CONTEXTUAL INFLUENCES ON VISUAL STATISTICAL LEARNING

by

Leeland Loew Rogers

A dissertation submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Psychological and Brain Sciences.

Spring 2022

CONTEXTUAL INFLUENCES ON VISUAL STATISTICAL LEARNING

by

Leeland Loew Rogers

Approved: ____

Tania Roth, Ph.D.
Chair of the Department of Psychological and Brain Sciences

Approved: ____

John A. Pelesko, Ph.D.
Dean of the College of Arts and Sciences

Approved: ____

Louis F. Rossi, Ph.D.
Vice Provost for Graduate and Professional Education and
Dean of the Graduate College

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ______________________________

Timothy J. Vickery, Ph.D.
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ______________________________

Irene P. Kan, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ______________________________

Paul Quinn, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ______________________________

Keith Schneider, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ______________________________

Peter Mende-Siedlecki, Ph.D.
Member of dissertation committee

ACKNOWLEDGMENTS

I would first like to thank my mentor, Dr. Timothy J. Vickery for his incredible mentorship during my time at the University of Delaware. His guidance has been invaluable to me and my career. Next, I would like to thank Dr. Irene P. Kan, who had supported me as my advisor during my master's thesis and who offers support again by serving on my dissertation committee. I would also like to thank the rest of my dissertation committee, Dr. Paul Quinn, Dr. Keith Schneider, and Dr. Peter Mende-Siedlecki for their support and guidance during my dissertation. I must also recognize three friends and colleagues of mine, Dr. Sarah Fairchild, Dr. Steve Beighley, and Dr. Su Hyoun Park, who have always uplifted me with their incredible support and friendship. I also thank my stepmother and my father, Linda and Bruce Rogers, who have always given me unconditional love and support. This thanks also extends to my late mother, Carolyn Loew, who gave everything she could to provide for me. Lastly, but certainly not least, I would like to thank my loving partner, Derek Hughes, who has been with me every step of the way offering endless support and encouragement.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

Visual statistical learning describes the unintentional extraction of statistical regularities from visual environments across time or space, and is typically studied using novel stimuli (e.g., symbols unfamiliar to participants). Additionally, familiarization procedures in experiments that have explored visual statistical learning tend to be passive or require only basic vigilance from participants. The natural visual world, however, is rich with a variety of complex visual stimuli, and we experience that world in the presence of goal-driven behavior including overt learning of other kinds. The present dissertation examines how visual statistical learning responds to such contextual factors. Chapter 2 finds that visual statistical learning is sometimes influenced by natural and artificial categories in the presence (and absence) of explicit learning about artificial categories. Chapter 3 examines the impacts of categories on visual statistical learning in the presence of systematic visual similarity manipulations, and also considers how visual similarity might be modulated by statistical learning. Chapter 4 focuses on how different familiarization tasks may influence the behavioral and neural correlates of visual statistical learning using brain imaging (fMRI). Together, the results from these experiments demonstrate that visual statistical learning is often altered

depending on contextual factors that would be expected to fluctuate in everyday contexts.

Chapter 1

INTRODUCTION

Learning is oftentimes thought of as an effortful process. For example, students may study to learn content before an exam, a person may recite a new acquaintance's name a few times to make sure they've learned it well enough not to forget, or a child will learn to balance a bike through repeated attempts at remaining upright. However, there are also many instances where learning is not an effortful or intentional process (it is 'incidental'): a road closure may force an individual to expand their spatial memory for the city they live in, or an individual may learn a stranger's name after hearing somebody else refer to that stranger by name. As we consider the complexity and mass of information we regularly engage with, and the various contexts in which we encounter that information along with our own specific goals as we navigate our world, many other examples of incidental learning are evident. The varieties of incidental learning are not fully known, and their role in cognition is still poorly understood, but laboratory examples of incidental learning have led to progress in cataloguing and understanding these phenomena. The present proposal focuses on a specific form of incidental learning that is involved in extracting regularities from the environment: statistical learning.

1.1 Incidental Learning

Anecdotal examples of incidental learning, where learning can occur without putting forth intention and effort to learn, are plentiful. For instance, people who drive may not always study the location they parked their car at the supermarket in an effort to remember later, but rather may rely upon an incidentally formed memory when they return to their vehicle. Although this ability may be taken for granted (e.g., failing to causally remember if you turned off the stove can have serious consequences), typical functioning for most people depends on the ability to passively process vast amounts of information on a day-to-day basis. Insight into exactly *what types of information* people can learn in this manner and *how* they accomplish this is the subject of intense scientific interest.

One early example scrutinizing an ability to incidentally learn (and subsequently benefit from) come from the work of Nissen and Bullemer (1987). In a serial reaction time task, participants viewed stimuli that appeared in four potential locations on a computer screen, and produced a speeded response based on the location of their appearance. Unbeknownst to participants, the stimuli could appear in a pre-determined sequence multiple times. As participants incidentally learned the sequence, reaction times decreased for stimuli within the later parts of the sequence. This pattern was observed for both neurotypical participants, some of whom could explicitly recall the sequence of stimuli at the end of the experiment when asked, as well as participants with anterograde amnesia associated with Korsakoff's syndrome who could not explicitly recall the sequence.

Additional insight into how we may benefit from incidental learning comes from Chun and Jiang (1998), where participants viewed visual search displays in which they had to identify a target within a set of distractors. Importantly, the configurations of targets and distractors (i.e., the relative locations of distractors and the target location) could repeat in some of the trials. Experimenters did not inform subjects about this repetition, but participants were able to incidentally learn these regularities and respond faster on subsequent trials; a phenomenon known as contextual cueing. Considering our understanding that we are faster to search and respond to familiar real-world contexts (e.g., from Biederman and colleagues (1982) work, a fire hydrant is easier located in a street scene than a living room scene), contextual cuing explains one way in which we may learn and adapt to processing visual scenes.

These examples of incidental learning are useful, but they are contingent upon task demands. In other words, in the work discussed thus far, incidental learning can benefit the subject according to task demands. The indices of learning sequences from Nissen and Bullemer (1987), as well as the global contexts from Chun and Jiang (1998), were increased accuracy and decreasing reaction times in response to patterns. Since subjects are motivated by accuracy and time spent on task, improvements in such provide a feedback signal that can lead directly to benefits in task performance during acquisition. This leaves open the question of what forms of incidental learning might occur independently of task (if there is even a task), and how might we benefit from such forms of incidental learning. To address these questions, we turn our discussion toward a phenomenon known as "statistical learning", which was initially pursued scientifically as a potential solution to long-standing problems in language learning. However, the phenomena discovered by those efforts proved to be ubiquitous in sensory and cognitive systems, leading to an explosion of research that crossed domains of Psychology and Neuroscience.

1.2 Statistical Learning

Across languages, adult speakers display the ability to discriminate individual units of language, even though such units often do not correlate well with clear physical signals that differentiate them from neighboring signals. How, then, is it possible for babies to acquire language without pre-existing knowledge about how individual units are separated? Saffran and colleagues (1996) argued that infants must learn to identify meaningful acoustic cues that designate word boundaries, because even when word boundaries are acoustically ambiguous, infants appear to be able to segment fluent speech into sounds. They hypothesized this was possible via the extraction of statistical regularities within speech sounds. To test the viability of this hypothesis, they presented babies with a continuous speech stream comprised of four three-syllable nonsense words (e.g., bidaku, padoti, golabu, etc.). These four words were presented in randomized order for two minutes, with no identifiable separation between words (e.g., bidakupodotigolabubitaku…). After this familiarization phase where babies were expected to learn the statistical regularities within the stream, Saffran and colleagues created foil words comprised of the first two syllables of one target word combined with the last syllable of another target word. For example, if “bidaku” and “padoti” were words in the original stream, associated foils would be “padoku” and “bidati”. Babies displayed shorter durations of fixation (listening) for the target words as compared to the foil words. Since babies show preference for novel information, the reduced listening time given to target triplets suggests that information about the sequential statistics of syllables had been learned.

This sort of learning may be critical for early language development, but similar evidence of statistical learning has been replicated in an adult population (Saffran et al., 1999). The abstraction of statistical regularities from spoken sounds at any age suggests that making sense of a complex auditory world may require the adaptation afforded by constant learning throughout the lifespan. An important question raised by this work was whether such learning was constrained to auditory stimuli, or whether it could occur in other sensory modalities or even at a higher, cognitive level that abstracted away from the sensory stimulus.

Fiser and Aslin (2001) were first to demonstrate evidence of statistical learning in the visual domain, by presenting participants with static displays of shapes within a grid (see Figure 1). Covertly structured pairs appeared regularly within the display; participants were not aware of the structured pairs while engaged with the task. Participants were only instructed to pay attention to the continuous stream of displays.

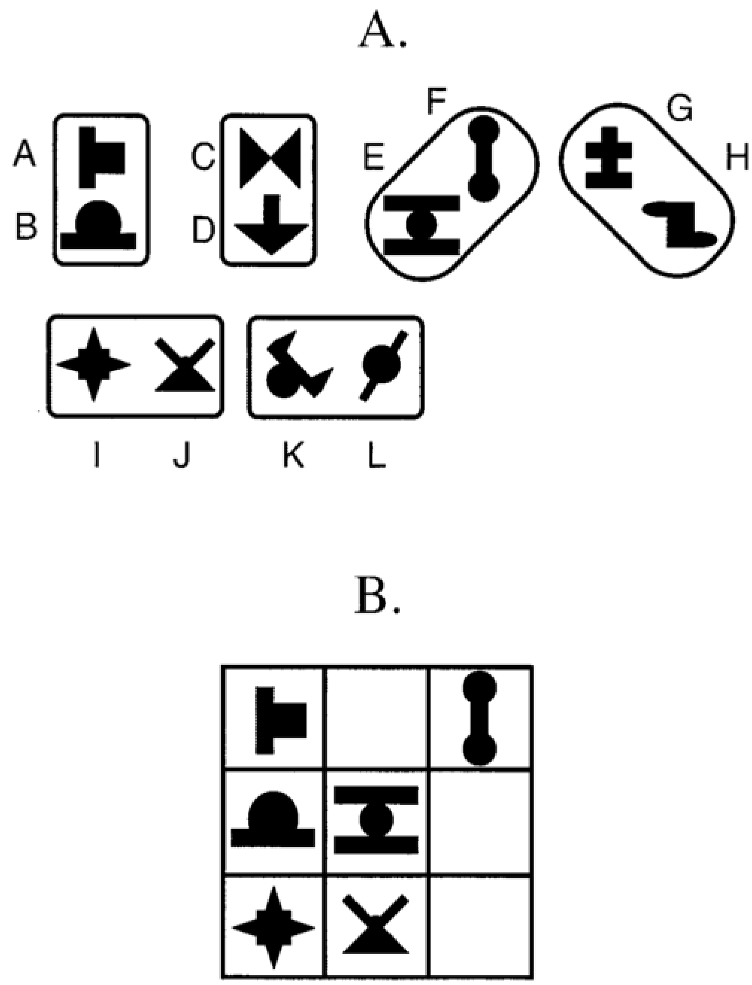

Figure 1. from Fiser and Aslin (2001) A) Structured pairs that appeared during the exposure phase. B) A sample display, of which there were 144 that subjects passively viewed during the familiarization phase, containing the structured pair “A-B”.

Participants viewed 144 displays with the embedded structured pairs, after which they were given a two alternative forced choice test task. Within the test, participants were shown one of the structured pairs and one non-base pair (see Figure 2) and were asked to choose the pair that appeared “more familiar”. In Experiment 1, the non-base pairs were constructed from shapes that were not presented as a structured pair both by identity and spatial location. In Experiment 2, however, non-base pair shapes had appeared in the tested cells of the grid during the familiarization phase equally as often as the structured pairs. This manipulation in the second experiment allowed for spatial frequency to be controlled while the non-base pairs were only constructed of shapes that had never appeared together.

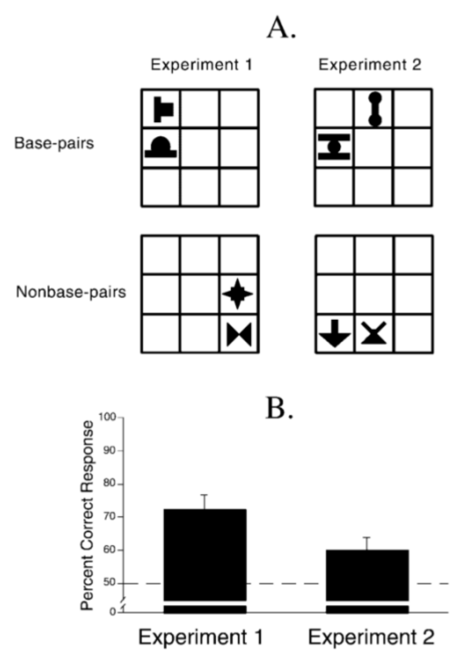

Figure 2 from Fiser and Aslin (2001). A) Example of a test-phase presentation of one base-pair and one nonbase-pair for Experiments 1 and 2. B) Participants’ accuracy in reporting the pairs that had previously occurred (Experiment 1) even when the spatial frequency of the shapes was held consistent (Experiment 2).

Despite matched joint probabilities of the non-base pairs and some of the structured pairs during familiarization, Fiser and Aslin (2001) found that participants still identified the structured pairs as more familiar over the frequency matched non-base pairs (Figure 2). These experiments suggest that humans are also able to automatically extract higher-order statistics from visual scenes.

Extending their findings from spatial visual statistical learning, Fiser and Aslin (2002) found evidence of temporal learning using higher-order temporal structures. In their experiment, participants watched a movie that presented 12 shapes covertly grouped into four triplets (Figure 3). While still unaware of the grouping, participants were again asked to simply passively view the presentation of the shapes.

