amso, 2012; Karuza et al., 2012; Turk-Browne, Scholl, Chun & Johnson, 2008; Turk-Browne, Scholl, Johnson & Chun, 2010: the hippocampus, basal ganglia and inferior frontal cortex are involved in learning from statistical regularities). Thus, we consider statistical learning to be representing a general and essential aspect of learning: the ability to pick up on statistical information in the environment and use it to shape one’s behavior.

²Unlike the numerous studies linking statistical learning to language development, additional studies are needed to more directly link visual statistical learning to visual development.

rapid and ubiquitous statistical learning abilities of both infants and adults in many different types of input, it is clear that understanding cognitive responses to statistical information is foundational to understanding many aspects of perception and cognition.

However, despite its clear importance, little is known of how statistical learning proceeds in our daily environments. A central problem is that it is unknown how statistical learning mechanism(s) operate given the richness of the statistical information that we routinely receive. In this paper, we focus on a particular type of *rich* information: **what happens when statistical information is present across multiple levels of abstraction?** How does learning proceed when an observer can learn that not only do categories of objects predict another, but also that individual objects within categories predict each other? This kind of overlapping statistical information is routinely presented to us. For example, the predictive relationship between dogs and leashes exists based on abstract categories as well as on the actual objects or exemplars seen in the world (e.g. dogs have their specific leashes). There are many other examples where both categories of objects and also the individual objects within these categories are predictive of each other: types of animals and their habitats (e.g., fish and bodies of water), chairs and tables, people and their personal affects (e.g., wallets, cell phones). Thus, when two groups of objects are paired in the world, there is visual statistical information from which one can learn, all the way from low-level visual similarities across exemplars within a category, to relationships between individual objects, to their semantic, abstract categories. While the current paper focuses on visual statistical learning, similar inputs are found in other domains; for example language input, where there is overlapping statistical information at the levels of phonology, lexical items, and syntactic information when one rehearses a single utterance.

The current paper addresses what is learned when there is predictive visual information across levels of abstraction. Is statistical learning *a priori* constrained to learn particular patterns? There are two types of competing theories on this: The first proposes that statistical learning is largely unconstrained; The second proposes that statistical learning is inherently constrained to learn from particular types of regularities.

By and large, previous research has supported the view that learning mechanism(s) are largely unconstrained: Statistical learning has been demonstrated in multiple sensory modalities (e.g., Conway & Christiansen, 2005) and across a wide range of perceptual input. For example, in the visual modality, learning can occur from sequences of gestures (Baldwin, Andersson, Saffran, & Myers, 2007) or abstract shapes (Fiser & Aslin, 2001) and from spatial and non-spatial information (Mayr, 1996). While the majority of these studies have focused on learning probabilistic relations between individual objects, there is evidence that learning can occur at higher levels of informational abstraction including based on categories of nonsense words (Saffran, 2002; Reeder, Newport & Aslin, 2013) or familiar semantic categories (Brady & Oliva, 2008). Overall, these studies, among many others, have led to the belief that statistical learning is largely unconstrained. That is, if there is any reliable probabilistic information in the environment, humans can learn from it regardless of modality or level of abstraction.

If statistical learning were unconstrained, how would learning proceed when a learner is faced with statistical information across multiple levels of abstraction simultaneously? A logical prediction is that an observer would learn across these multiple levels of representation *simultaneously* and in parallel. While this topic has been largely unexamined, there is a handful of evidence that learning can interact across levels of representation: Onnis, Waterfall and Edelman (2008) found that the narrative structure of successive utterances supported statistical learning of individual syllables in adult learners; similarly, Saffran (2001) found that syntactic information in the form of phrase structures aided learning at lower levels (i.e., learning across individual elements), and Koch & Hoffmann (2000) suggested that different levels of information can compete in a serial reaction time (SRT) task.

However, the evidence that statistical learning can occur for all types of input and all levels of abstraction has been obtained under very constrained experimental conditions which may not reveal how learning operates over the rich statistical input encountered “in the wild”. To illustrate, Brady and Oliva (2008) use a paradigm where the categories of scenes are predictive of picture order but individual scenes are not (e.g. beaches predict kitchens as categories of scenes but beach₁ does not predict kitchen₁). Participants show evidence of learning the relationship between abstract categories of scenes, but it is not possible to learn based on individual scenes because only category-level regularities are present. While studies such as Brady and Oliva (2008) provide essential existence proofs, statistical learning has not been investigated in the context of the rich input that an everyday learner encounters, where multiple types of statistical regularities are present simultaneously.

The alternative theory is that statistical learning does have some constraints and that these constraints bias what is learned when presented with rich input. These theoretical views come in two types: The first type directly concerns the question of domain-generalizability and asserts that learning is biased towards a particular type of input (e.g., auditory vs. visual; linguistic vs. non-linguistic).³ The question of domain-generalizability is not directly addressed in the current paper, as we do not compare learning across types of sensory input but learning across multiple levels of abstraction within largely the same sensory input. The second family of theories, directly addressed here, asserts that there is a bias in the types of statistical information that will be learned within a given input. Newport (1990) and Elman (1993) suggest that with language learning “less is more” and learners “start small,” respectively. Both of these views suggest that the most tractable patterns, or those connecting smaller units in the input, will be learned first. A recently emerging view considers the relationship between prediction and statistical learning and proposes that the goal of statistical learning is to acquire the most complete model of the environment and to enable the best possible prediction of the upcoming stimuli (Emberson, Liu & Zevin, 2013; Karuza, Farmer, Fine, Smith, and Jaeger, 2014; Lupyan, forthcoming). These views are

³Indeed there has been some evidence that statistical learning is not equivalent across stimulus modalities (Conway & Christiansen, 2005; Conway & Christiansen, 2006; Emberson, Conway & Christiansen, 2011) or types of sensory stimuli (e.g., Gebhart, Newport & Aslin, 2009 for the slow amount of time to learn from complex, non-musical, non-speech auditory stimuli compared to speech and tonal melodies). While some have argued that biases in learning according to the type of input is evidence for domain-specificity, by and large, experimental evidence suggests that statistical learning is domain-general and can be accomplished with a wide variety of stimuli. Indeed, recent theoretical work has argued both of domain-generalizability and modality-specificity (Frost, Armstrong, Siegelman & Christiansen, in press).

grounded in the theory of the brain as a predictive system with the goal of minimizing prediction error (e.g., predictive coding; Clark, 2013; Friston, 2005; reinforcement learning, associative learning and/or conditioning; Rescorla, 1988, Schultz, Dayan, & Montague, 1997). These views both provide a clearer prediction of what would be learned when encountering rich statistical input: in the former case, that learning would be constrained to either the smallest available units or simplest patterns in the input; in the latter case, the learner would learn from the statistical information that would provide the best possible prediction of future sensory input.

The current paper examines learning when participants are exposed to statistical information at multiple levels of abstraction. Investigating how learning proceeds over statistically rich input will allow some dissociation between constrained and unconstrained views of statistical learning. Specifically, do participants learn from the multiple levels of predictive dependencies simultaneously, as would be predicted by unconstrained theories of statistical learning, or are learners biased to learn at a certain level of abstraction? Furthermore, if learners are constrained, do they learn from the smallest and most tractable patterns and/or do they acquire representations that will best allow them to best reduce their prediction error for upcoming sensory input?

To address these questions, we devised a novel statistical learning task where predictive regularities are learnable and redundant across two levels of abstraction and investigated these questions in the domain of visual statistical learning of familiar objects and semantic categories. For example, in a stream of passively viewed pictures, the semantic categories (dogs, fish, flowers, birds) as well as the individual exemplars of these categories (e.g. dog₁, fish₁, dog₂, fish₂) were predictive of picture order (see Fig. 1). In other words, there was learnable statistical information at both the *category-level* and the *object-level* of the familiarization stream. We examined whether participants learned from both the statistical information linking the semantic categories as well as the individual objects, or whether they were biased to learn at a particular level of abstraction and if so, which level they are biased to.

Experiment 1: Establishing Category-Level Learning with Shuffled Exemplars

Before examining learning in the presence of two levels of abstraction, we first sought to replicate and extend the finding from Brady and Oliva (2008): namely that there is category-level learning when individual objects are not predictive of each other. In other words, we sought to find evidence of learning at the level of abstract semantic categories when individual objects are not predictive of picture order. We extended the previous findings from Brady and Oliva (2008) by employing new categories and new pictures. Similar to the categories employed in previous work, the categories used in the current experiment are initially processed at the basic-level or the level of the categorical regularities in the current task, rather than the subordinate level (e.g. basic level of dog as opposed to the subordinate level of beagle; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Following Brady and Oliva (2009), we chose categories that were already well-learned and readily recognized

by our participants (dogs, birds, flowers, and fish in the current study; kitchens, beaches, etc. in Brady & Oliva, 2008).

In addition, we examine a logical extension of the claim that participants are using statistical information to learn the structure of abstract categories: that their learning will be modulated by the typicality of the category exemplars. Specifically, participants were exposed to either typical or atypical exemplars (Fig.1, Fig.2, respectively). Typicality effects are considered “the strongest and most reliable effects in the categorization literature” (Murphy, 2002, p.22) with object typicality strongly modulating many categorization tasks from reaction time (RT) effects to ratings to learning new categories (see Murphy, 2002 for a comprehensive review, also Dale, Kehoe, & Spivey, 2007, Jolicoeur, Gluck, & Kosslyn, 1984). The importance of typicality effects in conceptual processing is so important that “[not finding] an effect of typicality ... suggests that there is something wrong with—or at least unusual about—the experiment.” (Murphy, 2002, p.24). Thus, finding an effect of typicality will strongly support that participants are learning at the category-level of representation.

Methods

All studies participants were either Cornell University or University of Rochester undergraduates who took part in exchange for course credit or money and provided informed consent consistent with the Cornell or University of Rochester Institutional Review Boards. The 16 participants in Experiment 1 (age: mean = 21.5, std = 1.06; 1 left handed, 11F) were randomly and evenly assigned to either the typical or atypical condition.

Familiarization—A statistically-structured familiarization sequence was presented, using PsyScope X B53 on a MacMini computer with a 17in CRT monitor or Psychtoolbox 3.0.10 for MATLAB R2013a on a 17in flat screen 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. Participants viewed either typical or atypical exemplars of this category. Exemplars (4 typical, 4 atypical for each category) were selected, based on pilot testing, to be both recognizable for their semantic category and typicality or atypicality, from an initial set of 12 candidate pictures per category. Stimuli are available at the first author’s website and as supplementary materials for this manuscript.

Typicality was counterbalanced across participants. Each participant viewed either the 4 typical or the 4 atypical exemplars for each category (bird₁, bird₂, etc., a total of 16 pictures). For each participant, the 4 categories were randomly grouped into two pairs (e.g., birds–dogs, fish–flowers) with the individual exemplars shuffled across presentation (e.g., bird₁–dog₄, bird₂–dog₁, bird₄–dog₄). Thus, the statistical information at the level of category is highly predictive of picture order (transition probabilities between category pairs are 100%) but the statistical information at the level of individual exemplars or objects is not predictive of picture order (transitional probabilities between pictures within a category pair are 25%). All individual exemplars or pictures were presented with equal probability. Simply instructed to look at the pictures, participants saw each category pair 112 times in random order.