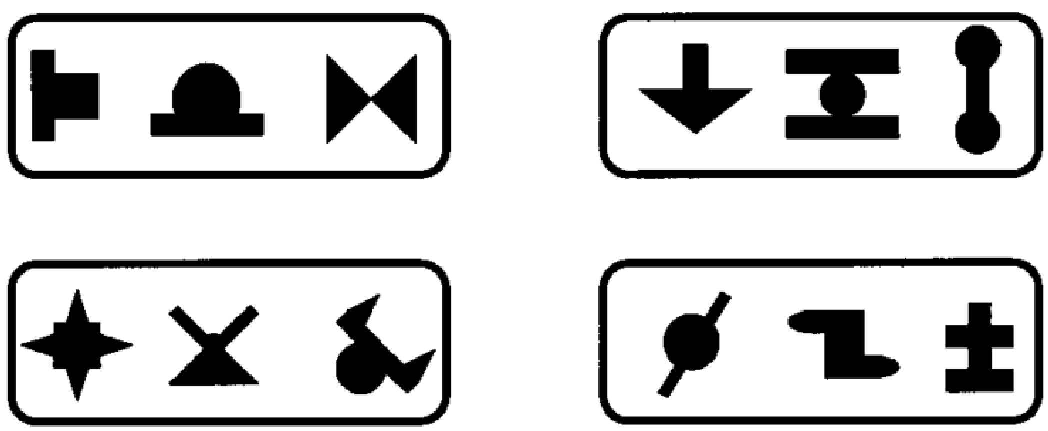

Figure 3 from Fiser and Aslin (2002). The 12 basic shapes participants were exposed to were grouped into triplets without participants' knowledge.

The movie displayed to participants showed a single shape at a time appearing from behind a vertical occluder from one side and returned behind the occluder on the other side (Figure 4). Each shape was on screen for one second and each triplet of shapes appeared a total 96 times in semirandom order to avoid triplet repetitions.

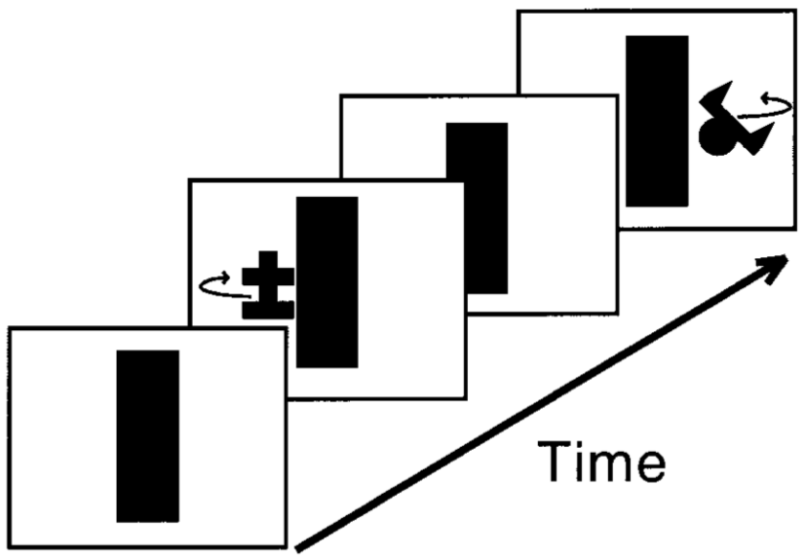

Figure 4 from Fiser and Aslin (2002). A movie was shown to participants to display a continuous stream of shapes. Unbeknownst to participants, shapes always appeared in their respective triplets and triplets were pseudorandomized to prevent triplet repeats. During presentation, any one shape was on screen for one second before being replaced by the next shape.

After familiarization, participants were again given a two-alternative forced choice task in which they had to choose the triplet that was more familiar. In the test task, they chose between a base triplet (from the presentation phase) and a nonbase triplet (constructed from characters that were not learned to be predictive of one another from the movie). Each triplet was displayed in the same context of the movie with one character appearing after another. Participants correctly chose the base triplet over the foil triplet with overall accuracy at 95%, unequivocally displaying an ability to learn temporal-order regularities in a visual context and in an unsupervised fashion. The findings from this study along with that of Fiser and Aslin (2001) suggest that, depending on the context in which statistical regularities are presented, participants can incidentally learn contingencies in both spatial and temporal domains.

1.3 Utility of Visual Statistical Learning

Thus far, we have considered visual statistical learning as a laboratory phenomenon indexed by subsequent recognition. What potential ecological usefulness could the visual system derive from such learning? Clues to the answer to this question could be found by considering what other systems may be supported by such stimulus-stimulus associations. Previously, we have discussed the application of auditory statistical learning to provide insight into how infants may segment speech patterns into meaningful and separable representations that serve as a basis for language (Romberg & Saffran, 2010; Saffran et al., 1996). Given the focus of the present dissertation on statistical learning in the visual domain, we must also consider similar applications of statistical learning to visual input. In other words, when we discuss visual statistical learning as something that "makes sense" of our visual world, what exactly do we mean?

Not unlike the ambiguity of auditory boundaries within a stream of speech sounds, like that used by Saffran and colleagues (1996), different pieces of visual information can also be challenging to distinguish. Typically, the parsing of different objects in a visual scene is thought to rely heavily on visual boundary information (Peterson, 1994; Riesenhuber & Poggio, 1999; Zhou et al., 2000) which is often contingent upon luminance contours (Kellman et al., 1983; Palmer et al., 1994; Spelke, 1990). But even then, there are plenty of cases where, with these low-level perceptual cues, we may fail to extract some object representations (Kellman & Shipley, 1991). Rather than focus on other visual features predictive of successful extractions of object representations from visual scenes, Lengyel and colleagues (2021) argue that consistent statistical properties within a scene underlie the formation of object representations.

To test this, participants first engaged in a spatial visual statistical learning task similar to Fiser and Aslin (2001) for participants to incidentally learn "visual chunks" of information. Most compellingly in their second experiment, they used a well-established object-based attention paradigm (Figure 5), which has previously been used to highlight how attention may benefit from a cue appearing within an object as compared to between objects. Thus, in a speeded response task where participants had to find a target after being exposed to a cue, Lengyel and colleagues (2021) were able to provide evidence of "chunk-based attention" that followed a pattern similar to the traditional "object-based attention" findings.

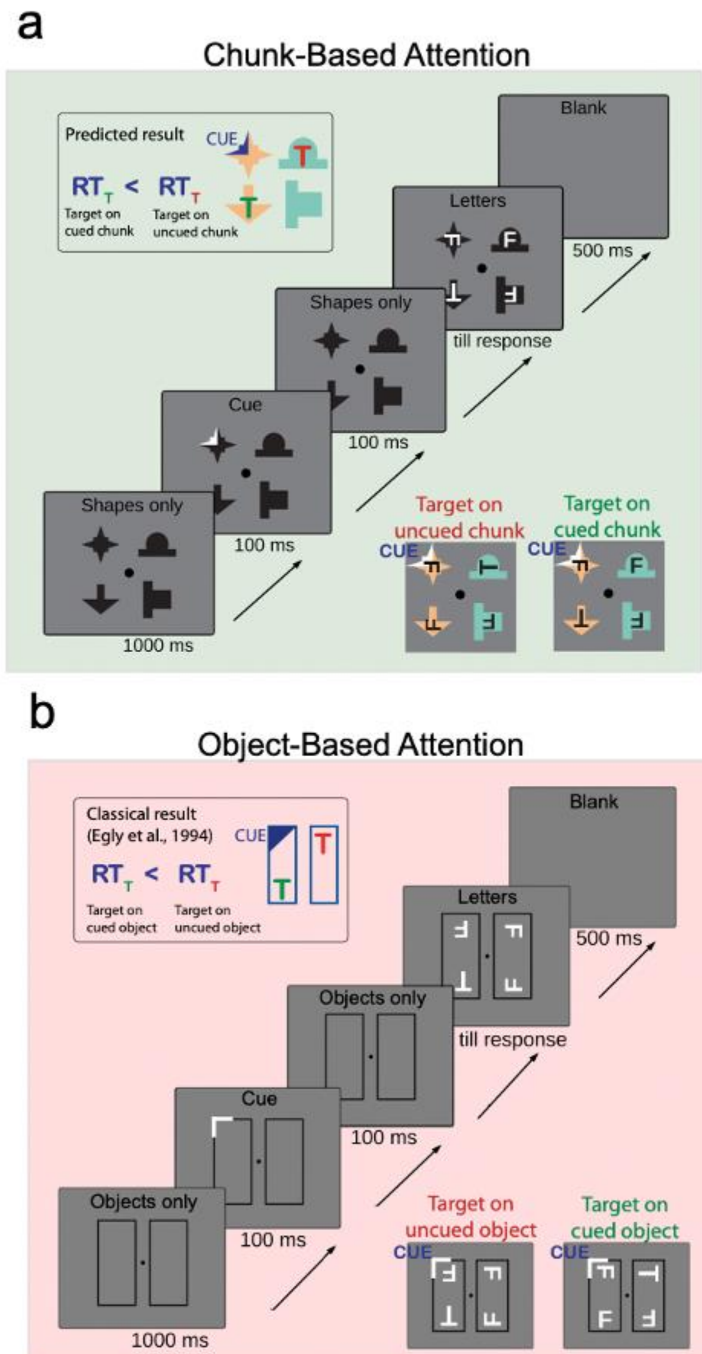

Figure 5 from Lengyel and colleagues (2021). a. The chunk-based attention task where "chunks" represented visually learned regularities. b. The traditional object-based attention task.

Taken together, the findings of Lengyel and colleagues (2021) provide insight into how visual information can be segmented into object-like "chunks" formed through visual statistical learning that does not rely on the presence of clear visual boundaries. This is comparable to Saffran and colleagues (1996) who highlighted infants' ability to segment fluent speech that also lacked clear auditory boundaries. Both of these cases are useful in considering how statistical learning benefits us in our ability to "make sense" of the world in different modalities. However, in considering the ways in which visual statistical learning is applicable to our day-to-day functioning, we are rarely exposed to statistical regularities in a context that demands passive observation (like that used in Lengyel and colleagues' passive viewing task). We oftentimes move through different environments with different goals in mind, and these multiple contexts we encounter throughout the day rarely present opportunities to receive visual information without some goal or task in mind. Thus, it is also worthwhile reviewing literature that identifies how such contextual influences can influence visual statistical learning.

1.4 Task Influences on Visual Statistical Learning

Evidence of statistical learning extending into the visual domain provides a basis from which we can consider how the human cognitive system is able to automatically and incidentally learn from everyday visual experiences. However, it is important to investigate how such learning responds to variations in task, stimulus attributes, and other manipulations of context during familiarization, to better understand its potential impact on everyday cognition. For example, human visual performance is known to be highly limited in certain ways – access to working memory and other higher cognitive functions faces a bottleneck in the form of selective attention (Akyürek et al., 2007).

Do humans learn any statistical regularity they are exposed to, or is learning constrained to the objects of selective attention? Turk-Browne and colleagues (2005) found evidence that visual statistical learning is modulated by selective attention. Using the same basic shapes from Fiser and Aslin (2001, 2002), participants viewed an interleaved stream of red and green shapes. They were told to monitor only one color for immediate repeats of a shape, while statistical structure was covertly present in the stream (see Figure 5 below from Turk-Browne and colleagues (2005).

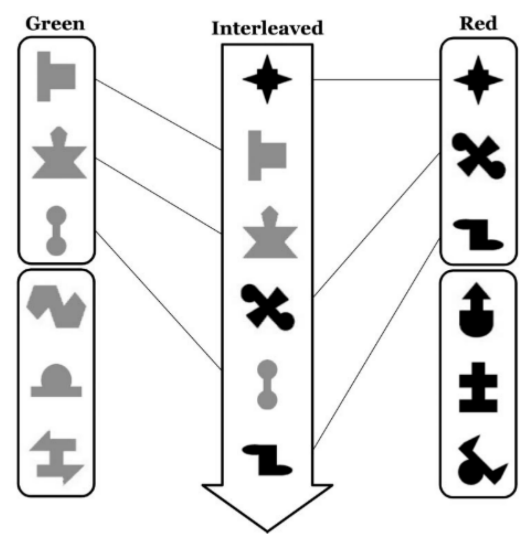

Figure 6 from Turk-Browne and colleagues (2005). Participants viewed a stream of images and were instructed to monitor one of two colors for immediate repeats. Regularities appeared within-color during the stream of images.

Turk-Browne and colleagues (2005) found that only regularities within the attended color were learned at above-chance levels. Thus, visual statistical learning requires some level of attentional selection to extract regularities; incidental learning of this kind does not appear to occur for ignored information. These findings provide useful insight about how contextual effects on attentional state (such as those driven by task demands) can influence visual statistical learning, and are echoed by other works that use the same basic shapes as Fiser and Aslin (2001, 2002) to investigate attention-related influences on learning (J. Zhao et al., 2013; Jiaying Zhao et al., 2011).

The use of basic shape stimuli in the previously discussed works allowed researchers to gain important insights into visual statistical learning while carefully controlling for extraneous influences. However, if our aim is to expand our understanding about how visual statistical learning may operate on everyday visual experiences, there appear to be many ways we can build from these foundational works. Specifically, we may consider how variations in learning contexts or prior experience with stimuli may shape learning. Prior work has found evidence of statistical learning using semantically rich stimuli (Brady & Oliva, 2008), and we can look to more recent work for insight into how visual statistical learning may operate over a combination of existing knowledge and task demands.

Faces and scenes are commonly encountered visual stimuli that are naturally complex (in comparison to basic shapes) and possess relatively rich semantic and categorical information. Vickery et al. (2018) explored how varying tasks demands may influence visual statistical learning. Their final experiment characterizes the importance of both pre-existing knowledge for stimuli (i.e., prior experience with faces and scenes) and the impact of task on visual statistical learning. Participants viewed a stream of male/female faces and indoor/outdoor scenes. Half of participants were assigned to a “jiggle” group and were tasked with pressing spacebar whenever an image jiggled back and forth. The other half of participants were assigned to a categorization group and were required to make a button press with one hand for male faces/outdoor scenes and a button press with the other hand to female faces/indoor scenes. Unbeknownst to participants, 32 images were organized into 16 pairs which allowed for four different pair types (see Figure 7a for a full visualization). Thus, both groups viewed the same stream of images (and occasional jiggle) but were given two different tasks.

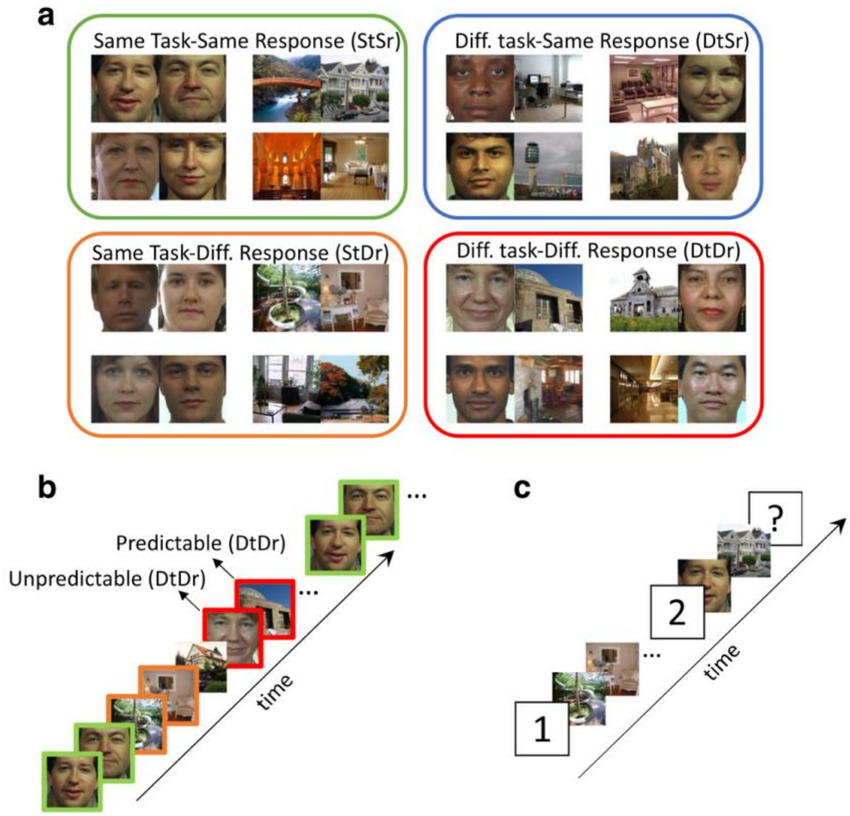

Figure 7 from Vickery and colleagues (2018). a) For the categorization group, pairs of images constituted four conditions dictated by same or different task (i.e., same or different category), and by same or different response (i.e., response made with the same hand or different hand). b) Example stream with images in pairs highlighted with a color border (participants did not see a color border). Only the second image of a pair could be predicted, and the magnitude of the predictive value between pairs was hypothesized to be modulated by the amount of task/response information shared between those images. c) Example of a recognition test trial at the end of the experiment.

Participants in the detection group displayed above-chance and roughly equal learning across conditions (Figure 8), while the categorization group displayed a larger effect of recognition for StSr pairs over all other pairs. Additionally, the only other condition to display above-chance recognition in the categorization group was DtSr pairs, which the authors took as evidence of response similarity acting as either a learning impetus of same-response pairs or a learning inhibitor for different response pairs. This work presents strong behavioral evidence that task demands influence visual statistical learning as well as how learning may be modulated by task-related or stimulus-related contingencies.

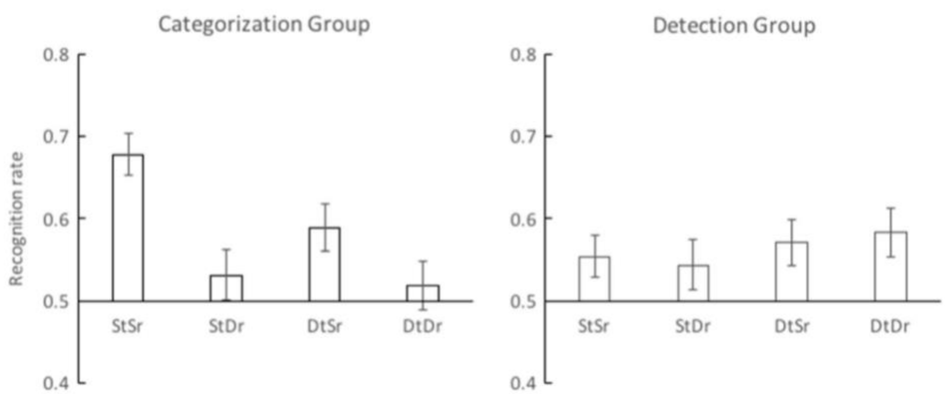

Figure 8 from Experiment 4 of Vickery et al. (2018). Equal above-chance learning is observed for the detection group, while the categorization group displayed greatest learning for Same-Task (i.e., same category) Same-Response pairs, above chance learning for Different-Task (i.e., different category) Same-Response pairs, and no significant learning for Same-Task Different Response pairs or Different-Task Different-Response pairs.

However, the introduction of complex stimuli such as faces and scenes does not leave the interpretation of the results unequivocal; how are the effects observed from Vickery and colleagues (2018) influenced by categorization as opposed to the simple fact that two face images or two scene images are more visually similar than across-category combinations, on average? Would such category-driven effects be sustained if participants were not explicitly making category-related decisions during familiarization? Does drawing attention to categories evoke similar or dissimilar mechanisms of learning and representation, as opposed to tasks that do not draw attention to categories? The present proposal seeks to gain insight into some of these questions using behavioral and neuroimaging methods. Therefore, it is also important to discuss a handful of fMRI studies that shed light on activity in brain areas that correlate with task demands or category-related influences on visual statistical learning.

1.5 Neural Correlates of Visual Statistical Learning

Thus far, we have discussed work that focuses on how task demands and categorizable stimuli may influence visual statistical learning, but there are still open questions about how exactly these influences impact learning. Some of these questions may be best answered with neuroimaging, but we must first consider typical findings associated with statistical learning and categorization. Turk-Browne and colleagues (2009) used fMRI while participants responded to whether basic shapes presented to them "jiggled" back and forth, obscuring the fact that on some runs the shapes were covertly structured into triplets while other runs contained no statistical structure. During structured streams, and consistent with other types of associative learning and memory, greater activity was observed in the medial temporal lobe (MTL) and striatum. Secondly, these results were partially supported by later work conducted by Turk-Browne and colleagues (2010), who investigated the neural underpinnings of visual statistical learning while participants engaged in a face/scene categorization task. Their study, which utilized event-related fMRI, required participants to respond to a stream of face and scene images while lying inside the scanner. The stream of images either appeared as a singleton (an unpaired face or scene image), or in a scene-face or face-scene pair (Figure 9), pseudorandomized to preserve pairs across runs. While viewing singletons and the first image of pairs, greater activity was observed in the right anterior hippocampus, which the authors argue may be related to a unique role of perceptual anticipation (but no activity in the medial temporal lobe by the same standards).

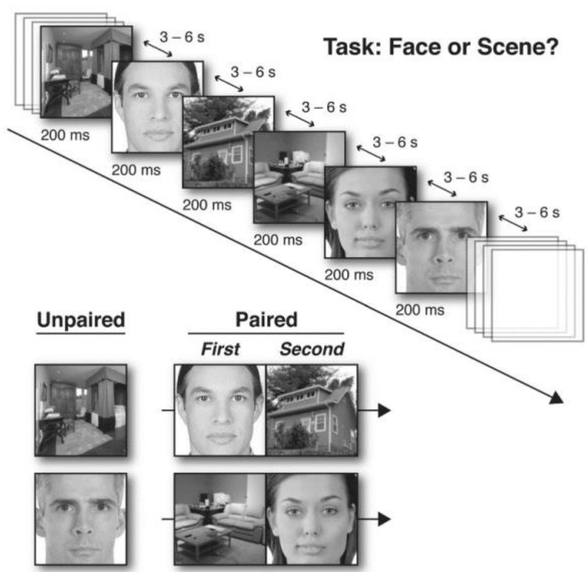

Figure 9 from Turk-Browne et al. (2010). The event-related design that included singletons and pairs of face-scene or scene-face stimuli.

Although the pairs used by Turk-Browne and colleagues (2010) did not contain any same-category pairs like those found in Vickery and colleagues (2018), the authors did find some interesting category-specific effects. In particular, bilateral activity in the parahippocampal place area (PPA) appeared to uniquely discern both faces and scenes that appeared either as a singleton, as the first image of a pair, or as the second image of a pair (Figure 10). This highlights an important interaction in category related activity and visual statistical learning; activity in the PPA is able to predict both the category of the upcoming stimulus and the amount of statistical information bound to that stimuli.

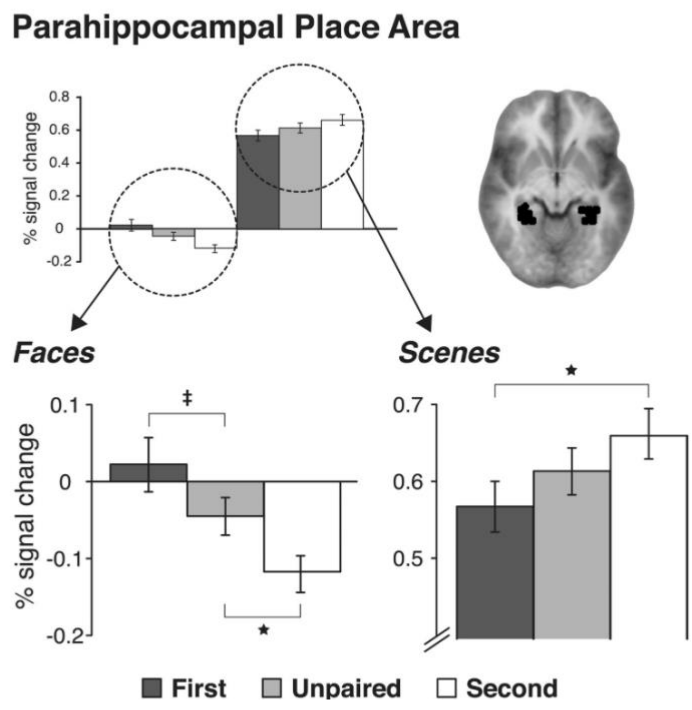

Figure 10 from Turk-Browne et al. (2010). Differential activity in the PPA distinguishes the category of image being viewed as well as the position in a statistical or non-statistical representation.

These works by Turk-Browne and colleagues (2009, 2010) highlight some fundamental insights into how we may consider future research into task demands that influence statistical learning. Vickery and colleagues (2018) explored how categorization of naturalistic stimuli may influence visual statistical learning, but no neuroimaging work has sought to uncover task-specific impacts while also controlling for some of the potentially confounding variables touched on previously (e.g., visual similarity between items of the same category). Additionally, and in anticipation of task-related influences on visual statistical learning, it is worthwhile considering the unique insights provided by Schapiro and colleagues (2012) and how patterns of activity within the hippocampus specifically are altered by visual statistical learning.

Schapiro and colleagues (2012) exposed participants to a 40-minute sequence of colorful fractals (Figure 11A, bottom) while being scanned using fMRI. During this exposure phase, participants were given a cover task where they had to respond to infrequent greyscale fractals. Similar to other statistical learning tasks, this cover task served to obscure the underlying structure of the fractals within the stream (Figure 11A, top); each fractal could be followed by the same fractal each occurrence, designating a “strong pair”, or could be followed by the same fractal 1/3 of the time designating a “weak pair”. Participants were also exposed to the fractal images in completely random order once at the beginning of this phase and once at the end of this phase. Finally, rather than attempting to test participants immediately after the final run where pair structure was disrupted, a separate behavioral study was conducted where the final block of familiarization (which possessed pair structure) was immediately followed by a test task that indicated participants had successfully learned the strong pairs better than the weak pairs.

Figure 11B displays some findings from Schapiro and colleagues (2012). Patterns of activity were extracted from each ROI and for each fractal image and correlated to produce the example correlation matrices (before learning appearing above, and after learning appearing below). Fractal images that constituted strong pairs predicted a significant correlation in activity in the hippocampus after learning had taken place. In other words, activity associated with the 1^{st} image of a pair was correlated with activity associated with the 2^{nd} image of a pair before and after learning took place (these correlation matrices were also computed for weak pairs and shuffled pairs, respectively). The pattern correlation between strong pairs of fractals was increased across the entire hippocampus using ROIs focused on the subiculum, CA1, and CA2/CA3/dentate gyrus (and also included the perirhinal cortex).

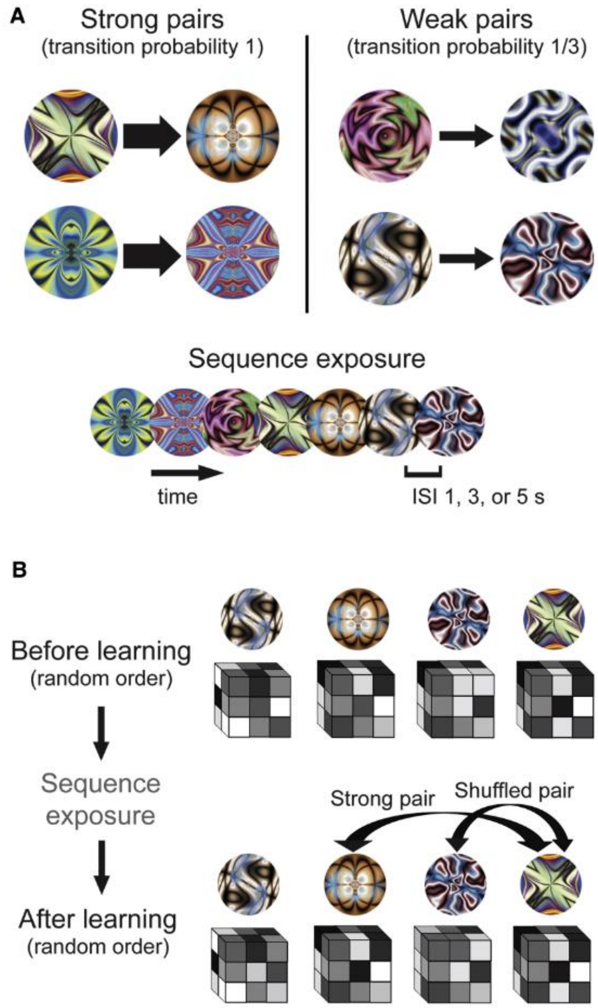

Figure 11 from Schapiro et al. (2012). (A) Strong pairs and weak pairs used. (B) Pattern similarity for each fractal before and after familiarization.