Testing—After familiarization, participants were tested by comparing two pairs of pictures presented sequentially, with 700ms between pictures in the same pair, and 1200ms separating the pairs. One pair was from the familiarization (e.g. birds-dogs), and one was a foil that violated the statistical information from the familiarization stream by mixing categories (e.g. birds-flowers). As with familiarization, all individual exemplars were shuffled across presentations and presented with equal frequency. The participants were instructed to choose which of the pairs seemed more familiar, based on the familiarization task, with no time constraint imposed. There were 64 test trials. After the experiment the participants completed a survey in which they rated both the typical and atypical pictures they had seen on a scale of 1–5 for “interestingness” and typicality, completed a post-test questionnaire about their test strategies and familiarity of the different categories of objects, and were debriefed. The post-test questionnaire included the question “Did you notice any patterns in the initial stream of pictures?” See Appendix 2 for a discussion of explicit knowledge.

Results and Discussion

Overall participants showed robust evidence of learning from the statistical regularities in the exposure stream. Compared to chance performance of 50%, mean performance was 72.7% (std = 25%; $t(15) = 3.67$, $p = 0.002$, see Fig.4). This finding suggests that statistical learning of abstract categories is a robust phenomenon extending to 4 previously uninvestigated categories.

We had additionally hypothesized that if participants are learning based on the statistical regularities of abstract semantic categories, they would show additional modulation of learning based on exemplar typicality. To this end, we divided participants between those exposed to typical and atypical exemplars and compared their performance at test. We continue to find robust evidence of learning in participants exposed to the typical exemplars (mean performance = 87.1%, std = 22%, $t(7) = 4.98$, $p = 0.001$) but there was no evidence of learning in participants exposed to the atypical exemplars (mean performance = 58.2%, std = 20%, $t(7) = 1.18$, $p = 0.27$). Moreover, an unpaired t-test revealed that atypicality significantly reduced learning in this task ($t(13.9) = 2.84$, $p = 0.013$). Furthermore, the size of this typicality effect indicates that it is extremely robust (Cohen’s $d = 1.16$). See Table 1 for summary of findings and Appendix 1 for reaction time analyses and Figure 4. The finding that exemplar typicality modulates learning outcomes in this task provides additional evidence that participants are indeed learning based on the statistical regularities across abstract semantic categories.

We also confirm that the typicality ratings differed across typical and atypical exemplars (typical exemplars: $M = 4.0$, range 3.1–4.4; std = 0.42; atypical exemplars: $M = 2.5$, range 1.7–3.1; std = 0.43). Mean typicality ratings between typical and atypical exemplars did not overlap within each category, and were significantly different overall ($t(30) = 9.79$, $p < 0.001$), confirming that participants perceived the typical and atypical exemplars as expected.

Experiment 2: Testing for Object-Level Learning

Building on Experiment 1, we examine whether participants continue to learn from statistical regularities between semantic categories when more concrete object-specific regularities are present. In a small change from the previous study, participants view streams of pictures where the individual objects or pictures, as well as the categories, are predictive of picture. To illustrate, while in Experiment 1 pictures of birds predicted pictures of dogs but no individual bird predicted any individual dog, in this experiment pictures of birds would still predict pictures of dogs but individual birds also predict individual dogs (e.g., $\text{bird}_1\text{-dog}_1$, $\text{bird}_2\text{-dog}_2$, $\text{bird}_3\text{-dog}_3$, see Figure 1). After exposure, participants are asked to distinguish pairs of pictures from familiarization from a foil pair that violates object-based regularities while maintaining categorical regularities (e.g. comparing $\text{bird}_1\text{-dog}_1$ with $\text{bird}_1\text{-dog}_2$, see top panel of Fig. 3). Since participants *require* object-specific knowledge of the familiarization stream in order to distinguish the foils from the pairs, the ability to distinguish pairs from foils would be evidence that participants can learn based on object-specific regularities.

Having established in Experiment 1 that participant learning was strongly modulated by the typicality of the exemplars that they saw, we again expose half of our participants to the same typical and half to the same atypical exemplars. Manipulating typicality provides a second yet well-established method to probe learning at the level of category; using typicality as a litmus test for category-level learning. In the current experiment, with individual exemplars or objects predictive as well as categories, we expect that if participants are still learning based on the categories of the objects that their learning will still be modulated by the typicality of the exemplars that they were exposed to. However, if they are learning based on the objects or exemplars themselves, we do not expect learning to be modulated by typicality.

Methods

The 22 participants in Experiment 2 (age: mean = 20.7, std = 1.75; 2 left handed; 10F) were randomly and evenly assigned to either the typical or atypical condition.

Familiarization—As with Experiment 1, there were 4 categories of pictures: birds, dogs, fish, and flowers. For each category, 4 different exemplars were used (dog_1 , dog_2 , etc.) with typical and atypical exemplars counterbalanced across participants. The pictures were then grouped into 8 pairs such that both the categories and the specific exemplars were predictive of picture order. For example, $\text{bird}_1\text{-dog}_1$ would always occur as a pair, as would $\text{bird}_2\text{-dog}_2$, $\text{bird}_3\text{-dog}_3$, and $\text{bird}_4\text{-dog}_4$ (Fig. 1). To control for any effect of specific pairings on learning, different categories and object pairings were employed across participants. Simply instructed to look at the pictures, participants saw each pair of objects 28 times in random order without repetition. This results in 112 repetitions of each pair of semantic categories, the equivalent exposure to Experiment 1.

Testing: The testing procedure was identical to that of Experiment 1 with the exception of stimuli: one pair from familiarization (e.g. $\text{bird}_1\text{-dog}_1$) was presented with one foil that only

violated the object-specific level of statistical regularities but maintained the relationship between the categories of the objects (e.g. bird₁-dog₂; Fig. 3).

Results and Discussion

The current experiment was designed such that only object-specific knowledge could distinguish pairs seen during familiarization and foils. Overall, participants reliably distinguished pairs from foils compared to chance (mean = 75.4%; std = 22.8; $t(21) = 5.24$, $p < 0.0001$) indicating that participants learned object-specific regularities even though more abstract, category-level regularities were also present (Fig. 4 but also Table 1).

We hypothesized that any contribution of categorical knowledge would be modulated by the typicality of the exemplars as it was in Experiment 1. As in Experiment 1, participants' typicality ratings were different across types of exemplars (atypical mean = 2.7, std = 1.4; typical mean = 3.8, std = 1.4; $t(30) = 4.00$, $p < 0.001$). However, unlike Experiment 1, we did not find an effect of typicality on learning outcomes. We report no effect of exemplar typicality at test with participants performing well whether they were familiarized with typical or atypical exemplars (Atypical mean = 77.0%, std = 21.5; Typical mean = 73.9%, std = 24.9, $t(21) = 0.315$, $p = 0.76$). Importantly, given the effect size of the typicality effect observed in Experiment 1 ($d = 1.16$) and the current sample size, the current t-test has a power of 0.74 (R library: pwr; Champely, 2012). The lack of typicality effect suggests that participants do not additionally learn at the more abstract category-level and only learn based on the object-specific regularities (and see Appendix 1 for Reaction Time analysis). However, Experiment 3 lends further evidence to this claim.

Experiment 3: Testing for Additional Category-Level Knowledge

Having established that participants can learn from object-specific statistical regularities in the presence of more abstract category-level statistical regularities, Experiment 3 examined whether learning occurs along multiple levels of abstraction simultaneously (i.e., objects: bird₁-dog₁; semantic categories: birds-dogs) or whether participants preferentially learn based on statistical regularities of objects as suggested by the lack of typicality effect in Experiment 2. To this end, we changed the foils used at test so that *both* object-specific and category-level patterns could be used to distinguish pairs seen during familiarization from foils (e.g. bird₁-dog₁ vs. bird₁-flower₃; Fig. 3). We hypothesized that if participants learn from regularities at *both* levels of abstraction, test performance should increase in Experiment 3 compared to Experiment 2 where only object-specific knowledge could be used at test. Conversely, failure to observe a significant increase in test performance suggests that learning does not occur at the abstract categorical level in addition to the object-specific knowledge demonstrated in Experiment 2.

Again, participants viewed objects that were either typical or atypical for their basic-level categories. In Experiment 2, we did not observe any effect of typicality. Given the strong typicality effect seen in Experiment 1, not finding a typicality effect in Experiment 2 suggests that participants are not learning at the category-level now that object-specific regularities are present.

However, in Experiment 2 categorical knowledge would have interfered with test performance. In the current experiment, categorical knowledge would be of benefit. Thus, we hypothesized 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 (16F, 1 left handed, age: mean = 19.6, std = 1.28, typical = 12) and randomly assigned to view either typical or atypical exemplars. The procedure in this experiment differed from Experiment 2 in only one respect: the foil pairs during the test violated the statistical structure of the familiarization sequence at the exemplar *and* the category level (Fig. 3).

Results and Discussion

We find significant learning overall in Experiment 3 compared to chance performance (mean = 67.7%, std = 21.2; $t(23) = 4.10$, $p < 0.0001$; Fig. 4). Again, we did not find any evidence that the typicality of exemplars affected learning outcomes: a t-test comparing performance based on object typicality revealed no effect of object typicality (atypical mean = 65.4%, std = 19.8; typical mean = 70.1%, std = 23.1; $t(23) = 0.53$, $p = 0.60$, Table 1). Based on the effect size of the typicality effect observed in Experiment 1 and the current sample size, this test has a power of 0.78. As in the previous experiments, typicality ratings confirmed that participants perceived the exemplars as expected (atypical mean = 2.5, std = 0.55; typical mean = 4.0, std = 0.45; $t(30) = 8.79$, $p < 0.001$).

We also hypothesized that if participants learned from both category-level and object-specific regularities, Experiment 3 performance would increase compared to Experiment 2. However, the means are numerically in the opposite direction of these predictions: average test performance in Experiment 2 is 75.4% and performance dropped in Experiment 3 to 67.7%. Results from both experiments were analyzed in a logistic regression considering the factors of Experiment (2 vs. 3), Typicality, and their interaction. This analysis revealed that the decrease in performance between Experiments 2 and 3 was significant ($\beta = -0.164$, $z(1,42) = -2.02$, $p = 0.043$). Note again that the means moved in the opposite direction from what would be predicted if participants were learning categorical information.

We hypothesized an increase in performance if participants are able to use both category-level and object-specific cues, and no change in performance if only object-specific information was present. However, a decrease in performance across these two experiments was unanticipated. One possibility is that participants learned exclusively based on object-level regularities and since the test added the presence of a second source of information that was not learned (i.e., category-level regularities), this resulted in greater difficulty at test. While future empirical work would be needed for this interpretation to be conclusive, one possibility for the reduced performance across Experiments 2 and 3 would also be consistent with the predominant learning of object-level regularities. However, another possibility for the decline in performance is simply between-subject variance. The comparison between foil-types was applied across participants in a non-randomized fashion, thus, this drop in

performance could be attributable to non-randomized between-subject variance. Indeed, future work would benefit from a within-subjects comparison across these types of foils.