Interestingly, this pattern was not observed for any other condition other than strong pairs (i.e., weak pairs or shuffled pairs), and the authors acknowledge the learned "strong" pairs as object representations formed by the statistical learning of temporal regularities. These insights have contributed to our understanding of the multiple learning systems at work within the hippocampus (McClelland et al., 1995), and with respect to the present dissertation's focus on multiple forms of learning, this work has also provided a basis upon which researchers have been able to model separate anatomical pathways within the hippocampus that are thought to be involved in different types of learning (Schapiro et al., 2012). This work will be considered further in Chapter 5, with our attention now shifting toward the multiple types of learning and stimulus features that could potentially influence visual statistical learning.

1.6 Visual Statistical Learning, Category Learning, and Similarity

Both Vickery and colleagues (2018) and Turk-Browne and colleagues (2010) provide insightful work into how prior knowledge and tasks generally influence visual statistical learning. However, these works also leave open several important avenues of research that can lead to a much more complete understanding about the interaction of context and visual statistical learning. For example, Vickery and colleagues found that differences in learning pair types depended on participants' instructions (respond to a jiggle or actively categorize the stimuli). Those differences found during categorization suggest the greatest learning for same-task same-response pairs (e.g., two face images paired together that require the same response), but as mentioned previously, couldn't it also be argued that two faces or two scenes share a great deal of visual similarity? How much of this effect is due to stimuli belonging to the same group of a well-learned category, and how much of it is due to the two images of a pair sharing a great deal of low-level visual features? This distinction has yet to be made.

As discussed previously, statistical learning is evident from infancy (Aslin et al., 1998), consistently engaged in processing visual stimuli in adulthood (Otsuka et al., 2013; Turk-Browne et al., 2005; Vickery et al., 2018), but it also continues to be apparent into late adulthood (Campbell et al., 2012). Despite the apparently lifelong utility of this learning phenomenon, the literature thus far has somewhat neglected how other kinds of experiences encountered throughout a lifetime may modulate visual statistical learning. Categories seem to play an important role in visual statistical learning, as evidenced by the work of Vickery and colleagues (2018) and Turk-Browne and colleagues (2010), but several fundamental questions remain if we aim to grasp how category learning and visual statistical learning interact. For this reason, it is critical to first consider the similarities and differences between visual statistical learning and category learning.

Not unlike visual statistical learning, category learning plays an important role in processing the massive amounts of stimuli we are faced with on day-to-day basis, and in gaining a means to generalize from specific learning episodes to different circumstances. Category learning may be described as our ability to group things (e.g., stimuli) based on some shared feature or attribute. This is also thought to be essential for survival in several different ways (e.g., is this something that is good or bad to eat?) and is most often studied in the lab in contexts where participants are directed to learn, feedback is provided to participants, and learning goals are made explicitly clear (Ashby & Maddox, 2005). This is in stark contrast to statistical learning, in which 1) all information about statistical regularities may be hidden and 2) unlike other forms of incidental learning (Chun & Jiang, 1998; Nissen & Bullemer, 1987), even when participants are engaged in a task irrelevant to statistical regularities, statistical learning will extract those regularities.

Additionally, while visual statistical learning may be considered a form of associative learning, research into category learning is diverse and may extend beyond simple grouping of stimuli and apply more broadly to concept formation (Ashby & Maddox, 2005; Richler & Palmeri, 2014). Critically, category learning is often studied in the context of distances within continuous similarity spaces (Shepard, 1987), with boundaries imposed on similarity spaces, or based on resemblance in terms of the presence of various discrete characteristics (Tversky, 1977). Consider Figure 12 from Richler and Palmeri (2014) for an example of stimuli used in visual category learning tasks. Perceptual differences in these categorization tasks can vary on few dimensions, (b) providing an example of two-dimensional stimuli varying in shape, size, and shading only, or multiple dimensions like (f) where a three dimensional “greeble” can also vary in body/appendage shape.

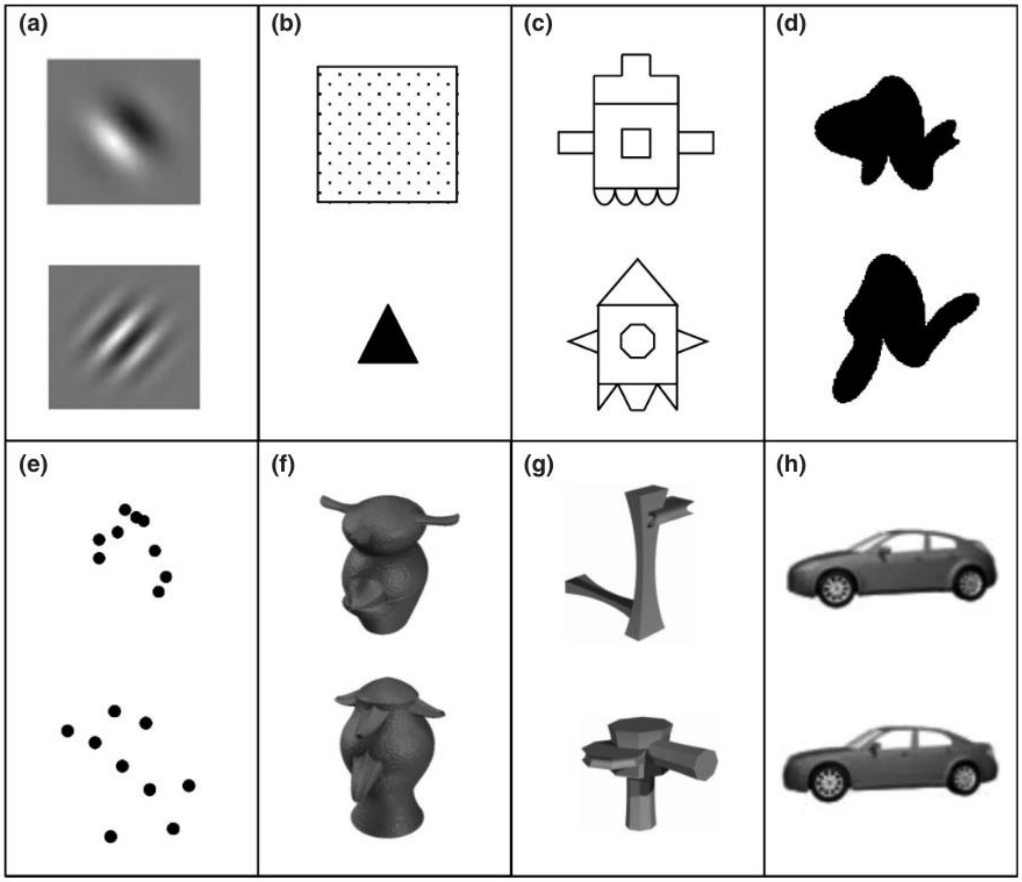

Figure 12 from Richler and Palmeri (2014). Eight examples of object sets used in visual category learning experiments varying in many or few dimensions.

While perceptual categorization operates systematically over multiple visual features and dimensions, categories can also be abstract or “ad hoc” (Barsalou, 1983). For example, “things that need to be donated to charity” or “stuff that I will take to the library today" are both concept-driven categories that are subjectively applied to stimuli that possess any number of visual features.

Categorization is an invaluable tool by which we may organize and make sense of the world on a moment-to-moment basis. It enables us to generalize properties from one instance to another, to achieve invariance with respect to feature variation that does not impact categorical membership, and to track collections bound together by common goals, origins, and hidden properties. In recognizing this, we may be able to ask how both category learning, or more specifically, perceptual category learning and all the goal-directed effort that goes into it, may work in tandem with visual statistical learning.

Apart from the handful of studies previously discussed, statistical learning is typically studied using novel stimuli that do not fall into obvious previously learned categories, or do not have an obvious arrangement in a feature space. Additionally, visual statistical learning has been discussed as being unique from other types of incidental learning; extracting regularities from a visual environment can occur even when it does not benefit the observer. This leaves open an important question: how might statistical learning be altered as a function of categorical knowledge? If categorization and statistical learning both play critical roles in cognitively organizing the visual world, it stands to reason that an incidental learning process like visual statistical learning may be modulated by explicit and/or pre-existing knowledge, such as category information. More so, considering how categories are so often defined by distinct visual features, but visual features have traditionally been controlled in the statistical learning literature, how might stimulus-level visual differences interact with statistical learning? With an understanding of the previously reviewed works that used visually distinct categorizable stimuli in a statistical learning context, scrutinizing these interactions will require refined control over the information shared between structured representations, as well as control over the specific context by which this information is presented. To begin, it is worthwhile raising several questions considering the aforementioned research that focuses on influences on visual statistical learning.

First, how much of the effect of category on visual statistical learning has to do with the act of categorizing stimuli into one group or another, and how much of it has to do with a lifetime of experience we have with features that constitute stimuli such as faces and scenes? Likewise, how much of the enhanced learning for same category pairs, as observed in Vickery et al. (2018) is due to the perceptual similarity shared between the stimuli within a pair? How might perceptual-driven categories drive statistical learning, or how might statistical learning inversely impact perceptual similarity? Finally, if context influences visual statistical learning, how much of the prior neuroimaging work are supported by a scenario where the neural correlates of learning are compared between differing tasks?

1.7 Overview of the Present Dissertation

The present dissertation puts forth a set of experiments that systematically controlled for prior knowledge, perceptual similarity, and effects elicited by task. In

Chapter 2, a series of experiments are presented to disentangle task-related influences of visual statistical learning from visual similarity between images of the same natural category group (e.g., better learning for face-face pairs cannot yet be determined to be due to the visual similarities shared between two faces or the fact that the two images fall within the same pre-learned "face" category). This was accomplished by controlling for prior knowledge by having participants learn arbitrary category groups while diminishing similarity-driven differences across stimuli through the use of fractal images that are visually unique from one another but do not easily fall into perceptual categories like faces and scenes. Chapter 3 used basic shape stimuli designed to be visually similar or dissimilar to determine the role of perceptual similarity in visual statistical learning while probing any inverse impact of statistical learning on perceptual judgements. Finally, Chapter 4 adopted methods from Chapter 2 to control for prior knowledge and visual similarity but uses fMRI to attempt to scrutinize the neural correlates of potential task-related differences in visual statistical learning with a multi-task within-subjects design.

Chapter 2

MODULATION OF STATISTICAL LEARNING BY NATURAL AND ARTIFICIAL CATEGORIES

2.1 Experiment 1

The work of Vickery and colleagues (2018) provide substantial evidence that task demands and contexts may influence visual statistical learning. Same-category pairs (e.g., face-face or scene-scene pairs) were learned significantly better than different-category pairs (e.g., face-scene or scene-face pairs). However, while category membership is often dictated by visual similarity, it remains to be seen how much of this effect on visual statistical learning is due to stimuli being members of a well-learned category group and how much is due to the fact that these stimuli can share a great deal of low-level visual information. We designed a paradigm in which visual stimuli do not fall into previously learned natural category groupings. Participants must explicitly learn category group information for each image within a task designed to promote visual statistical learning. In other words, participants are engaged in a simple rule-based learning task where an arbitrary category grouping for each image must be learned. This experiment seeks to answer two important questions. First, in a task where participants must effortfully memorize arbitrary category group information, will incidental learning (i.e., visual statistical learning) persist? To date, there is no evidence suggesting visual statistical learning may occur simultaneously with effortful explicit learning. Secondly, will learned arbitrary category groupings modulate visual statistical learning similarly to natural category groupings (e.g., faces and scenes) with which participants have a lifetime of experience? If visual statistical learning persists in a task dominated by explicit learning, and if newly learned category groupings impact learning similar to previously learned natural category groupings, we may expect better learning for same-category pairs as compared to different-category pairs similar to that of Vickery and colleagues (2018).

2.1.1 Participants

A total of 30 University of Delaware Students participated in the study for partial completion of general course requirements.

2.1.2 Materials and Procedure

Stimuli consisted of 32 fractal images that were covertly placed into 16 pairs. Half of the images were randomly assigned to one category group (a “z” category which required a z response on the keyboard) and the remaining half were randomly assigned to the other category group (an “m” category which required an m response on the keyboard). Examples of fractal images and groupings are displayed in Figure 13. The 16 pairs of images consisted of 8 same-category pairs (e.g., both images within the pair belonged to the “z” category) and 8 different-category pairs (e.g., one image within the pair belonged to the “m” category while the other image belonged to the “z” category). Once in pairs, each image was presented onscreen for 1 second and participants were required to respond before the image disappeared. Feedback was provided in the form of a green fixation circle for correct responses that remained on screen for 1 second. If incorrect, participants were presented with a red fixation circle for 1.5 seconds. All images appeared in their respective pairs, and to further obfuscate participant awareness of the structure of the images, pairs were pseudo-randomized such no pair could immediately repeat or repeat with a single intervening pair (Figure 13). Participants viewed each pair four times per block with six total blocks of training. After this initial familiarization phase, participants completed a surprise test phase identical to that of Vickery and colleagues (2018), wherein each pair was presented temporally along with a foil comprised of two images that did not appear together. The order of the target/foil two alternative forced choice test was counterbalanced, and participants made their selection for all pair options before the experiment was concluded.