Additionally, we confirm that across both experiments there is no main effect of typicality of exemplars and no interaction between these factors ($z_s < 0.97$, $p_s > 0.3$). Since we know the effect size of the typicality effect in Exp. 1, we can more precisely determine the power of the current test: for a model with 2 predictors and a sample size of 46, the current test is very well powered (power = 0.99). Thus, test performance is statistically equivalent across experiments and across typical and atypical exemplars, indicating that participants likely did not acquire category-level knowledge during exposure to the familiarization stream.

Overall, these results demonstrate that violating both object-specific and category-level regularities learned during familiarization does not boost test performance compared to a test where only object-level regularities are violated. These findings suggest that participants do not learn from predictive regularities at the category-level, in addition to the less abstract, object-specific patterns found consistently across both experiments. However, part of this argument arises from a null effect: no typicality effect across both Experiments 2 and 3. While we have demonstrated that these tests are well-powered, it cannot be ruled out that participants might acquire subtle category-level information in the current task that is not picked up by the current tests.

We investigated two different signatures of category-level knowledge across Experiments 2 and 3 in well-powered tests and find no evidence of abstract learning (i.e., difference in performance between Experiments 2 and 3, and typicality effects robustly found in Experiment 1). Future work could build from this foundation and investigate another hallmark of categories: Generalization. The question of whether learners can generalize during statistical learning or at test, from what they have learned in a statistical learning paradigm, has been the focus of much theoretical debate (e.g., Marcus, Vijayan, Rao & Vishton, 1999, review by Aslin & Newport, 2012). It is possible that phenomena like rule learning are simply the ability to do statistical learning over abstract categories as demonstrated in Experiment 1 as well as Brady and Oliva (2007), and would be an interesting avenue for future work. However, investigations of generalization at test should be careful to be able to dissociate whether participants are applying category level knowledge, gained during learning, to new exemplars or simply generalizing at test based on object-level knowledge. These are importantly different mechanisms: The latter arises from an extrapolation to known categories from previously learned object-pairings and importantly is not evidence of category-level learning, and the former arises from the learning based on categories which contain the property of generalization. These different mechanisms will be difficult but essential to disentangle in this line of work. Thus, in addition to evidence presented in the current paper, future work will be beneficial to provide additional evidence that participants are not learning based on category-level regularities at test.

What is clear is that participants' learning performances are predominantly driven by object-specific regularities even when category-level information is still present. This is certainly not evidence that there could not be other circumstances when category-level information

may arise even in the presence of object-specific information. For example, as the number of objects for each category increases, the memory demand on participants will also increase and may possibly result in category-level learning. Another possibility is that when participants are presented with novel objects and categories, participants will also exhibit a bias towards category-level learning. Finally, it is possible that patterns of learning will be different at different stages of development: Would infants who do not have strong lexical knowledge exhibit stronger category-level learning? These are all important avenues for future work.

Recent work by Turk-Browne and colleagues have examined neural responses during statistical learning tasks, at times, using the semantic categories of face and scene in their tasks (e.g., Kim, Lewis-Peacock, Norman & Turk-Browne, 2014; Turk-Browne, et al, 2010). This previous work differs from the current experiments in that category is not consistently manipulated throughout the task (e.g., faces predict houses and houses also predict faces) making it difficult to consider how category and specific exemplars play a role in this task. However, an interesting and potentially fruitful area of future research would be to investigate whether neural responses to overlapping category and exemplar information mirror the current behavioral information, or whether there are differences and the brain retains category-level predictions even when behavior appears to be dominated by object-specific regularities.

Experiment 4: Learning from Snapshots of Individual Objects

The findings from the previous experiments suggest that participants learn preferentially based on statistical regularities of individual objects. Indeed, when object-specific regularities are present, participants appear to be biased toward learning these regularities over regularities of semantic categories. However, the evidence is ambiguous as to whether they are preferentially learning based on individual objects or whether participants are learning based on the most specific statistical regularities possible (e.g., the ones that will best predict the upcoming sensory input). The final experiment in this paper is designed to tease apart these two possibilities. We examined whether participants learn based on the individual *snapshots* of objects or at a higher level of abstraction, based on the objects themselves. Similar to Experiments 2 and 3, individual objects were paired and the individual snapshots of these objects were also predictive of picture order. Snapshots varied in viewpoint and/or state of the object. For simplicity, we will refer to these differences as different *pictures* or *snapshots* of an object. As an example, a familiarization stream presenting a bird-dog pair would be instantiated as bird_picture₁-dog_picture₁, bird_picture₂-dog_picture₂. Importantly, the dog-bird pair is presented at a higher level of abstraction than the individual snapshots of the dog and bird. All participants received the same familiarization. At test, participants compared familiar pairs either against foils that violated only the individual snapshots of the objects (e.g., bird_picture₁-dog_picture₁ vs. bird_picture₁-dog_picture₂, similar to Experiment 2), or against foils that violated statistical regularities of the paired objects (e.g., bird_picture₁-dog_picture₁ vs. bird_picture₁-flower_picture₂, similar to Experiment 3). See Figure 6 for a depiction of the method of this experiment. Note that unlike the previous 3 experiments, we are not manipulating typicality of objects in this experiment.⁴

Methods

An additional 16 participants were recruited (mean age = 21, std = 1.09, all right handed, 12 female). All aspects of the methods were the same as in Experiments 2 and 3 with the exception of changes in the pictures viewed. Sixteen new pictures were selected such that 4 snapshots of the same object were presented for 4 objects. These objects belong to the same semantic categories employed in the previous experiments (Figure 5). The familiarization stream was organized such that individual objects predicted each other but the individual pictures of these objects also predicted each other (Figure 6). This structure of familiarization is analogous to Experiments 2 and 3 but with the levels of objects shifted down in abstraction from semantic categories and individual objects (Exp. 2 and 3) to individual objects and specific views of these objects in the current experiments. As with Experiments 2 and 3, the 8 pairs of objects and individual views were randomized across participants and received the same amount of exposure. After familiarization, participants compared pairs from familiarization to one of two foil types: either foils which violated the snapshot or picture level statistical regularities but not the object-level regularities (similar to Exp. 2), or foils which violated both object-level and snapshot level regularities.⁵

though building from the previous three experiments, we do not anticipate a significant contribution of category-level learning as a result of this experiment's familiarization stream.

Results

While overall participants showed evidence of learning (mean performance = 62.9%, std = 19%; $t(15) = 2.74$, $p = 0.015$, Fig.7), we find large differences in learning depending on the type of foil participants were presented with at test. When participants were asked to compare familiar pairs against foils that changed which objects were being paired, they showed evidence of robust learning (mean performance = 72.8%, std = 21%; $t(7) = 3.13$, $p = 0.017$, bottom foil type in Figure 6, Fig.7). However, if participants were asked to evaluate the same pair with a foil that only changed the views of the paired objects (therefore isolating the viewpoint specific information), participants showed no evidence of discrimination (mean performance = 52.9%, std = 10%; $t(7) = 0.806$, $p = 0.45$, top foil type in Figure 6, Fig.7). There is a significant difference in learning outcomes across foil types ($t(10.26) = 2.44$, $p = 0.034$, Table 1). See Appendix 1 for reaction time analysis.

⁴Since typicality effects are a hallmark of category-level processing, we employed typicality to provide convergent evidence as to when participants are learning at the category-level vs. learning at the object-level. We are not probing learning at the category-level in this experiment but rather at less abstract levels of representation (object-level and snapshot-level) and thus typicality effects would not provide additional information. While it is still possible that typicality could affect learning at more specific levels of representation (e.g., typical views or states of an object) and could be an interesting area for future work, these kinds of effects are not germane to the current study.

⁵It is important to note that the latter foil types also violate the basic-level categories. In order to include foils that did not violate basic-level categories one would have to introduce a second object for each basic-level category (e.g., snapshots of two different dogs) during familiarization which will require changing the overall statistical structure of the familiarization retained across Experiments 2–4 but could be an interesting avenue for future study. Nevertheless, the current experimental design does leave open the possibility that participants are using category-level knowledge and/or object-level knowledge to discriminate this foil type at test. Building from the results from Experiments 2 and 3, it seems unlikely that participants will be employing category-level knowledge in this task. Specifically, in two different measures, we fail to find evidence that category-level knowledge affects performance at test. If anything, we found that the inclusion of category-level information lead to a significant decrease in performance at test. Thus, if category-level information is having any effect here it is likely to decrease test performance which would work against any evidence of object-level learning in the current experiment. However, we find strong evidence of object-level learning in Experiment 2 suggesting that this type of learning is present in the current experiment.

Discussion

These results clearly support the view that participants learn based on the statistical regularities of individual objects and not based on the most specific statistical information present in the familiarization stream - the snapshots of individual objects.

The inability of participants to learn based on individual pictures or snapshots of the objects is surprising given that participants viewed each individual pair of pictures 28 times over the course of exposure. Importantly, while different snapshots of the objects are subtle in their visual differences, previous work has shown that adult participants robustly represent these subtle visual details in visual memory. Brady, Konkle, Alvarez and Oliva (2008) examined adults' visual memory of objects after exposure to 2,500 distinct pictures over 5.5 hours. Despite this incredibly demanding task, participants remembered subtle state changes (e.g., a dresser with doors open vs. closed) to the objects they viewed at an astonishing 87% accuracy. Thus, object states can be robustly stored in long-term visual memory. Moreover, participants were as accurate at distinguishing changes in state of objects as they were in detecting exemplar differences (e.g., two different remote controls). Thus, in the two types of pictures employed in the current study (exemplars of categories and different snapshots of single objects), Brady et al (2008) find equivalent and very high accuracy in visual long-term memory.

Despite evidence that different snapshots of objects are robustly encoded in visual long-term memory, we did not find evidence that participants used this information to learn in a statistical learning task. A crucial difference between the tasks has to do with the density and repetition of visual information for a single object. In the current study, four different snapshots of a single object were repeated 28 times each whereas Brady et al (2008) presented "objects from mostly distinct basic-level categories [in order to] minimize conceptual interference" (p. 14326). Thus, Brady et al (2008) employed discrete visual inputs to maximize detailed visual memory whereas a statistical learning task like the current one, by definition, requires integration of across multiple repetitions and variable inputs across time and/or space. Perhaps these different task demands tap into different types of learning and memory or even visual processing, which modulate the maintenance of object-state information vs. the abstraction over object-state information in favor of learning at the level of the objects themselves. Another notable difference between the current study and the one from Brady et al (2008) is the type of category (animate vs. inanimate).

Another recent study found evidence that implicit visual spatial learning is *not* learned in a viewpoint invariant fashion. In a very different task and domain, Jiang and Swallow (2014) examined statistical learning of scenes in a visual search task when participants changed their viewpoint relative to a static visual scene. The authors consistently found that participants did not learn when their viewpoint changed but could learn from the same statistical regularities when they either did not change their viewpoint or the visual scene changed along with their position, obviating the need for any visual transformation of the scene. As with the Brady et al (2008) study, there are a number of differences between the current task and Jiang and Swallow's (2014), which might explain these different findings and should be explored in future work. In Jiang and Swallow (2014), the tasks were designed such that learning based on viewpoint invariance would require participants to

perform *addition* visual transformation of the visual input (i.e., encode the input independent of their viewpoint) adding a level of difficulty to the task. Indeed, when the task was made substantially easier, participants did show the ability to learn from statistical regularities in a viewpoint-invariant fashion. In the current experiment, it is intuitive that segregation of the multiple pictures of a single object into discrete units is more effortful than transformation of these pictures into a single object representation. Moreover the learning task is rendered easier if learning based on individual objects (2 pairs) vs. pictures (8 pairs).⁶ Thus, across these two sets of results, task difficulty maps onto level of visual encoding (viewpoint or state specific vs. invariant) in opposite directions. There could also be important differences in the ecological demands of each task; Visual spatial implicit learning and statistical learning of animate objects might place different demands on one's visual and learning systems.