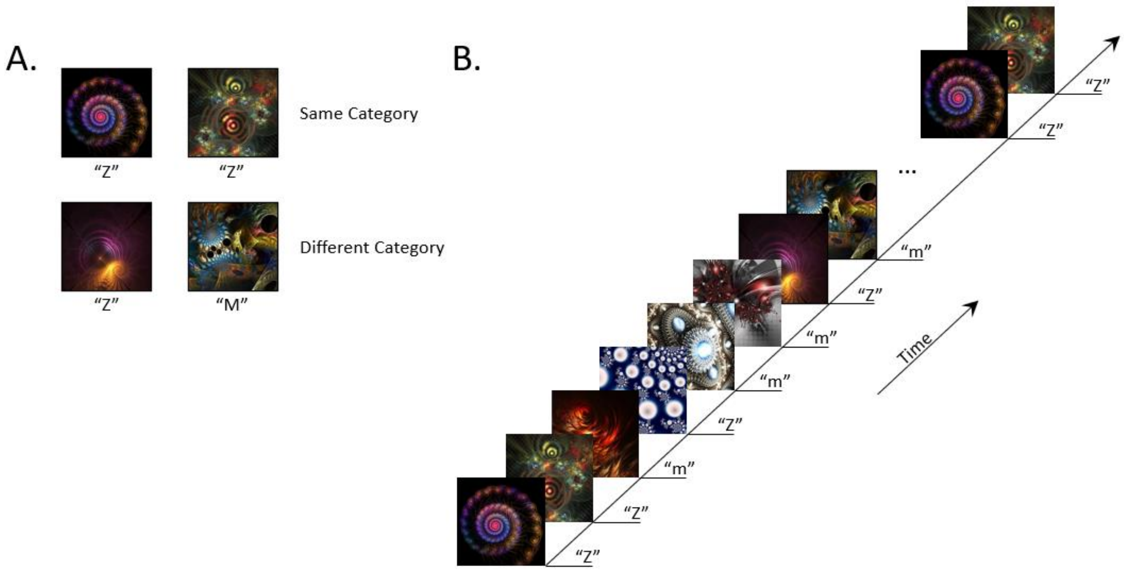

Figure 13. A) Example stimuli and pairings for Experiment 1. B) Example sequence of stimulus presentation during familiarization for Experiment 1.

2.1.3 Results

Learning for categories was significantly above chance by the end of block one $t(29)=2.08$, $p = .046$, Cohen's $d = 0.38$) and remained above chance for the remaining blocks (all $p < .001$). During the test phase, overall visual statistical learning was observed significantly above chance ($t(29) = 4.42$, $p < .001$, Cohen's $d = 0.81$). Learning for both different-category pairs ($t(29) = 2.19$, $p = .037$, Cohen's $d = 0.38$)

and same-category pairs ($t(29) = 5.1$, $p < .001$, Cohen's $d = 0.93$) was evident at above-chance levels. A comparison of conditions revealed that same-category pairs were learned significantly better than different category pairs $t(29)=2.173$, $p = .011$, Cohen's $d = 0.5$).

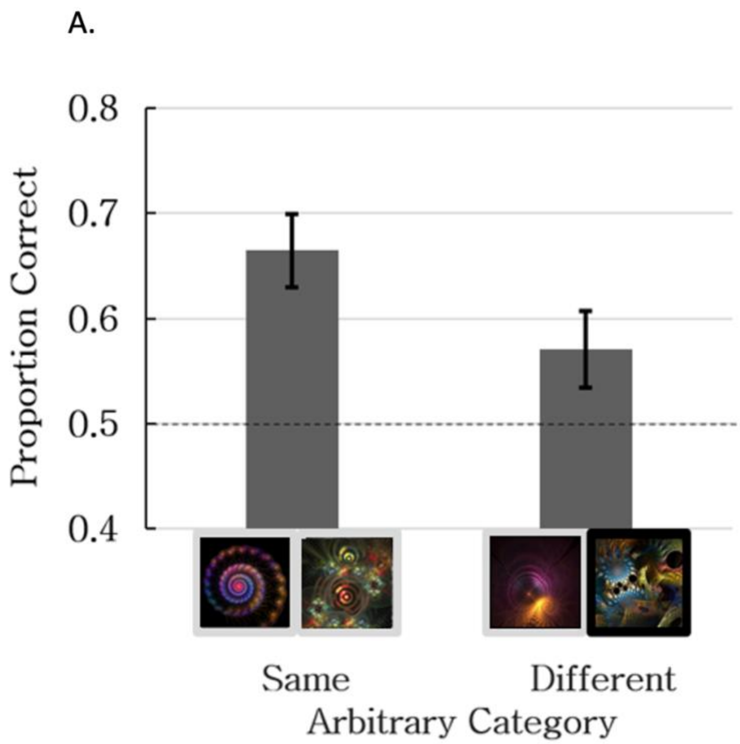

Figure 14. Preliminary data for Experiment 1. a) Mean category learning accuracy by block. b) Mean accuracy of recalling the target pairs during the test phase, broken down by same-category and different-category pairs.

2.1.4 Discussion

Overall, visual statistical learning was evident even in the presence of an explicit category-learning task. In other words, visual statistical learning persisted despite the cover task involving explicit, goal-directed learning of category information that was irrelevant to the regularities embedded within the stream. Additionally, learning the arbitrary category associated with each image while being exposed to statistical regularities also produced a difference in learning depending on learned category membership. Same-category pairs of images were better learned than different-category pairs of images, suggesting an interaction between visual statistical learning and arbitrary category learning. Together, these results suggest that newly learned category information for stimuli that do not already possess category information (but are still visually complex) influences statistical learning in such a way that shared category information leads to the greatest amount of learning for statistical regularities.

2.2 Experiment 2

Experiment 1 provided supporting evidence replicating that of (Vickery et al., 2018); pairs comprised of the same category group were learned better than pairs comprised of different category groups, even after controlling for well-learned and naturally occurring category groupings such as faces or scenes. Experiment 2 was designed to take a step back toward that of Vickery and colleagues while using arbitrary category groupings to avoid artificially drawing participants' attention to stimulus-level attributes. In other words, we replaced our fractal image set with faces and scenes but continued to assign stimuli to arbitrary category groupings. With this design we aimed to provide converging evidence to that expected from Experiment 1 (i.e., better learning for same category pairs) while also exploring the impact of category groupings that have already been learned (e.g., faces and scenes). If the task calls for explicit learning of novel category groupings, and the stimuli consist of images that already fit into pre-learned category groupings that consistently share perceptual similarity with other pre-learned group members, we expected an additive effect wherein image pairs consisting of congruent arbitrary and natural category groupings will be incidentally learned better than all other conditions. Additionally, we hypothesized that learning for pairs will decrease as the shared category group information is partially reduced (e.g., for congruent/incongruent arbitrary/natural pair mixtures) and reduced the most for pairs where no category group information is shared between the items of a pair. Thus, we did not anticipate the presence of any interaction wherein same natural category information led to increased learning for different arbitrary category information but not same arbitrary category information, or vice versa for arbitrary category information driving differential learning for natural category pairings.

2.2.1 Participants

A total of 30 University of Delaware Students participated in the study for partial completion of general course requirements.

2.2.2 Materials and Procedure

Stimuli again consisted of 32 images that were covertly placed into 16 pairs. Half of the images were randomly assigned to one category group (a “z” category which require a z response on the keyboard) and the remaining half were randomly assigned to the other category group (an “m” category requiring an m response on the keyboard). In lieu of fractal images, 16 face images and 16 scene images were used. Faces appeared as all male or all female (to avoid unintentional inter-gender categorization) and scenes appeared as all indoor or all outdoor. Gender and scene location were counterbalanced between participants. Examples of scene and face images and groupings are displayed in Figure 15. The 16 pairs of images consisted of 8 same-category pairs (e.g., both images within the pair belonged to the “z” category) and 8 different-category pairs (e.g., one image within the pair belonged to the “m” category while the other image belonged to the “z” category). Overall, there was an equal number of face/scene and category “z”/“m” members within each condition. Once in pairs, each image was presented onscreen for 1 second and participants were required to respond before the image disappeared. Feedback was provided in the form of a green fixation circle for correct responses that remained on screen for 1 second. If incorrect, participants were presented with a red fixation circle for 1.5 seconds. All images appeared in their respective pairs, and to further obfuscate participant awareness of the structure of the images, pairs were again pseudo-randomized such that no pair could immediately repeat or repeat with a single intervening pair (consistent with Experiment 1, Figure 13) Participants viewed each pair four times per block with six total blocks of training. After this initial familiarization phase, participants completed a surprise test phase identical to that of Experiment 1.

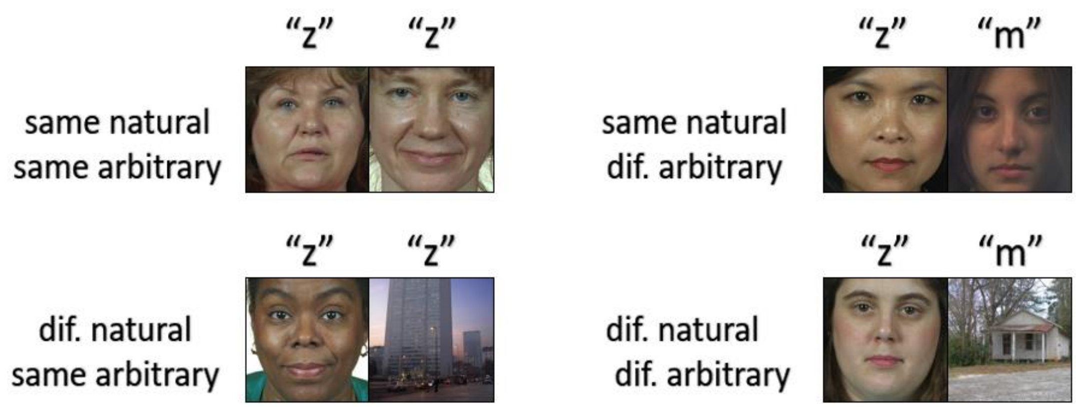

Figure 15. Four pair types from Experiment 2, presented primarily using face stimuli. An equal number of face and scene stimuli were used (i.e., there was a same-natural pair consisting of two scenes, etc.).

2.2.3 Results

Learning for categories was not above chance by the end of block one ($t(29) = -0.545$, $p = .59$, Cohen's $d = -0.01$) but quickly reached above chance levels by block two ($t(29) = 7.29$, $p < .001$, Cohen's $d = 1.33$) and remained above chance for the remaining blocks (all $p < .001$). A 2 (same/different natural category) x 2 (same/different arbitrary category) repeated measures ANOVA revealed a main effect of natural category ($F(1, 29)=6.97$, $p = .013$, $\eta_p^2 = .19$), and a main effect of arbitrary category ($F(1, 29)=17.65$, $p < .001$, $\eta_p^2 = .38$). In both cases, same-pairs were learned better than different-pairs, but no significant interaction was observed $F(1, 29)=0.07$, $p = .8$, $\eta_p^2 = .002$). Four Bonferroni corrected t-tests against chance revealed non-significant learning for different-natural different-arbitrary pairs $t(29)=2.23$, $p = .033$, Cohen's $d = 0.41$) but significant learning in all other conditions (same-natural same-arbitrary, $t(29)=8.89$, $p < .001$, Cohen's $d = 1.62$, same-natural different-arbitrary, $t(29)=4.44$, $p < .001$, Cohen's $d = 0.81$, different-natural same-arbitrary, $t(29)=5.08$, $p < .001$, Cohen's $d = 0.93$.

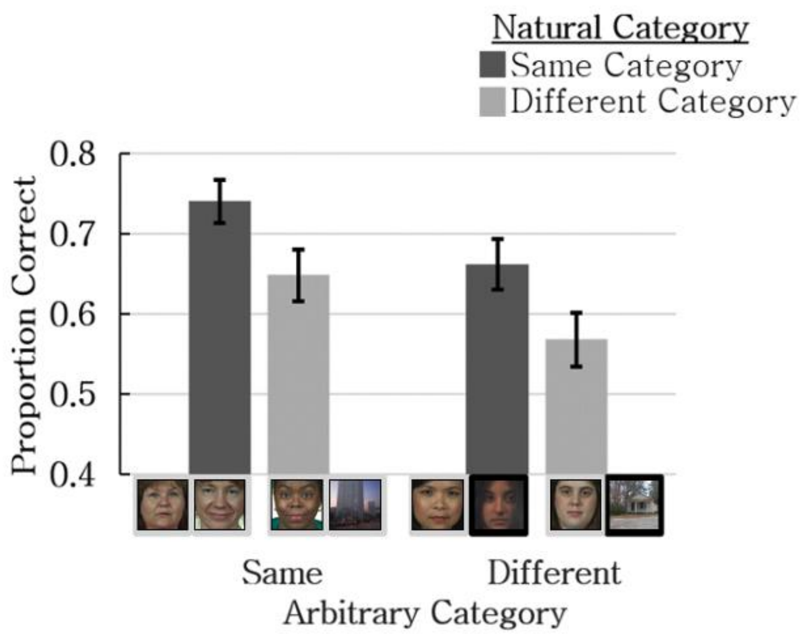

Figure 16. Preliminary data suggest greatest learning for pairs that shared the maximum amount of category group information. Pairs that only share one category group dimension show decreased learning, with the least amount of learning for pairs that share no category group membership.

2.2.4 Discussion

Similar to Experiment 1, visual statistical learning was observed overall in Experiment 2. The impact of arbitrary category was replicated with same arbitrary category pairs learned better than different arbitrary category pairs. Additionally, we observed an impact of natural category membership despite natural categories being wholly irrelevant to the task. Similar to the pattern observed for arbitrary category pairs, same natural category pairs were learned better than different natural category pairs. However, there was no interaction between arbitrary and natural category membership. This suggests the total information shared between two images may predict how well a visual statistical representation is learned, and that this information can be both novel (as in the case with the arbitrary categories) and preexisting (i.e., natural face/scene categories).

2.3 Experiment 3A

Experiment 1 and Experiment 2 both investigated the impact of category information shared between items of a statistical representation. However, we cannot say with certainty that arbitrary category information is comparable to natural category information in terms of being like a stimulus-level influence that is driven by perceptual similarly, at least in terms of its influence on visual statistical learning. Thus far, arbitrary categories have been learned during the presentation of statistical regularities. This unequivocally has tied arbitrary category information to task demands; participants were required to actively learn category information while, at the same time, incidentally learning statistical information. Given evidence from Experiment 2 that natural categories (which possess naturally occurring visual features that contribute to category discrimination, unlike the fractal images) are able to impact statistical learning without a task that requires participants to use the natural category information, it would be prudent to ensure this pattern is maintained with the use of artificial categories by temporally separating category learning from visual statistical learning. To examine this, we first trained participants to learn arbitrary category information in the absence of statistical regularity. Then, participants were exposed to statistical regularities in the context of a jitter-detection task. Figure 17 provides a visualization of this procedure.

2.3.1 Participants

A total of 30 University of Delaware Students participated in the study for partial completion of general course requirements.

2.3.2 Materials and Procedure

Similar to Experiment 1, stimuli consisted of 16 fractal images that were covertly placed into 8 pairs after the categorization phase was complete. Half of the images were randomly assigned to one category group (a “z” category which require a z response on the keyboard) and the remaining half were randomly assigned to the other category group (an “m” category requiring an m response in the keyboard). The

8 pairs of images consisted of 4 same-category pairs (e.g., both images within the pair belonged to the “z” category) and 4 different-category pairs (e.g., one image within the pair belonged to the “m” category while the other image belonged to the “z” category). Images were not organized into pairs during the initial categorization phase where participants learned the arbitrary category for each image. Each image within the categorization phase was presented onscreen for 1 second and participants were required to respond before the image disappeared. Feedback was provided in the form of a green fixation circle for correct responses that remains on screen for 1 second. If incorrect, participants were presented with a red fixation circle for 1.5 seconds. To equate Experiment 3 with Experiments 1 and 2 in terms of stimulus exposure, participants viewed each fractal image four times per block with three total blocks of training. The remaining three blocks of training were reserved for the familiarization phase where participants would be exposed to statistical regularities. This would match the amount of time participants encountered each stimulus to Experiment 1 while allowing participants to learn categories or statistical contingencies by using less stimuli overall (16 images as opposed to 32 images).

After this training phase, participants began the familiarization phase where all images were organized into pairs. Pairs were pseudo-randomized such no pair can immediately repeat or repeat with a single intervening pair. While participants viewed the images, they were required to respond to jitter events (a slight back-and-forth movement of an image) by pressing the spacebar, and the timing of stimulus presentation and feedback was consistent with the training phase (Figure 17). This jitter detection task was modeled after that which was used in Vickery and colleagues (2018). Participants again viewed each pair four times per block with three total blocks of training. After the familiarization phase, participants completed the same surprise test task identical to that used in Experiment 1.

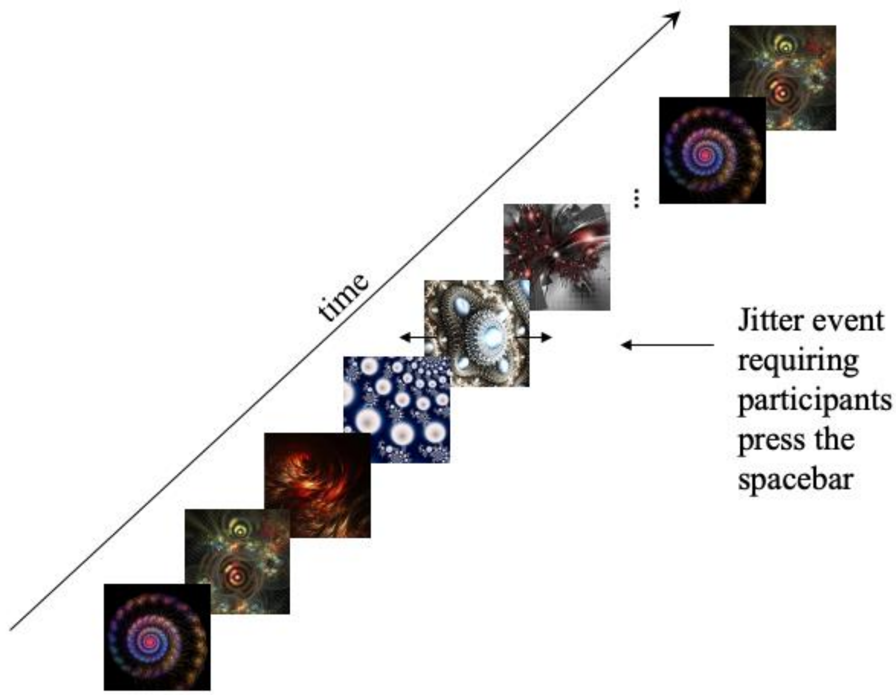

Figure 17. After category learning, participants were exposed to regularities within the context of a jitter detection task. Participants watched images go by and pressed spacebar whenever an image quickly moved left and right.

2.3.3 Results

Overall, learning for categories was significant by the end of the category-learning phase ($t(29)$ =7.32, $p < .001$, Cohen's d = 1.34) and participants had performed well when detecting jitter events (proportion of hits, M = 0.85, SD = 0.27). Evidence from the test phase, however, suggested that neither same-category pairs ($t(29)$=-0.152, p = .88, Cohen's d = -0.03) nor different-category pairs ($t(29)$ = -0.289, p = .78, Cohen's d = -0.05) were learned at above-chance levels (Figure 18). There was also no observed difference in learning of same-category pairs and different-category pairs ($t(29)$=0.127, p = .033, Cohen's d = 0.41).

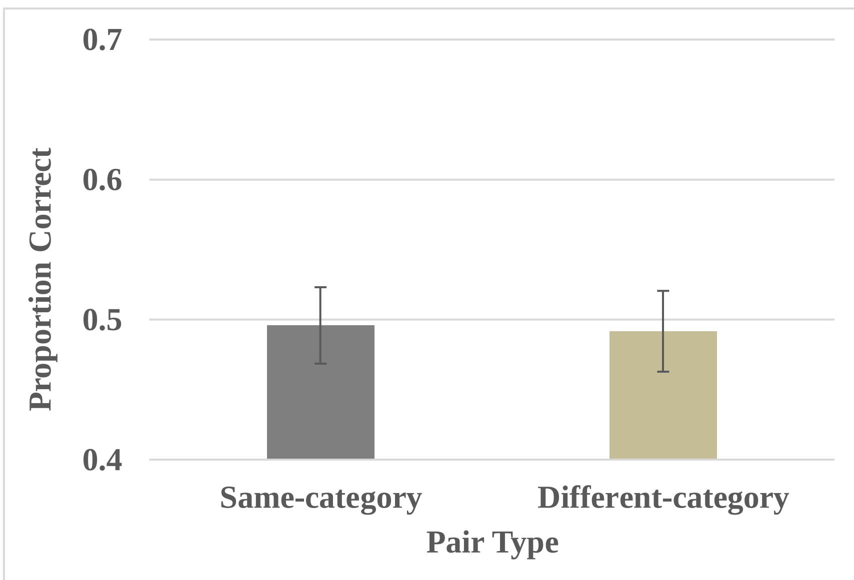

Figure 18. Test data from Experiment 3A.

2.3.4 Discussion

Overall, statistical learning was not evident in Experiment 3A. Despite the fact that participants viewed stimuli the same amount of times as in Experiment 1, and that we reduced the overall number of pairs from 16 down to 8 (because the 6 blocks of training were now split to allow for 3 blocks of category learning sequentially preceding 3 blocks of statistical learning), we failed to observe a pattern even remotely similar to Experiment 1. There are several possibilities worth considering about why this turned out to be the case.

First, and directly relating to the impetus driving Experiment 3, this could suggest that novel information (i.e., arbitrary category) does not impact visual statistical learning unless it is intrinsically task-relevant when visual statistical learning is occurring. However, this seems unlikely given the complete absence of evidence for visual statistical learning. Experiment 1 provided evidence that statistical learning can occur at the same time as explicit learning of category information, but the separation of explicit learning should not entirely eliminate visual statistical learning in Experiment 3. One plausible explanation would point toward our unique decision to use multiple tasks in a within-subjects design.

In Experiment 3, we used an arbitrary categorization task (that did not possess statistical regularities) followed by a jitter detection task (that did possessed statistical regularities). Our jitter detection task was modeled after Vickery and colleagues' (2018) detection task. In their fourth experiment, they examined the impact of task on visual statistical learning. As discussed previously, they used a categorization task and a jitter detection task, but their design was between subjects; one group from Vickery and colleague's experiment engaged in categorization while the other engaged in jitter detection. This is in stark contrast to Experiment 3, where participants first engaged in an arbitrary categorization task before being exposed to regularities within the jitter detection task. With the category learning task demanding that participants learn through trial and error (they begin the first block at chance accuracy due to the random assignment of arbitrary category), it is likely we had exhausted participants' willingness to consider the stimuli they were viewing in the experiment in the context of the jitter-detection task which is, by comparison to the arbitrary categorization task, effortless to perform well in (i.e., they only need to be set for jitter-related motion, they do not need to consider the content of each image). Given that one of the original statistical learning paradigms by Fiser and Aslin (2002) only required participants to passively view stimuli (and no task was present that required a response), it seems possible that our within-subjects design maybe have discouraged statistical learning by only introducing statistical regularities after participants had engaged in a cognitively challenging task. However, results may unexpectedly contribute to our understanding about inconsistencies in the visual statistical learning literature (Musz et al., 2015). We will speculate more on this in Chapter 5.

2.4 Experiment 3B

Given the unexpected failure to produce any evidence of visual statistical learning in Experiment 3A, Experiment 3B was included to investigate the original question posed in Experiment 3A while attempting to avoid the within-subject cross-task pitfalls we speculate to have driven the results of Experiment 3A. We have done this in two ways. First, we have moderately increased the overall amount of training participants receive. If we continue to see a lack of evidence of visual statistical learning, we will be able to say with greater certainty that it was not due to a lack of opportunity to extract the regularities presented within the experiment. Second, we have changed the jitter-detection task to a one-back task. We have previously speculated that part of the reason why visual statistical learning was not observed in Experiment 3A was due to participants' being influenced to consider each image as shallowly as possible (i.e., attending to motion cues only) after being exhausted by the categorization task (Himberger et al., 2019). Given the limited training and the fact that statistical regularities only came after the categorization task (i.e., we do not flip between blocks of categorization and blocks of regularity familiarization training), changing the jitter detection task to a 1-back task encourages participants to keep a previous stimulus in mind and consider how it matches the stimulus presented on any given trial. In other words, participants cannot merely be "set" for motion to more easily complete the task.

If, with the present set of alterations to our design, we continue to observe a complete lack of evidence for visual statistical learning, it may suggest that visual statistical learning is a form of incidental learning that may not occur as efficiently as it would if not immediately following an explicit learning task. In other words, if visual statistical learning is combined with explicit learning like in Experiments 1 and

2, or if participants have an opportunity to switch back and forth between blocks involving some explicit learning task (which has yet to be explored), visual statistical learning may persist, but a single instance immediately following explicit learning may temporarily inhibit this particular form of incidental learning. We considered this scenario to be unlikely, but possible.

A more likely possibility is that the use of a demanding categorization task preceding a relatively passive jitter-detection task influences participants to engage in the experiment in a way that is different from a scenario where they only receive the jitter-detection task. Again, previous evidence of visual statistical learning was observed using only a jitter-detection task(Vickery et al., 2018), or even no task at all other than to passively watch the images go by (Fiser & Aslin, 2002). Changing the task used to familiarize participants with statistical regularities to one that encourages consideration of the features within the image should preserve an opportunity for visual statistical learning to occur.

2.4.1 Participants

A total of 81 online Mturk participants completed the study. A priori cutoffs were established such that participants who failed to respond correctly to a minimum of 70% of the categorization trails in the final block and/or respond correctly to a minimum of 70% of n-back events were excluded from analyses. This left us with a total of 54 participants.

2.4.2 Materials and Procedure

Similar to Experiment 3A, stimuli consisted of 16 fractal images that were covertly placed into 8 pairs after the categorization phase was complete. Half of the images were randomly assigned to one category group (a "z" category which require a z response on the keyboard) and the remaining half were randomly assigned to the other category group (an "m" category requiring an m response in the keyboard). The 8 pairs of images consisted of 4 same-category pairs (e.g., both images within the pair belonged to the "z" category) and 4 different-category pairs (e.g., one image within the pair belonged to the "m" category while the other image belonged to the "z" category). Images were not organized into pairs during the initial categorization phase where participants learned the arbitrary category for each image. Each image within the categorization phase was presented onscreen for 1.5 seconds and participants were required to respond before the end of 2 seconds after stimulus onset. Responses immediately terminated the trial. Feedback was provided in the form of a green "Correct!" for correct responses that remained on screen for 0.5 seconds. If incorrect, participants were presented with a red "Incorrect, the answer was [correct answer appeared]" for 1.5 seconds. Non-responses were also given feedback in the form of a message informing them that their response was too slow, with a 1.5 second delay. Each image appeared 4 times per block, and participants completed 5 blocks of categorization training before moving on to the 1-back familiarization task.

When participants began the familiarization phase, all images were organized into pairs. Pairs were pseudo-randomized such that no pair could immediately repeat

(ABAB, for pair AB) nor could two back-to-back pairs repeat in sequence (ABCDABCD, for pairs AB and CD). While participants viewed the images, they were required to respond to the presence or absence of one-back events by pressing the “n” key on the keyboard (indicating a non-repeat) or by pressing the “m” key on the keyboard (indicating a repeat/one-back event), and the timing of stimulus presentation and feedback was consistent with the categorization phase (except that responses did not terminate trials, and correct responses were not given feedback – only false alarms and misses were given feedback). Each image appeared 4 times per block (not counting 1-back events), and participants completed 5 blocks of the 1-back familiarization phase before beginning the test task, which was identical to that used in Experiment 1. One-back events only occurred for the 2nd item in each pair, and occurred for each such item once in each familiarization block.