More broadly, viewpoint invariance is a seminal and theoretically important question in visual object processing. Indeed, a number of studies have found that visual processing and visual memory is highly sensitive to viewpoint differences as demonstrated behaviorally (Brady et al., 2008; Tarr, Williams, Hayward & Gauthier, 1998). In the last 15 years, the question of viewpoint encoding has received considerable attention in relation to mapping the nature of the visual system in the human brain using fMRI, and as some of the cortical regions involved in visual object perception (e.g., the lateral occipital cortex/complex or LOC) have been directly implicated in sequential visual statistical learning (Turk-Browne et al., 2009), this literature is particularly relevant to the current findings. In a seminal study, Vuilleumier, Henson, Driver and Dolan (2002), employed repetition suppression to uncover regions that are sensitive to viewpoint changes and those that exhibit invariance to viewpoint changes. They found that a number of regions typically associated with object processing exhibit viewpoint sensitivity (as measured by differences in adaptation to objects with and without viewpoint changes): notably the right LOC, the right fusiform and the posterior parietal cortex. A region in the left fusiform exhibited viewpoint independence (equal adaptation to objects whether it was a repeated or different viewpoints). Thus, many visual regions are sensitive to viewpoint changes - many more than exhibit complete viewpoint independence. Moreover, Vuilleumier et al (2002) found that only the left inferior frontal cortex exhibited repetition suppression for objects of the same basic-categories ("with the same labels" as they stated). However, other studies have confirmed that the LOC and other visual regions do not represent objects in a completely viewpoint dependent fashion (e.g., Konen & Kastner, 2008): Multivoxel pattern analysis (MVPA) has revealed that the LOC contains object-specific information which is invariant across viewpoints (more specifically in the posterior region of the LOC) and also that these same areas of the LOC encode some amount of category-level information (Eger, Ashburner, Haynes, Dolan & Rees, 2008). Interestingly, MVPA also revealed viewpoint invariant information in earlier visual regions (i.e., ones not typically believed to encode invariant object information). Yet other studies have found that representations of object shape and viewpoint are distributed broadly throughout the visual system including in multiple pathways (Konen & Kastner, 2008). In sum, these findings present the picture that visual object regions such as LOC, which has

⁶Though, we do not find that participants exhibit the same bias towards "ease" across object-level and category-level representations.

been directly implicated in visual statistical learning, are both sensitive and insensitive to viewpoint changes, as well as sensitive and insensitive to category level information.

However, it is not clear how these results correspond to the activity of the visual system in a statistical learning task. Specifically, the kind of repetition employed in fMRI adaptation studies (e.g., Konen & Kastner, 2008; Vuilleumier et al., 2002) is qualitatively different from the kind of repetition employed in a statistical learning experiment even at the most basic level. In fMRI adaptation studies, simple repetition results in suppression, whereas activity in visual regions *increases* during the time-course of statistical learning (e.g., Turk-Browne et al. 2009, see Summerfield, Trittschuh, Monti, Mesulam & Egner, 2008 for work on the relationship between predictability, repetition and neural activity, and Karuza, Emberson & Aslin, 2014 for a discussion of considering neuroimaging data in relation to the time-course of learning). It is also possible that during a statistical learning task, where participants are asked to integrate information across many dozens of presentations, the type of information coded in regions like the LOC shifts with the level of predictions being generated by other learning and memory systems such as the medial temporal lobe (Turk-Browne et al., 2010). This is an important area for future work.

General Discussion

Our daily sensory input – the information that any statistical learning mechanism(s) must use to learn about the structure of the surrounding environment – delivers much *richer* statistical information than the tasks that are typically employed in the laboratory. We investigated the nature of statistical learning when the statistical information being encountered is richer and more complex than in previous studies, and specifically, when there is learnable statistical information present at multiple levels of abstraction: from specific snapshots of an object to the semantic category that the object belongs to; from the least to the most abstract relative to the sensory input.

We report two major findings: 1) There are marked differences in the evidence for abstract level learning (the level of semantic or basic-level categories) when category-level information is presented in the absence vs. the presence of object-specific regularities. Specifically, participants clearly learn from statistics at the most abstract level (the level of semantic categories) when more specific statistical information is *not* present. This is evidenced in part by a robust typicality effect in Experiment 1. However, when object-specific information is introduced, evidence of learning at the category-level disappears (i.e., the typicality effect) or exhibits effects opposite from the hypothesized direction (i.e., poorer learning when foils violate both category-level and object-level regularities) suggesting that participants are not using both types of information at test. In combination with evidence for robust object-level learning, these results suggest that participants exhibit stronger, if not exclusive, learning of object-level regularities. 2) Learners have a bias towards learning from statistical information across individual objects, which is not a bias towards learning the most specific pattern in the input. Indeed, we find that when participants could learn either that individual objects and/or that individual snapshots of these objects predict each other, participants do not show any sensitivity to the individual snapshots but instead only to the statistical information at the level of the individual objects.

Is there evidence of learning at both the category-level and object-level when both types of statistical information are present? We investigated simultaneous category-level and object-level learning in two distinct ways: First, we investigated the presence of the typicality effect robustly found in Experiment 1 when only category-level regularities were present. Typicality effects are considered “the strongest and most reliable effects in the categorization literature” (Murphy, 2002, p.22). In a well-powered test, we find no suggestion of a typicality effect when object-level regularities were introduced across two experiments (Experiments 2 and 3). Thus, a well-recognized signature of category-level processing that would be present during familiarization is present when category-level information is presented alone, and is absent in two separate groups when object-specific information is presented. Second, we hypothesized that if participants were learning at both levels of abstraction, they would benefit from foils which violate both category-level and object-level regularities (i.e., comparing Experiment 2 and 3), but instead we find a significant decrease in learning across these experiments. Broadly, evidence of a decrease in performance when going from one cue to two cues suggests that both types of cues are *not* being learned during familiarization. Despite this robust evidence, as discussed earlier, there are additional ways that the presence of category-level learning could also be investigated, and the current results do not conclusively show that there is no subtle category-level learning present in the current experiment. However, at the very least, the current results demonstrate that category-level learning is negatively affected by the presence of object-specific regularities. Overall the current study establishes that statistical learning does not proceed uniformly across all levels of abstraction and does exhibit biases towards object-level regularities in the presence of both more abstract, and less abstract statistical information.

These results directly bear on the question of the nature of the mechanisms supporting statistical learning as well as provide an answer to what is learned when encountering rich statistical input similar to our daily experiences. Contrary to the view that statistical learning is largely unconstrained, our results show some level of constraint on statistical learning whereby less abstract statistical relationships appear to be preferentially learned. However, we do not show that participants learn based on the *least* abstract patterns investigated, as participants clearly do not exhibit evidence of learning based on individual snapshots of objects. Thus, across four experiments, we find evidence that participants are biased toward object-level regularities and that participants do not learn equally well across multiple levels of abstraction presented as would be predicted if statistical learning were entirely unconstrained. Thus, our results clearly support the view that there are constraints on the mechanism(s) supporting statistical learning. It is important to note that it is *possible* to obtain evidence of statistical learning at both the semantic and object-levels in separate experiments (i.e., during restricted exposure designed to provide existence proofs of learning from a particular type of regularities) but it is only when we probe learning in a rich context where both of these levels of abstraction are learnable simultaneously that individuals’ learning outcomes dramatically shift towards object-specific regularities. This demonstrates emerging biases or constraints in statistical learning when more naturalistic input is provided to the learner.

Importantly, evidence of constraints in statistical learning does not directly challenge the assumption of domain-generalty in this learning mechanism. For example, it is possible for

constraints to be general across domains. Indeed, recent work has directly argued that even modality-specificity (a topic not tackled in the current paper, see Introduction for review and brief discussion) is compatible with a domain-general view of statistical learning (Frost et al., in press). Thus, future theoretical work will benefit from a clear discussion on how domain-general and specificity/constraints of statistical learning interrelate.

However, our results do not unambiguously support the view that statistical learning mechanisms assemble the most complete mental model of the environment and enable the best possible prediction of the upcoming sensory input: we fail to find evidence of learning between individual snapshots of objects but rather only the relationships between the objects themselves. In this case, learning based on the objects did not entail reducing prediction error of upcoming sensory input as much as was possible in the current experiments, as would be predicted by a predictive coding accounts. One possible explanation is that individual snapshots of objects rarely, if ever, are predictive in our everyday lives, and thus we have developed a strong prior against learning from this kind of statistical information. This would effectively bias statistical learning mechanism(s) against learning at this level of regularity even if they were aimed at reducing prediction error as much as possible.

Current results also do not clearly fall in line with the “less is more” (Newport, 1990) or “starting small” (Elman, 1993) theories of language learning (and by extension statistical learning). While it is not clear exactly what a “starting small” prediction would be for this particular task, it seems unlikely that it would predict that participants would learn only based on object-level regularities when they entail both learning more pairs and at a lower level of abstraction in some circumstances (8 pairs of individual objects vs. 2 pairs of categories; Experiments 2 and 3) and learning fewer pairs and at a higher level of abstraction in others (2 pairs of objects; 8 pairs of views on these objects; Experiment 4). One possibility could be that objects are both the most natural and the smallest unit in these stimuli and thus learning based on them as opposed to more abstract (categories) or less abstract (snapshots) information could be considered the simplest learning to do in the current task. This is a similar explanation to what is presented above from the predictive coding perspective.

In both the case of “less is more” and the predictive coding theories of statistical learning, the current results could fit with these constraint-based views of statistical learning in the following way: Perhaps learning is biased towards the smallest unit, most simple pattern, or towards producing the lowest prediction error, but it is also affected by biases which suppress learning based on certain types of statistical information (e.g., learning of individual snapshots of objects). A bias of this sort could arise from previous learning experiences. This type of explanation certainly seems feasible, however, one of the strengths of these proposals is their, relatively, clear predictions about what will be learned under different conditions; invoking *a posteriori* explanations for when learning cannot proceed, without a way to quantify or necessarily predict these cases, reduces the explanatory power of these views of learning.

Importantly, the idea that prior learning experience might bias learning towards individual objects does generate specific hypotheses concerning the development of statistical learning.

Specifically, the “less is more” and “starting small” views do not presume that all learners are the same and specifically propose that younger learners might be more likely to start small. While less expressly developmental, a more Bayesian view of learning allows different strengths of priors to affect any kind of change of a model in response to experience. Thus, if adults’ priors are weighted against considering snapshots of objects as important statistical information to accommodate in their model, then it would be crucial to investigate this question in a developmental context. Specifically, it might be that infants or young children are able to learn based on the individual snapshots of objects in the current type of exposure while adults are not.

Behavioral and neural responses to statistical information have quickly become a foundational concept in many domains of cognitive neuroscience and cognitive psychology. As reviewed in the introduction, domains from language comprehension in adults to visual development are believed to be heavily influenced by exposure to statistical information. Thus, the question of what is learned when an individual encounters rich statistical information is crucial to address for a number of diverse lines of research. The interpretation of the current results and the corresponding verification of similar constraints on statistical learning in other domains is important future research. For example, if different levels of abstraction are systematically investigated in language comprehension, are individuals biased to respond to information at the lexical level? Given the importance of phonological information to language comprehension, perhaps phonological information is also learnable and even trumps lexical learning? Is this the same when participants are learning from tokens from a familiar vs. an unfamiliar language?

While additional investigations must be carried out in other domains, the implications of the current work are already clear: statistical learning cannot be assumed to be unconstrained in even if statistical learning has been established at multiple levels of abstraction. Given that statistical learning has been established across diverse inputs in many sensory modalities (see Introduction for a review), it has been tempting to assume that statistical learning is largely unconstrained and that demonstrations of learning in restricted circumstances will scale up to learning in more complex situations. The current results demonstrate that this is not the case for familiar visual objects and should caution the assumption of unconstrained statistical learning to researchers in other domains.

Finally, another important area of future research is when and how participants overcome the bias to learn at the object-level and learn at the level of categories. Current work suggests that it is a lack of learnability at the object-level that results in this learning shift: In Experiment 1, category-level learning was induced by scrambling the specific exemplars, which makes object-level regularities highly unreliable. Category-level learning might also be achieved by having a large number of object-level pairs (i.e., increasing the number of exemplars so that the number of object-level pairs increases). It is also unclear whether category-level and object-level statistical learning would be supported by similar or different neural mechanisms: Vuilleumier et al. (2002) found that category-level repetition suppression only occurred in the left prefrontal cortex while repetition suppression at the object-level (and viewpoint level) occurred throughout the ventral and dorsal visual streams. However, other studies have suggested that category level information is present in the LOC

(e.g., Eger et al., 2008). Both of these cortical regions have been implicated in statistical learning (Karuza et al., 2013, Turk-Browne et al., 2009) however it is unclear whether both would be equally involved in the learning of category-level or object-level regularities. An investigation of the neural mechanisms supporting category-level vs. object-level regularities would also be interesting in relation to theoretical questions of whether statistical learning is supported by single amodal learning systems or whether domain-general learning abilities are achieved through diverse mechanisms (e.g., Frost et al., in press).

In sum, we employed a novel task to examine how statistical learning proceeds when information along multiple levels of abstract is learnable. Three levels of abstraction were investigated: the most abstract being semantic categories, then, the individual objects themselves, and the least abstract, the individual snapshots or pictures of the objects. We find that learning is robustly biased towards learning from the statistical regularities present between individual objects whether learning at this level results in the best prediction of upcoming sensory input (Experiments 2 and 3) or when it results in a suboptimal prediction of successive pictures (Experiment 4). This suggests a bias in the mechanisms of statistical learning for a particular level of abstraction and provides contradictory evidence for the majority of theories of the mechanisms supporting statistical learning. This work has immediate implications for many areas of cognitive psychology and cognitive neuroscience where statistical information is believed to influence behavior: statistical learning cannot be assumed to be unconstrained in even if statistical learning has been established at multiple levels of abstraction.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We'd like to thank Drs. Dima Amso, Rick Dale, Jordan DeLong, David Field, Thomas Farmer, Gary Lupyan, Adele Goldberg, Erika Bergelson, Toben Mintz and three anonymous Reviewers for helpful conversations and/or comments on the manuscript. We'd also like to thank Esteban Buz and Dave Kleinschmidt for their (statistical) support. Thank you to Claire Schmidt, Andrew Webb, Joey Ciufo, Mary Marchetti, Haley Weaver and Camila Rivero for help with data collection and, in particular, we thank Dr. Michael Spivey for his support during initial data collection at Cornell University. This work is supported by a Canadian Institution of Health Research post-doctoral fellowship (201210MFE- 290131-231192) and an NIH K99 award (HD076166-01A1) to LLE.

References

- Baldwin D, Andersson A, Saffran J, Myers M. Segmenting dynamic human action via statistical structure. *Cognition*. 2007; 106:1382–1407. [PubMed: 18035346]
- Brady TF, Konkle T, Alvarez GA, Oliva A. Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*. 2008; 105(38):14325–14329.
- Brady TF, Oliva A. Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent. *Psychological Science*. 2008; 19:678–685. [PubMed: 18727783]
- Champely, S. pwr: Basic functions for power analysis. R package version 1.1.1. 2012. <http://CRAN.R-project.org/package=pwr>
- Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*. 2013; 36:181–253. [PubMed: 23663408]

- Conway C, Christiansen MH. Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*. 2005; 31:24–39.
- Dale R, Kehoe CE, Spivey MJ. Graded motor responses in the time course of categorizing atypical exemplars. *Memory and Cognition*. 2007; 35:15–28. [PubMed: 17533876]
- Eger E, Ashburner J, Haynes J-D, Dolan RJ, Rees G. fMRI activity patterns in human LOC carry information about object exemplars within categories. *Journal of Cognitive Neuroscience*. 2008; 20:356–370. [PubMed: 18275340]
- Emberson LL, Amso A. Learning to sample: Eye tracking and fMRI indices of Changes in Object Perception. *Journal of Cognitive Neuroscience*. 2012; 24:2030–2042. [PubMed: 22721373]
- Emberson LL, Conway C, Christiansen MC. Timing is everything: change in presentation rate have opposite effects on auditory and visual implicit statistical learning. *Quarterly Journal of Experimental Psychology*. 2011; 64:1021–1040.
- Emberson LL, Liu R, Zevin J. Is statistical learning constrained by lower level perceptual organization? *Cognition*. 2013; 128:82–102. [PubMed: 23618755]
- Fine AB, Jaeger TF. Evidence for error-based implicit learning in adult language processing. *Cognitive Science*. 2013:578–591. [PubMed: 23363004]
- Fiser J, Aslin RA. Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*. 2001; 12:499–504. [PubMed: 11760138]
- Friston K. A theory of cortical responses. *Philosophical transactions of the Royal Society of London Series B. Biological Sciences*. 2005; 360:815–36. [PubMed: 15937014]
- Frost R, Armstrong BC, Siegelman N, Christiansen MH. Domain generality versus modality specificity: the paradox of statistical learning. *Trends in Cognitive Sciences*. in press.
- Gebhart A, Newport E, Aslin R. Statistical learning of adjacent and nonadjacent dependencies among nonlinguistic sounds. *Psychonomic Bulletin & Review*. 2009; 16:486–490. [PubMed: 19451373]
- Hunt RH, Aslin RN. Statistical Learning in a Serial Reaction Time Task: Access to Separable Statistical Cues by Individual Learners. *Journal of Experimental Psychology*. 2001; 130:658–680. [PubMed: 11757874]
- Jiang YV, Swallow KM. Changing viewer perspectives reveals constraints to implicit statistical learning. *Journal of Vision*. 2014; 14(12):1–16.
- Jolicoeur P, Gluck MA, Kosslyn SM. Pictures and names: Making the connection. *Cognitive Psychology*. 1984; 16:243–275. [PubMed: 6734136]
- Karuz EA, Emberson LL, Aslin RN. Invited Review: Combining fMRI and behavioral measures to examine the process of human learning. *Neurobiology of Learning and Memory*. 2014; 109:193–206. [PubMed: 24076012]
- Karuz EA, Farmer, TA.; Fine, AB.; Smith, FX.; Jaeger, TF. On-line measures of prediction in a self-paced statistical learning task. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*; 2014.
- Karuz EA, Newport EL, Aslin RN, Starling SJ, Tivarus ME, Bavelier D. The neural correlated of statistical learning in a word segmentation task: An fMRI study. *Brain and Language: Special Issue*. 2013; 127:46–54.
- Keele SW, Mayr U, Ivry R, Hazeltine E. The Cognitive and Neural Architecture of Sequence Representation. *Psychological Review*. 2003; 110:316–339. [PubMed: 12747526]
- Kim G, Lewis-Peacock JA, Norman KA, Turk-Browne NB. Pruning memories by context-based prediction error. *Proceedings of the National Academy of Sciences*. 2014; 111(24):8997–9002.
- Kleinschmidt DF, Jaeger TF. Speech recognition: Recognizing the familiar, generalizing to the similar, and adapting to the novel. (accepted).
- Koch I, Hoffmann J. Patterns, chunks, and hierarchies in serial-reaction-time tasks. *Psychological Research*. 2000; 63:22–35. [PubMed: 10743384]
- Konen CS, Kastner S. Two hierarchically organized neural systems for object information in human visual cortex. *Nature Neuroscience*. 2008; 11:224–231. [PubMed: 18193041]
- Lupyan G. Cognitive penetrability of perception in the age of prediction: Predictive systems are penetrable systems. *Review of Philosophy and Psychology*. forthcoming

- Marcus GF, Vijayan S, Rao SB, Vishton PM. Rule learning by seven-month old infants. *Science*. 1999; 283:77–80. [PubMed: 9872745]
- Mayr U. Spatial Attention and Implicit Sequence Learning: Evidence for Independent Learning of Spatial and Nonspatial Sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1996; 22:350–364.
- Murphy, GL. *The Big Book of Concepts*. Cambridge, MA: The MIT Press; 2002.
- Onnis L, Waterfall H, Edelman S. Learn locally, act globally: learning language from variation set cues. *Cognition*. 2008; 109:423–430. [PubMed: 19019350]
- Perruchet P, Pacton S. Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences*. 2006; 10:233–238. [PubMed: 16616590]
- Purves D, Wojtach W, Lotto RB. Understanding vision in wholly empirical terms. *Proceedings of the National Academy of Science*. 2011; 108:15588–15595.
- Purves D, Monson BB, Sundararajan J, Wojtach W. How biological vision succeeds in the physical world. *Proceedings of the National Academy of Science*. 2014; 111:4750–4755.
- Reeder PA, Newport EL, Aslin RN. From shared contexts to syntactic categories: The role of distributional information in learning linguistic form-classes. *Cognitive Psychology*. 2013; 66:30–54. [PubMed: 23089290]
- Rescorla RA. Behavioral Studies of Pavlovian Conditioning. *Annual Reviews in Neuroscience*. 1988; 11:329–352.
- Romberg AR, Saffran JR. Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science*. 2010; 1:906–914. [PubMed: 21666883]
- Rosch E, Mervis C, Gray W, Johnson D, Boyes-Braem P. Basic objects in natural categories. *Cognitive Psychology*. 1976; 8:382–439.
- Saffran JR. Words in a sea of sounds: the output of infant statistical learning. *Cognition*. 2001; 81:149–169. [PubMed: 11376640]
- Saffran JR. Constraints on statistical language learning. *Journal of Memory and Language*. 2002; 47:172–196.
- Schultz W, Dayan P, Montague PR. A neural substrate for prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Smith JG, McDowall J. The implicit sequence learning deficit in patients with Parkinson's disease: A matter of impaired sequence integration? *Neuropsychologia*. 2006; 44:275–288. [PubMed: 15964035]
- Summerfield C, Trittschuh EH, Monti JM, Mesulam M-M, Egner T. Neural repetition suppression reflected fulfilled perceptual expectations. *Nature Neuroscience*. 2008; 11:1004–1006. [PubMed: 19160497]
- Tarr MJ, Williams P, Hayward WG, Gauthier I. Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience*. 1998; 1:275–277. [PubMed: 10195159]
- Thiessen ED, Erickson LC. Beyond Word Segmentation A Two-Process Account of Statistical Learning. *Current Directions in Psychological Science*. 2013; 22:239–243.
- Turk-Browne NB, Scholl B, Chun MM, Johnson M. Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*. 2009; 21:1934–1945. [PubMed: 18823241]
- Turk-Browne NB, Scholl B, Johnson M, Chun MM. Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience*. 2010; 30:11177–11187. [PubMed: 20720125]
- Vuilleumier P, Henson RN, Driver J, Dolan RJ. Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*. 2002; 5:491–499. [PubMed: 11967545]

Appendix 1: Reaction time analyses for Experiments 1 through 3

While it is important to note that there are reasons to approach this analysis with caution, reaction times (RT) provide additional evidence that category-level learning occurs only in Experiment 1, when object-level statistical information is not present and not during

Experiments 2 and 3. The prediction is that if object-typicality affects learning outcomes, that one would expect slower RTs in the atypical condition compared to the typical condition.