2.4.3 Results

Participants had reliably learned the categories for each stimulus by the final block ($t(53) = 35.6$, $p < .001$, Cohen’s $d = 4.84$). During the one-back training, participants reliably detected the majority of one-back events (for proportion of hits, $M = 0.87$ $SD = 0.07$) while making relatively few false alarms (for proportion of false alarms, $M = 0.01$, $SD = 0.01$). Same-category pairs ($t(53) = 4.29$, $p < .001$, Cohen’s $d = 0.583$) were learned at above-chance levels but different-category pairs were not ($t(53) = 1.32$, $p = .191$, Cohen’s $d = 0.18$). There was also a significant difference in learning of same-category pairs and different-category pairs ($t(53) = 2.07$, $p = .043$, Cohen's $d = 0.282$, Figure 19).

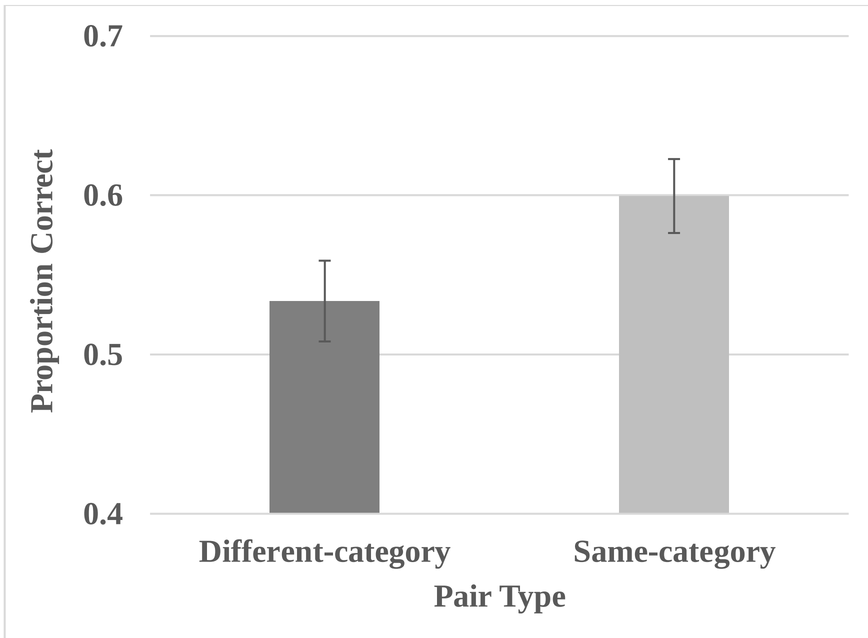

Figure 19. Test data from Experiment 3B.

2.4.4 Discussion

With the additional training blocks added to the categorization phase and the familiarization phase, and with the task switched to a one-back task from a jitter-detection task, we were able to gather evidence that both observes statistical learning and supports the findings from Experiments 1 and 2. Under these conditions where novel category learning is separated from visual statistical learning (it specifically precedes it), a difference in learning based on the newly learned categories was evident. Supporting our initial findings, same-category pairs were learned better than different-category pairs. Interestingly, only same-category pairs were learned at above-chance levels. Whether this is due to some of our earlier concerns about the familiarization phase following a challenging trial-and-error category learning phase, or due to our now successful attempt at separating category learning from visual statistical learning while replicating Experiments 1 and 2, remains to be seen. What we can conclude, however, is that we are still able to view evidence of category membership impacting visual statistical learning by temporally separating novel category learning from visual statistical learning. Although task demands clearly have an impact on what is visually statistically learned, the present findings suggest that it is the category information, rather than some combination of category information the task demands imposed during category learning, that may impact visual statistical learning.

Chapter 3

PERCEPTUAL SIMILARITY AND VISUAL STATISTICAL LEARNING

3.1 Experiment 4

Experiments 1 through 3 were designed to scrutinize the role of stimulus similarity and examine how pre-existing category group information or task may influence visual statistical learning. This left open an important question: how do perceptually similar or dissimilar stimuli dictate what is visually statistically learned? Stimuli that are perceptually similar, albeit incidentally learned, may be bound together more easily during visual statistical learning. After all, perceptual similarity can serve as a driving force in the learning of natural category groupings, even from early ages (Bomba & Siqueland, 1983; Quinn, 1987; Strauss, 1979; Younger & Gotlieb, 1988). In other words, perhaps in the absence of semantic or naturally occurring category group information, perceptually similar pairs will be learned better than perceptually dissimilar pairs in a visual statistical learning task.

Experiment 4 used a categorization task that was additionally driven by carefully controlled perceptual differences in basic shapes devoid of the high-level information faces and scenes possess. In other words, participants learned perceptual category groupings that were designed to be specifically driven by perceptual similarity (using shapes designed to be visually continuous, see Figure 20), and eventually engaged in a statistical learning task while categorizing the stimuli into their respective perceptual groups. With a lack of any immediately clear visual boundary between shapes (in contrast to the obvious differences between a face and a scene), participants effortfully had to learn perceptual category groups. As the perceptual category boundary arbitrarily and randomly splits stimuli into two groups, participants were actively engaged in learning perceptual category groupings that obfuscated the underlying statistical structure of the images.

We hypothesized the greatest learning for same-category perceptually similar pairs with a stepwise decrease for every pair that does not share either categorical information, perceptual similarity, or both. In other words, same-category perceptually similar pairs may be learned best, different-category perceptually dissimilar pairs would be learned the least, and same-category perceptually different/different category perceptually similar pairs would fall somewhere in-between.

3.1.1 Participants

A total of 30 University of Delaware Students participated in the study for partial completion of general course requirements.

3.1.2 Materials and Procedure

Stimuli consisted of 360 basic shape images that lie on a continuum (Li et al., 2020). The visual continuum of stimuli was randomly be divided in half (Figure 20), which designated the two category groupings. Participants first engaged in a category group training task that does not possess any regularity information. A single stimulus appeared on screen and participants had three seconds to press one of two buttons to indicate which category group the stimulus belongs in. No time penalty was imposed on participants for incorrect responses like in Experiments 1 through 3. Every 50 trials, participants were informed of their average accuracy for the last 50 trials. They continued to engage in this training task until 10 minutes have passed or until they reach 90% accuracy within the previous 50 trials.

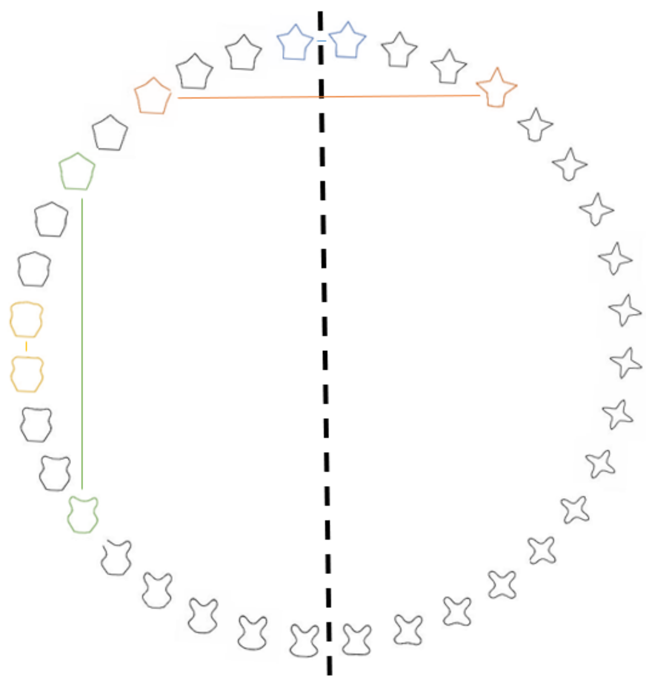

Figure 20. Basic shape stimuli continuum used in Experiment 4. The dotted line represents the category group boundary. Perceptually similar pairs within (highlighted in yellow) and between categories (blue) were matched in terms of similarity, as were perceptually dissimilar pairs within (green) and between (orange) categories. All shapes appeared in black for participants.

After the initial category group training, participants engaged in the same task at a faster pace (they had up to 2 seconds to respond to the stimulus, down from 3 seconds) and with all images covertly presented in pairs. A total of 16 images were used: 8 were from one category and 8 were from the other category. Eight pairs were created constituting four different pair conditions: perceptually similar same category pairs, perceptually dissimilar same category pairs, perceptually similar different category pairs, and perceptually dissimilar different category pairs. Each stimulus appeared on screen for one second with an intervening blank inter-trial interval that lasted for one second. Altogether, participants had two seconds to respond accurately before the next image appeared. If participants were inaccurate, they were presented with an additional screen for one second that informed them they were incorrect/too slow in responding. Each block of this training included 96 trials with a total of 6 blocks. After the second training task was complete, participants engaged in a two alternative forced choice task similar to those used in Experiments 1 through 3.

3.1.3 Results

Category learning was evident within the initial category group training ($t(29) = 27.5$, $p < .001$, Cohen's $d = 5.01$) and persisted throughout the second phase of category training that included statistical regularities ($t(29) = 51.7$, $p < .001$, Cohen's $d = 9.44$). For the test phase, a 2 (same/different arbitrary category) x 2 (perceptually similar/dissimilar) repeated measures ANOVA revealed a main effect of arbitrary category ($F(1, 29) = 5.76$, $p = .023$, $\eta_p^2 = .17$), a main effect of perceptual similarity

(F(1, 29) = 6.1, $p < .02$, $\eta_p^2 = .17$), with a interaction observed between perceptual similarity and arbitrary category (F(1, 29) = 5.97, $p = .02$, $\eta_p^2 = .171$). Four Bonferroni corrected t-tests against chance revealed significant learning across all four conditions (same-category similar pairs, $t(29) = 3.03$, $p = .005$, Cohen's $d = 0.55$, same-category dissimilar pairs, $t(29) = 3.13$, $p = .004$, Cohen's $d = 0.57$, different-category similar pairs, $t(29) = 8.14$, $p < .001$, Cohen's $d = 1.49$, different-category dissimilar pairs, $t(29) = 4.69$, $p < .001$, Cohen's $d = 0.86$).

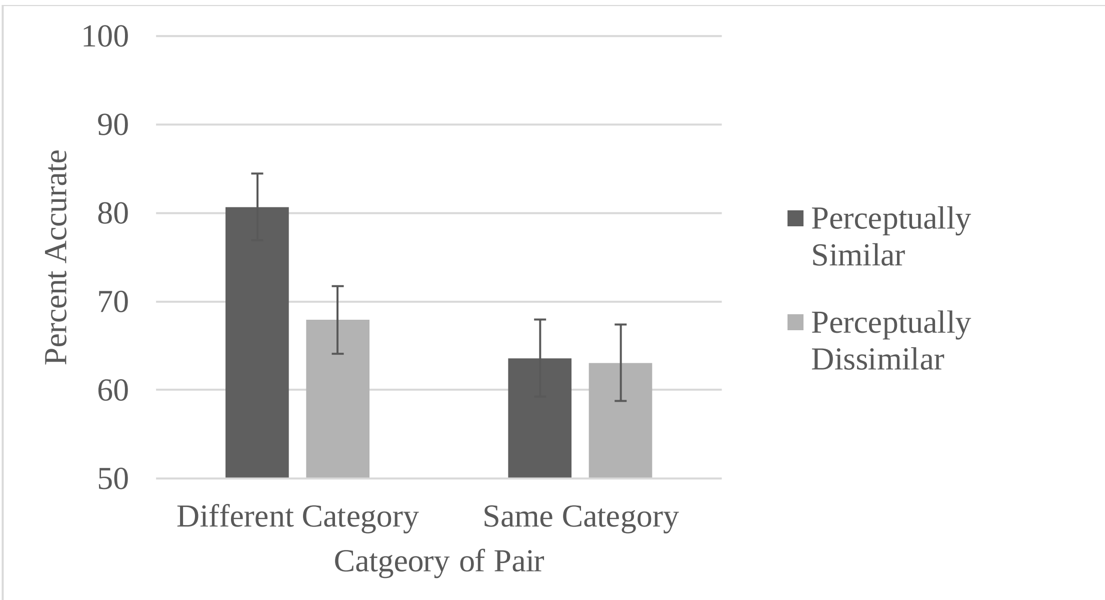

Figure 21. Test data from Experiment 4

3.1.4 Discussion

While visual statistical learning was evident overall as well as within each condition, results from Experiment 4 suggest that perceptually similar pairs of items from different categories are learned better than all other conditions. Why would this pattern emerge, rather than the pattern we hypothesized where a benefit for same-category/perceptually similar pairs would be observed (just as same natural/arbitrary categories contributed to visual statistical learning in Chapter 2)? One potential explanation would again call into consideration the impact of the task demands driving participant behavior during learning.

In Experiment 3, we argued that statistical learning may only occur when participants are effortfully engaging in the task. Thus, among other explanations, we primarily argue that we failed to observe statistical learning when statistical regularities were introduced after a cognitively demanding categorization task was completed and a simple jitter-detection task was introduced. Although the results from Experiment 4 were unexpected, they may serve to support this explanation. In other words, perhaps the pattern of data observed here are best characterized as resulting from a design that required differing amounts of engagement from participants between conditions (i.e., some trial conditions were easier for participants than others).

Specifically, as participants were learning the category for each stimulus, the most challenging stimuli to correctly respond to would naturally appear at the category boundary. Based upon participants' experience with the categories at any given time throughout the experiment, participants would have to more carefully scrutinize stimuli at the category boundaries to determine which category it belonged to. Additionally, for those pairs that were perceptually similar but crossed category boundaries, participants would have to carefully consider which category each stimulus best fits to prevent mis-categorization. Thus, this increased engagement with stimuli appearing at the category boundaries, or pairs of stimuli that crossed the category boundaries, may have driven the observed higher rates of learning for different-category pairs (as the category boundary was responsible for the various levels of difficulty between trials and conditions). See Figure 22 for an example. This account holds additional credence when considering prior work that has shown how discriminability for stimuli between categories has been observed to be better compared to discriminability for stimuli within categories (Etcoff & Magee, 1992; Liberman et al., 1957), even if categories are learned during the course of an experiment (Goldstone, 1994; Goldstone et al., 2001).