A posteriori, we analyzed RTs during the learning test (i.e., after familiarization). RT is calculated as the amount of time from the presentation of the prompt screen asking participants to select which pair is most familiar to them until their response. These prompts were not designed for RT studies. The timing of this prompt after the final stimulus is entirely predictable. The ability to predict when this screen would appear and what question would be asked allows participants to have potentially very low RTs on certain trials. These responses were also non-speeded: participants could take as long as they like and were not encouraged to be as quick as possible like most RT experiments. Also, we are only able to compare RTs across groups of participants: participants were only exposed to typical *or* atypical exemplars in a between-subjects design. Given large differences in RTs across participants, it is problematic to compare RTs between subjects both because of the large amount of variability it introduces which can potentially obscure real results, and consequently, given current methods, it will remain a question whether any significant differences could be a result of a difference in the mean RTs of a given set of participants and not a result of task manipulations.

With these strong caveats in mind, we used t-tests to compare RTs in Experiment 1 between participants who viewed typical and atypical exemplars, and found that RTs showed similar results to the learning outcomes with regard to typicality: participants who viewed typical exemplars responded significantly faster (mean = 364ms) than those who viewed atypical exemplars (mean = 771ms) according to an unpaired t-test ($t(14) = 2.82, p = 0.014$). This finding further supports the view that there is category level learning in Experiment 1, as an effect of typicality (e.g., slower processing for atypical vs. typical category exemplars) is a specific hallmark of category-level, abstract processing. However, since this finding such a result of a between-subjects comparison, it cannot be ruled out that this is a result of differences in mean RTs across groups of participants.

Crucially, this RT difference is not present in Experiments 2 and 3. In Experiment 2, response time was not affected by typicality (Typical mean = 730ms, std = 221ms; Atypical mean = 640ms, std = 238ms; $t(20) = 0.926, p = 0.37$), nor in Experiment 3, (Typical mean = 733ms, std = 211ms; Atypical mean = 925ms, std = 630ms; $t(22) = -1.00, p = 0.33$). A two-way ANOVA (Experiment 2 vs. 3, Typicality) was performed on response times, revealing no main effect of Experiment ($F(1,42) = 1.70, p = 0.20$), nor typicality of exemplars ($F(1,42) = 0.261, p = 0.61$). As with the accuracy results, the response time results suggest that there was no acquisition of category-level knowledge when individual exemplars or pictures are predictive of picture order.

To confirm that the difference in RTs in Experiment 1 is not due to the differences in learning outcomes across typical and atypical conditions (e.g., more incorrect responses), we examined reaction times in Experiment 4 where participants showed significant differentiation of foils and pairs when the foils violated the object-level regularities but not when the foils violated only the viewpoints (a learned and a non-learned condition). There is

no difference in object typicality across these groups and thus, there is no prediction for differences in RTs. However, if the differences in RTs in Experiment 1 are a result of different learning outcomes, one should also expect differences in RTs in the current experiment. A t-test performed on the reaction times from Experiment 4 revealed no difference between participants who were tested on foils that violated only the viewpoints (mean = 625ms, std = 181ms) and those who were tested on foils that violated both the viewpoints and the objects (mean = 535ms, std = 178ms; $t(14) = -1.00$; $p = 0.33$) providing additional support that category-level learning is influencing RTs in Experiment 1 but that in Experiment 2 and 3 there is no additional category-level learning taking place.

Appendix 2: Explicit Knowledge

While the current experiment was not designed to test the contribution of implicit vs. explicit knowledge in the current task, we did ask a question in the post-test that can provide a glimpse into whether participants developed explicit knowledge as a result of familiarization. Specifically, this question is “Did you notice any patterns in the initial stream of pictures?”

This question is insufficient to provide a full and unbiased picture into the role of explicit knowledge in the current study for two main reasons: First, answers to this question were difficult to ascertain as participants sometimes simply wrote “yes,” but as the question is quite underdetermined it is difficult to know whether this answer indicated the type of explicit awareness that would be relevant for the current task. For example, one participant wrote “yes but then it disappeared” indicating that they did not have explicit knowledge of the patterns in the familiarization stream. However, to be conservative, any participant who simply wrote “yes” was considered to have expressed explicit knowledge of the patterns in the current experiment. Thus, the number of participants with explicit knowledge is likely over-estimated in the current experiments.

Second, it is possible that upon being asked this question, participants may have generalized from their object-specific knowledge to categories. Specifically, asking participants to reflect on their experience and to verbally express their knowledge would likely result in a bias towards reporting category-level information as it is easier to express. Moreover, it is conceivable that recalling any knowledge could result in integration across similar types of pairs (e.g., $\text{dog}_1\text{-fish}_1$, $\text{dog}_2\text{-fish}_2$) to result in category-level information that was not present during familiarization and did not drive performance during the statistical learning test, but could be expressed as such in response to this question.

Despite the notable limitations of this data, the patterns of explicit knowledge that we observed are interesting and worthy of follow-ups that are more specifically designed to tackle this issue (e.g., manipulating the use of explicit strategies before familiarization, better designed probes to reveal the presence or absence of explicit knowledge after familiarization, and best test such as confidence judgment during test trials or more precisely formulated questions).

Experiment 1

We found that 9 out of 16 participants answered yes to the question concerning knowledge of patterns. Some were able to articulate some kind of explicit knowledge about the sequence (e.g., “fish, dog, bird, flower” or “each fish was followed by a bird and each flower was followed by a dog”). All participants who received typical exemplars were able to articulate the pattern to at least some degree, compared to only one participant who received atypical exemplars [“dog or other animal (not sea creature) was always followed by sea creature or flower”]. This pattern clearly indicates that there is some relationship between object typicality and explicit knowledge in the current task. Correspondingly, participants who demonstrated explicit knowledge also demonstrated better performance at test ($t(13.8) = -2.46, p = 0.029$) but again, this comparison is essentially the comparison between typical and atypical exemplar groups presented above. Interestingly, the single participant from the atypical group who demonstrated explicit knowledge did not have particularly good test performance (accuracy = 61%). This was neither the highest score for the atypical group (highest score = 69%) nor above average for the entire group (including typical and typical exemplars).

Experiment 2

Fifteen participants reported evidence of explicit knowledge via the post-test questionnaires. Unlike in Exp. 1, participants with explicit knowledge were more evenly distributed between those exposed to typical vs. atypical exemplars (Typical: 8; Atypical: 7). 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.”). Despite their explicit knowledge being unable to help them at test, these participants still demonstrated a high average accuracy that was significantly above chance level performance (83.4%, $SD = 22.4, t(14) = 14.34, p < 0.001$) and significantly better than those without explicit knowledge ($t(19.5) = -3.4395, p = 0.003$). Importantly, participants who do not express any explicit knowledge still show learning above chance performance ($t(6) = 12.89, p < 0.001$). This pattern of results suggests a dissociation between the explicit knowledge and performance at test, because the knowledge that “dogs follow fish” will not help you to differentiate familiar pairs from foils which do not violate the category-level information. Thus, while the ability to express explicit knowledge at test increases test performance, it is unlikely that explicit knowledge directly leads to good performance in our statistical learning test.

Experiment 3

Nine participants demonstrated explicit knowledge of the structure of the familiarization stream with participants distributed across typical and atypical exposures (5 typical, 4 atypical). Participants with explicit knowledge performed highly at test (92.2%, $SD = 10.9\%, t(8) = 25.28, p < 0.001$). Since explicit knowledge could support performance at test in the current experiment but not in the previous one, we confirmed that performance in those with explicit knowledge of the task did not differ across experiments ($t(16.19) = -1.55,$

$p = 0.14$). Thus, even though a subset of participants could articulate explicit knowledge at the category-level regularities in both experiments, there is no evidence that this explicit knowledge preferentially led to better test performance in Experiment 3, when category-level knowledge could support better test performance.

Experiment 4

Nine participants demonstrated some explicit knowledge of the statistical regularities. Interestingly, roughly half of the participants continued to express this knowledge in terms of categories (e.g., “flowers and puppies usually went together & birds & fish usually went together”), while the others expressed these regularities in terms of objects (e.g., “dog always came after rose & bird after fish. Sometimes fish and rose were back to back”). In this Experiment, all participants received the same exposure but received one of two types of foils. Participants with explicit knowledge were distributed across both types of foils (object-level foils: 4: snapshot-level foils: 5). Unlike the previous experiments, we did not find high test performance in those with explicit knowledge (67.7%, $SD = 23.8\%$, $t(8) = 2.23$, $p = 0.056$). While it is very difficult to draw conclusions from the small numbers of participants, we examined performance based on explicit knowledge based on foil type. We find a significant difference in learning across foil types ($t(4.7) = 2.63$, $p = 0.05$) with greater learning when the participants are comparing familiar pairs to foils that violate object-level regularities (86.0%, $SD = 22.0\%$, $t(3) = 3.27$, $p = 0.047$) than to foils that only violate snapshot-level regularities (53.1%, $SD = 13.1\%$, $t(4) = 0.533$, $p = 0.623$).

Discussion

Thus, we find that participants can develop explicit knowledge in this task. However, these posthoc analyses should be interpreted with caution for a number of reasons: First, this task was not designed to measure the effects of explicit knowledge and we were unable to ascertain the veracity of explicit knowledge for each participant. Specifically, because of the nature of the debriefing questions, we have likely overestimated the number of participants with explicit knowledge or possibly induced or biased this knowledge during the post-test.

Second, the findings concerning explicit knowledge appear to be highly idiosyncratic: In Exp. 1, participants with explicit knowledge by and large were exposed to typical exemplars but in Exp.s 2 and 3, participants with explicit knowledge evenly came from typical and atypical exposure groups; Participants with explicit knowledge demonstrated very high levels of accuracy in the first three experiments but not in the final experiment. This last finding is particularly surprising because this is the only experiment where we find evidence that explicit knowledge seems to relate to the structure of the foils at test.