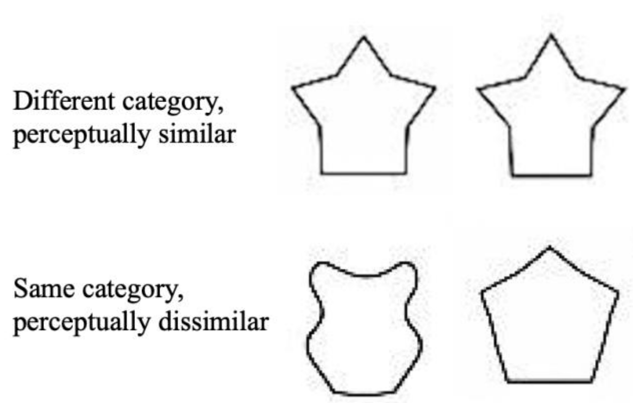

Figure 22. Side-by-side example of how the category boundary dictates a need for increased discriminability for pairs that cross category boundaries.

3.2 Experiment 5

Given categorization has a clear impact on perceptual discrimination ability, Experiment 5 is designed to explore the inverse question of Experiment 4: how might visual statistical learning shape perceptual similarity judgments? Using the similarity space from Experiment 4, we first created a series of pairs of stimuli that are matched in terms of visual similarity. Participants were then exposed to all of the stimuli used to create the pairs, but only half of the stimuli appeared in their respective pair while participants were exposed to a one-back task. The remaining stimuli did not appear as pairs but were rather shuffled into the trial sequence as singletons. This provided us with our two critical conditions that were matched on their similarity but differed in their capacity to elicit visual statistical learning: grouped (structured) pairs of stimuli can be visually statistically learned, and ungrouped (unstructured) pairs of stimuli that cannot be visually statistically learned

Following the familiarization phase, we then asked participants to rate pairs of images on how similar or dissimilar they are using a 9-point scale. Importantly, the unstructured singletons were returned into their paired form so that participants could provide ratings for both grouped pairs and ungrouped pairs. If statistical associations influence the perceptual space, we may expect pairs learned during familiarization to appear perceptually more different compared to pairs that do not possess any statistical associations.

3.2.1 Participants

A total of 30 University of Delaware Students participated in the study for partial completion of general course requirements.

3.2.2 Materials and Procedure

The materials used in for Experiment 5 are the same as that used in Experiment 4, while the familiarization phase used a one-back detection task (Figure 23). Using the similarity space from Experiment 4, we created 8 pairs of stimuli matched by similarity between pairs. During familiarization, half of the images were shown grouped within their respective pairs. Figure 24 provides a visualization of pairings drawn from the similarity space. All images appeared 4 times per block and participants completed 4 blocks of familiarization.

Following familiarization, participants began making similarity judgements on pairs of stimuli using a 9-point scale. Within this similarity judgment task, participants made similarity judgments on all 8 pairs of stimuli. Images of a pair appeared one after another with stimulus presentation timing matched to the familiarization phase (one second on, one second off). After each image is shown, participants entered their response on a scale from 1 (very similar) to 9 (not similar at all).

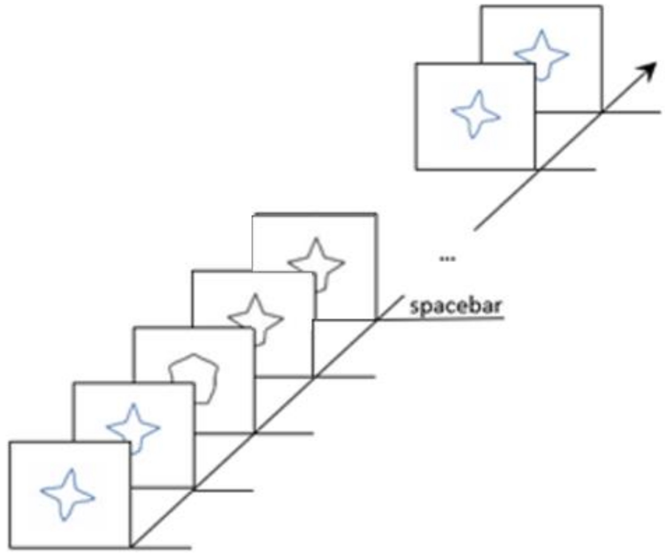

Figure 23. Example trial sequence for Experiment 5. Blue shapes represent grouped stimuli (pairs) that appeared together during familiarization. Participants pressed spacebar whenever a one-back event occurred.

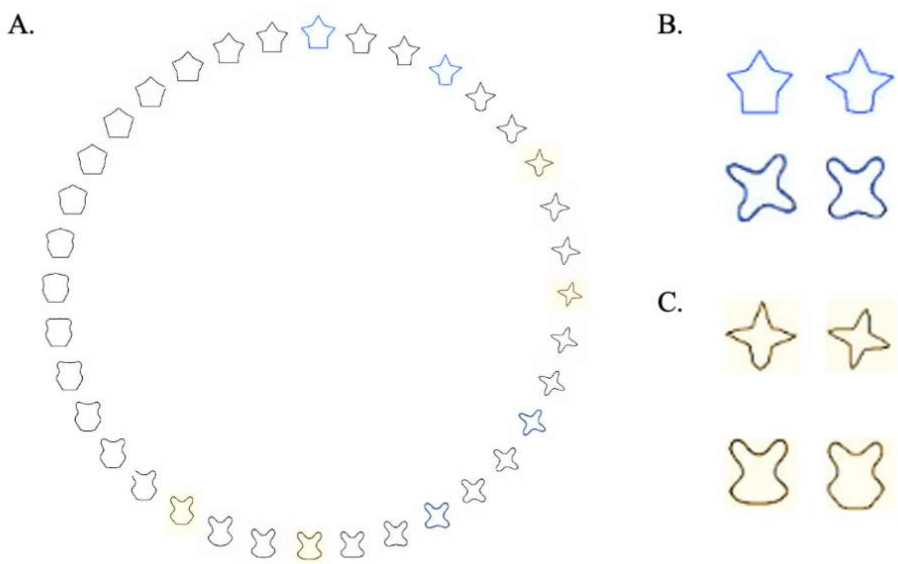

Figure 24. A. Example pairings drawn from the similarity space. B. Grouped (paired) stimuli that appeared during familiarization. C. Ungrouped stimuli that only appeared as pairs during the test phase. Stimuli did not appear highlighted.

3.2.3 Results

Although we did not set a priori cutoff criteria for Experiment 5, participants reliably detected the majority of 1-back events (for proportion of hits, $M = 0.66$ $SD = 0.17$) while making relatively few false alarms (for proportion of false alarms, $M = 11.12$, $SD = 0.07$). A comparison of similarity ratings for grouped (i.e., visually statistically learned) vs ungrouped dissimilar pairs yielded no observable differences ($t(29) = 0.505$, $p = .62$, Cohen's $d = 0.09$).

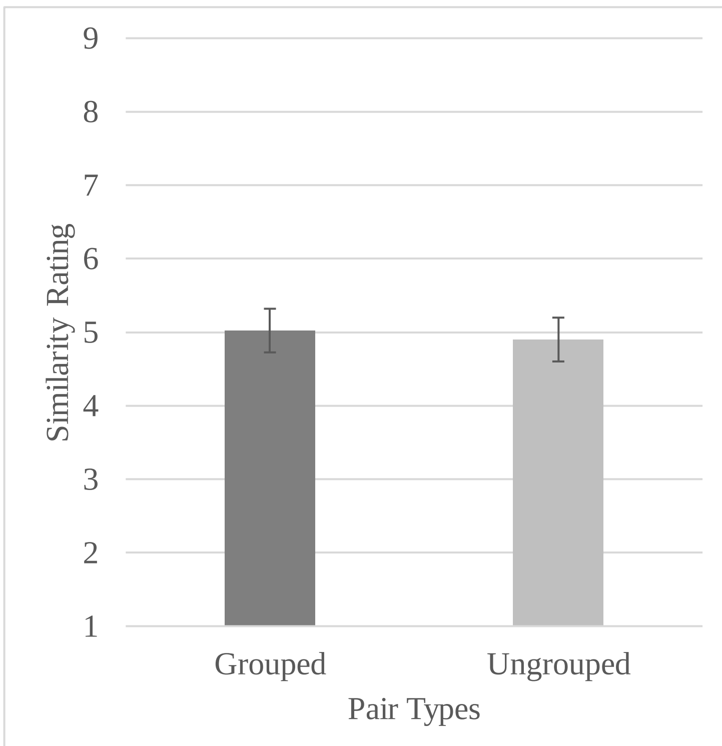

Figure 25. Similarity rating data for Experiment 5.

3.2.4 Discussion

Visual statistical learning does not appear to influence similarity ratings. Although we were unable to test for visual statistical learning because the similarity rating task explicitly revealed the presence of pairs, it is unlikely that pair learning failed to occur since the 1-back task typically elicits robust pair learning. Nevertheless, it is possible that pair learning failed to occur, thus a follow-up could address this by including more exposure and a pair testing stage in leu of a similarity judgement task for a matched group.

Experiment 4 and Experiment 5 attempted to scrutinize ways in which visual statistical learning and perceptual similarity may interact by using basic shape stimuli. It may be worthwhile for future work to consider the use of more complex stimuli, not unlike those used in Experiments 1 through 3. If visual statistical learning does influence perceptual similarity judgements and Experiment 5 was unable to tap into this effect, it may have been because the simple shapes we used do not offer enough featural variety for visual statistical learning to interact with, or for participants to make reasonable similarity judgements on (i.e., any two of the most different stimuli from our set are simple shapes constructed from one continuous line).

Chapter 4

NEURAL SIGNATURES OF VISUAL STATISTICAL LEARNING AS SHAPED BY TASK

4.1 Experiment 6

Experiments 1 and 2 illustrate that category training has a profound influence on patterns of visual statistical learning. Prior work (Vickery et al., 2019) also suggests that different tasks, such as simple detection tasks, do not lead to such patterns. These observations raise the question: does attention to categories evoke similar or distinct mechanisms of statistical learning? While it is difficult to access such distinctions, neuroimaging can help by suggesting whether different or similar neural consequences of learning in different contexts occurs. Experiment 6 is designed to examine neural correlates of visual statistical learning within subjects and between learning contexts. All participants completed the same experiment (except counterbalancing of the order of conditions between subjects) and were exposed to statistical regularities within two different learning contexts. Like Experiment 3, we examined the impact of the learning context in a separate, non-learning context. After learning is complete, participants were scanned using fMRI while engaged in a task that is not like either of the two previous tasks (they passively viewed images from the learning contexts) and were not actively learning new information.

A number of studies have used fMRI to explore differences in neural correlates relating to statistical learning, and while some evidence seems to replicate across studies, there are certainly some inconsistencies between studies that adopt different protocols (Turk-Browne et al., 2009, 2010), perhaps due to variations in the training paradigm. One clear path to clarifying these inconsistencies is to investigate the neural correlates of how different learning contexts influence visual statistical learning. Prior work has highlighted several cases where differing task demands lead to differences in visual statistical learning (Vickery et al., 2018; Turk-Browne et al 2005), as well as the evidence from the previous chapters. Using fMRI, Experiment 6 is designed to probe differences in incidental learning driven by different task demands. Using a within-subjects design, participants engaged in a visual statistical learning paradigm by going through alternating blocks of two different tasks, providing two different learning contexts.

In this experiment, all participants engaged in both a categorization task (Turk-Browne et al., 2010; Vickery et al., 2018) and a 1-back task (Turk-Browne et al., 2005) that were associated with both paired images and singleton images, unique to task context. Both tasks were exactly the same (i.e., repeat images also occurred during categorization), the only difference between the tasks are the instructions participants receive and how they respond to each image. Additionally, all images consisted of fractal images in order to control for prior natural category knowledge (i.e., instead of using faces and/or scenes). Participants were exposed to singletons and pairs in these two different contexts for approximately 1.5 hours before entering the scanner. In order to examine effects arising from two distinct on-line task demands, all previously learned pairs and singletons were intermixed into a single event-related stream once participants enter the scanner. Participants were told to carefully focus in on and watch each image go by slowly, allowing for a neutral context unrelated to either of the learning contexts.

If visual statistical learning shares a common neural underpinning across tasks, we may expect to see greater hippocampal activity for paired images learned during either task when compared to activity elicited by singleton images. If not, or additionally, we may expect to see some differences driven by the task. It may be possible for some brain regions associated with statistical learning (e.g., lateral occipital cortex or medial temporal lobe) to differentiate first, second, and singleton images learned during categorization, even when prior knowledge is controlled for (i.e., participants do not have experience with fractal images like they have experience with faces or scenes). Critically, our design focuses on the neural correlates of statistical learning *after* they have been learned (the reviewed works measured brain responses while task and learning was ongoing), and outside of their learned context (during a neutral passive viewing task).

4.1.1 Participants

20 participants were paid $20/hour for their time spent inside the scanner and $10/hour for their time spent outside of the scanner. Scanning typically lasted 1.5 hours and non-scanning time typically lasted 2 hours.

4.1.2 Materials and Procedure

A total of 96 fractal images were used. These fractals were split between two tasks: the categorization task and one-back task (Figure 26). For the categorization task, 32 of the 48 images were assigned into pairs consisting of 8 same-category pairs and 8 different-category pairs. The remaining 16 images served as singletons that were not paired with another image. For the 1-back task, 32 of the 48 images were assigned to pairs irrespective of category (because category was not present nor relevant in the one-back task). As with the categorization task, the remaining 16 images also served as singletons that were not paired with another image.