Finally, and perhaps most revealing, we find that participants with explicit knowledge of semantic categories have excellent performance in a test where reliance on category-level information would have resulted in chance level performance and indeed, comparing between experiments where category level knowledge would either help at test or not (Exp.s 2 and 3), we find no difference in performance for those with explicit knowledge.

Together, these findings suggest that explicit knowledge is at best indirectly related to test performance. One possibility that is broadly consistent with these findings is that participants who achieve a high level of proficiency arising from implicit statistical learning mechanisms develop explicit knowledge or are able to integrate across their implicitly learned representations to state an extrapolation or generalization of this knowledge when asked in the post-test questionnaire. This may be more pronounced in the current task as the experimental design necessitated a relatively few number of pictures and also employed familiar and easily nameable semantic categories. Future work would be needed to tease apart these various factors to more directly consider the role of implicit and explicit learning mechanisms in the current task.

Appendix 2

Raw data associated with this article can be found at the Open Science Framework: <https://osf.io/uqe5x/>.

Highlights

- Examined learning of statistical regularities at multiple levels of abstraction
- Learning at the most abstract level occurs only in the absence of other information
- Learning preferentially occurs based on objects not at other levels of abstraction

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

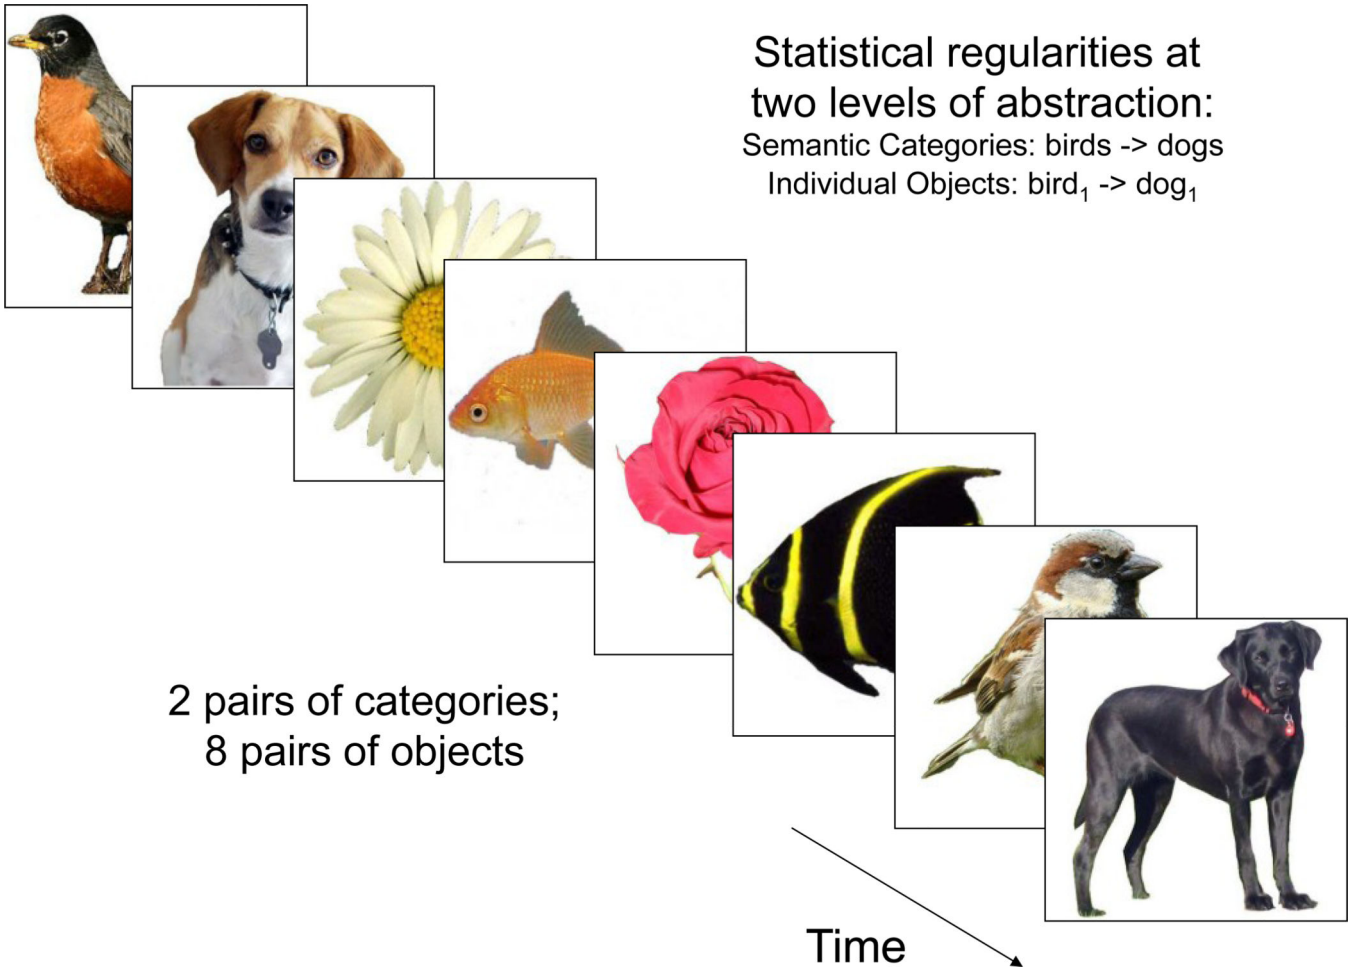


Fig. 1. A sample familiarization stream from Experiments 2 and 3. The same familiarization was given to participants in both experiments. Pictures were organized into pairs of categories (e.g. birds -> dogs) as well as specific objects within these categories (e.g. robin -> beagle). Thus, predictive regularities were redundant across multiple levels of abstraction resulting in two pairs of categories and eight pairs of objects or exemplars of these categories. In the sample stream, birds predict dogs and flowers predict fish.

All Atypical Exemplars




Fig. 2.

All atypical exemplars used in the current paper, organized by category (from left: dog, flower, fish, bird). Half of participants received familiarization with atypical exemplars, pictured here, and half were familiarized with typical exemplars depicted in Figs. 1 and 3. For a given subject, the test phase was conducted with the same pictures as familiarization (i.e., participants familiarized with atypical pictures were presented with the same atypical pictures during test). We hypothesized that any category-level learning would be modulated by typicality of exemplars, with atypical exemplars resulting in weaker category-level learning. A number of the typical exemplars are presented in Figures 1 and 3. The set of stimuli are available as supplementary material and at the first author's website.

Pair from Familiarization

Foil at Test

Experiment 2
Testing for object-
level learning



Experiment 3
Testing for additional
categorical-level learning




Fig. 3. Method differences for Experiments 2 and 3. A sample pair from familiarization is compared to a sample foil from each Experiment. The familiarization is consistent across experiments, represented by the same familiarization pair. The sole difference between experiments was the composition of foils used at test. In Experiment 2, foils were designed to assess learning at the object or exemplar-specific level only. In this case, the category-level relationship of birds predicting dogs is held constant but the specific dog is changed to violate the object-specific regularities. In Experiment 3, foils allow for knowledge at both the object and category regularities to influence test performance.

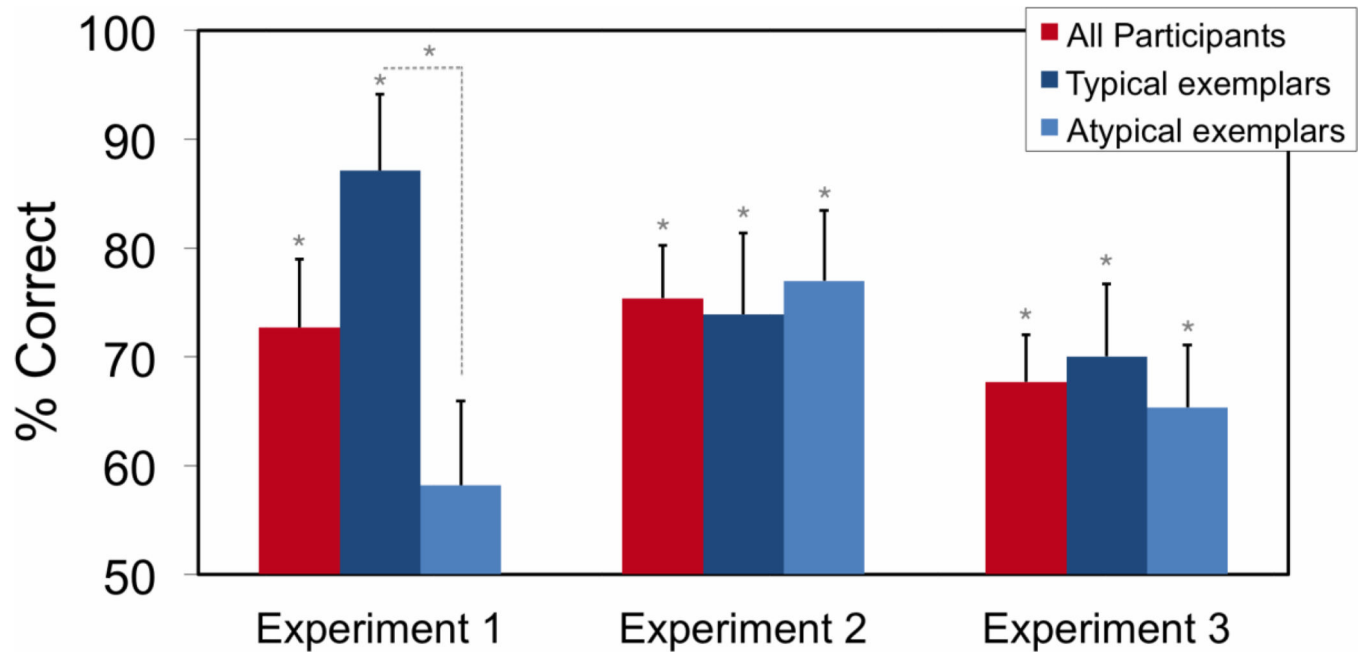


Fig 4.

Results for Experiments 1 through 3: red bars present performance for all participants and the blue bars present means for participants grouped by whether they viewed typical or atypical exemplars. To demonstrate significant learning, we compare performance in each group to chance performance of 50% (the origin of the y-axis) and significance in this test is indicated as an asterisk above the error bar. Differences across groups within an experiment (e.g., the typicality effect in Experiment 1) are indicated by dotted grey lines with an accompanying asterisk. Error bars represent standard error of the mean.

Four Objects: Four Snapshots Each




Fig. 5.

In Exp. 4, new stimuli were employed. Participants saw four different *snapshots* from four objects, derived from the same semantic categories as previous experiments. Each snapshot varied in viewpoint and/or state.

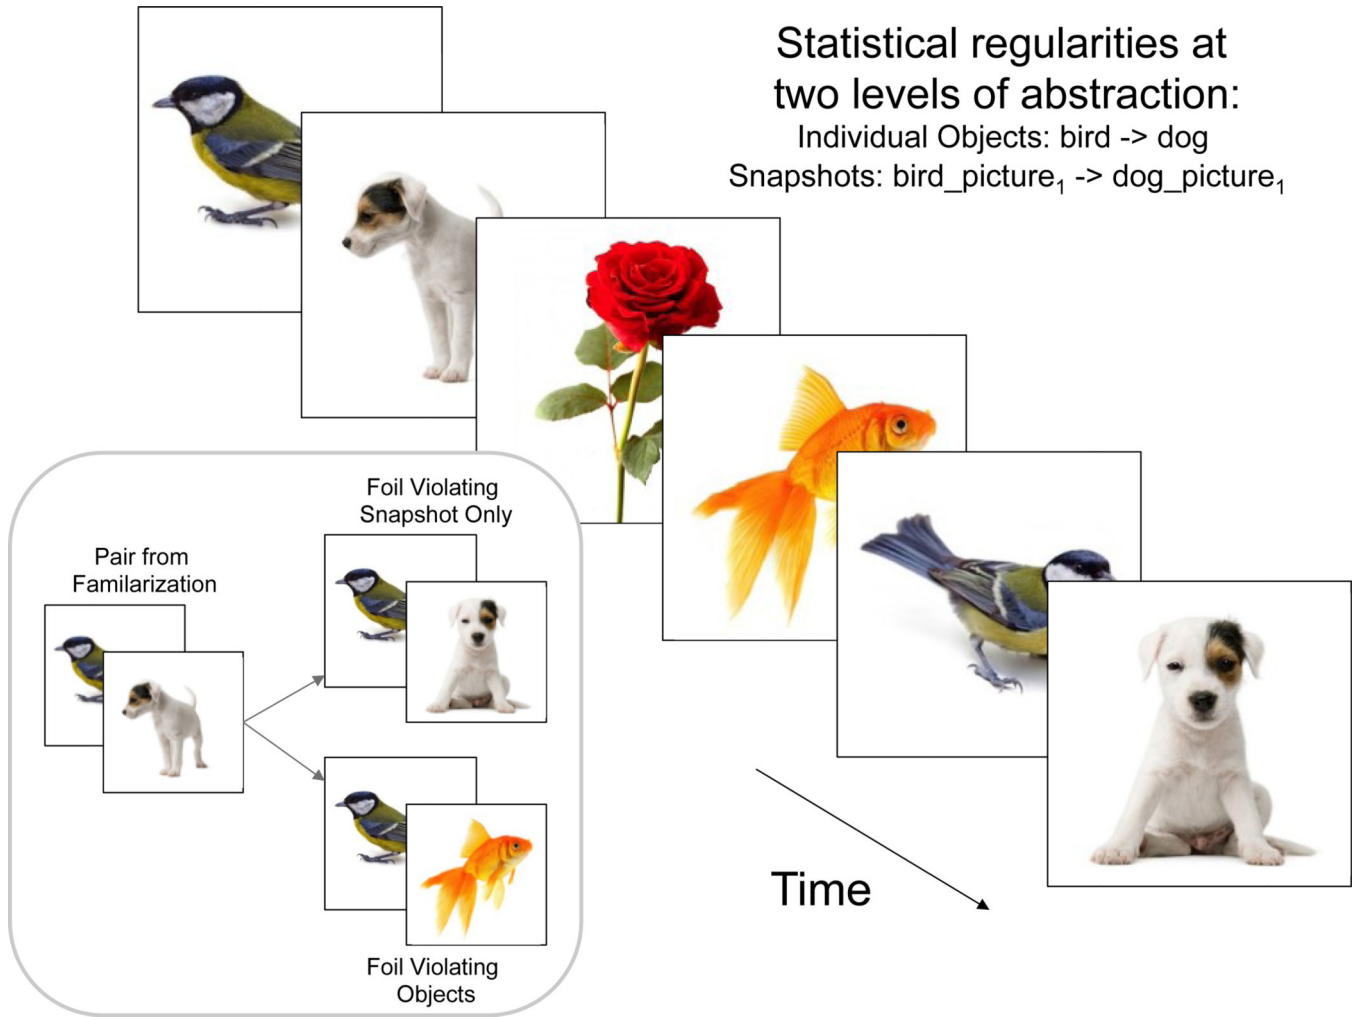


Fig. 6. A sample familiarization stream from Experiment 4. All participants received the same familiarization: pictures were organized into pairs of objects (e.g. bird -> dog) but the snapshots of these objects were also in pairs (e.g., picture 1 of the bird -> picture 2 of the dog). Thus, predictive regularities were redundant across multiple levels of abstraction resulting in two pairs of objects and eight pairs of pictures or snapshots of these objects. Participants either compared pairs from the familiarization to foils which only violated the viewpoint statistical regularities (e.g., picture 1 of the bird -> picture 2 of the dog) or foils which violated the objects (e.g., picture 1 of the bird > picture 1 of the fish).

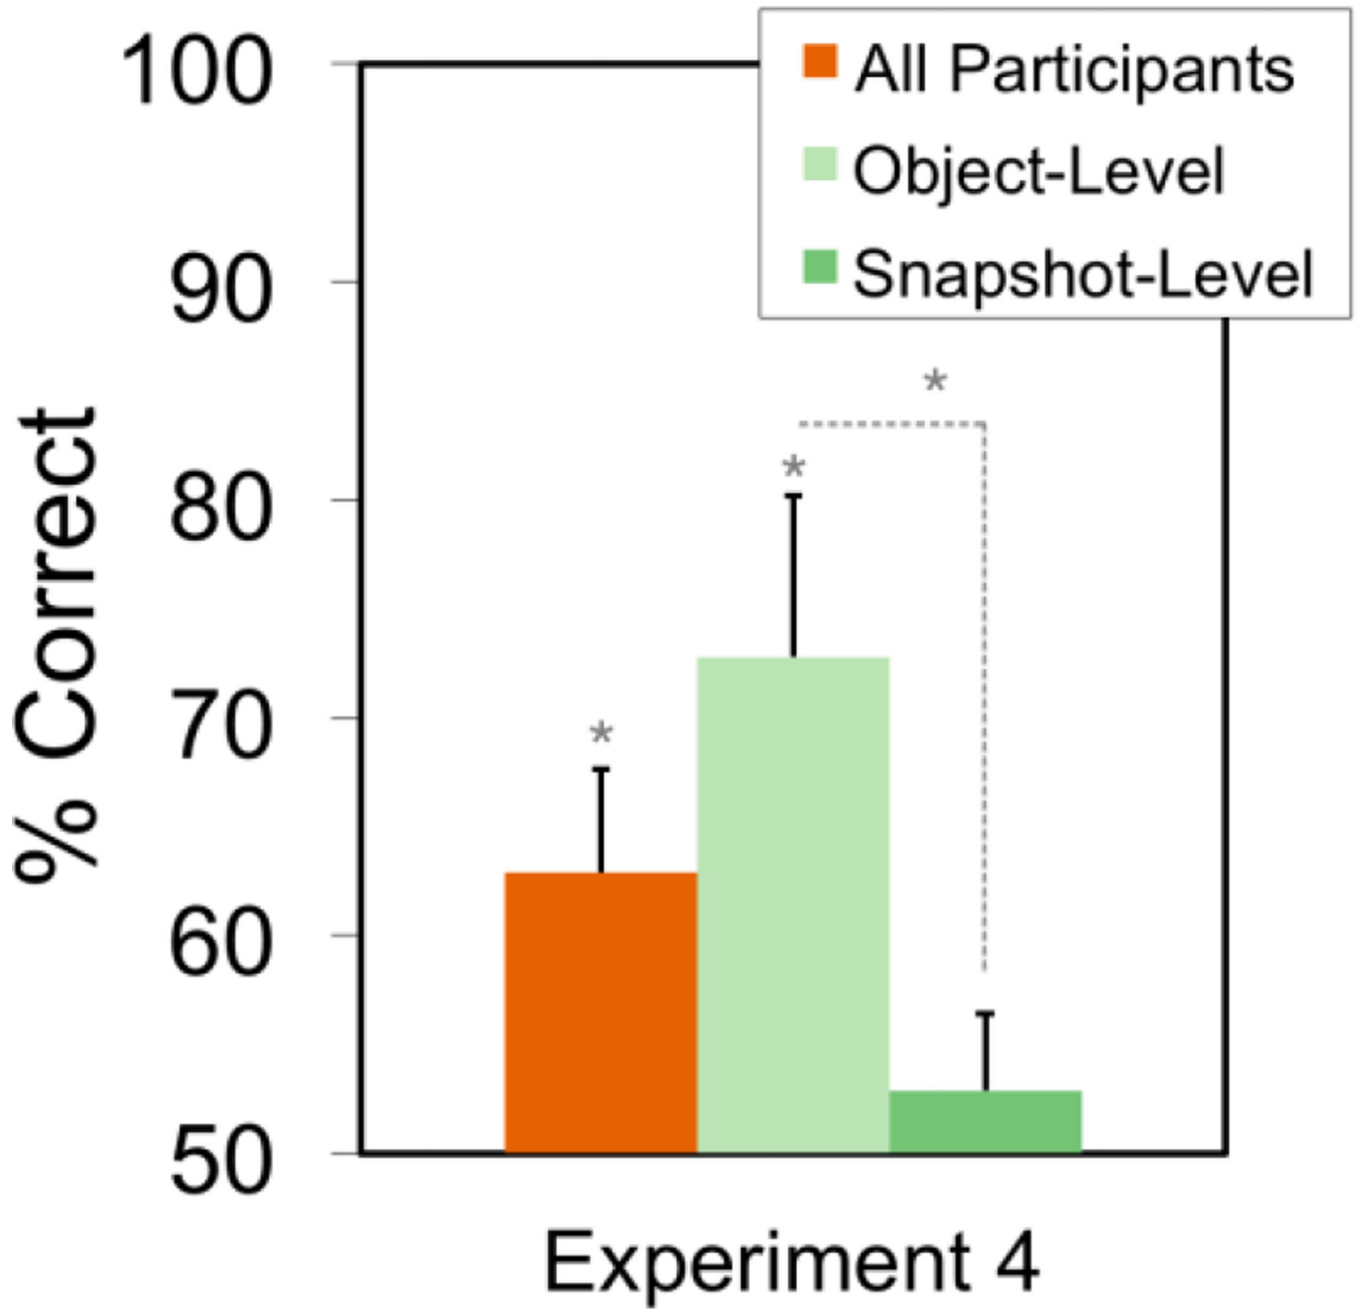


Fig. 7. Results from Experiment 4: Orange bar depicts performance for all participants. Green bars present performance by foil type: Foils tested either for Object-level knowledge (bird_view₁-dog_view₁ vs. bird_view₁-fish_view₂) or for Snapshot-level knowledge (e.g., bird_view₁-dog_view₁ vs. bird_view₁-dog_view₂).

Table 1

Summary of learning outcomes from all experiments in relation to key theoretical predictions.

Experiment	Level of Statistical Regularities	What Level Was Learned	Typicality Effect?	Gain the Best Predictions?
1	Category only	Category	Yes	Yes
2	Category and individual object	Individual object	No	Yes
3	Category and individual object	Individual object	No	Yes
4	Individual object and snapshot	Individual object	n/a	No

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript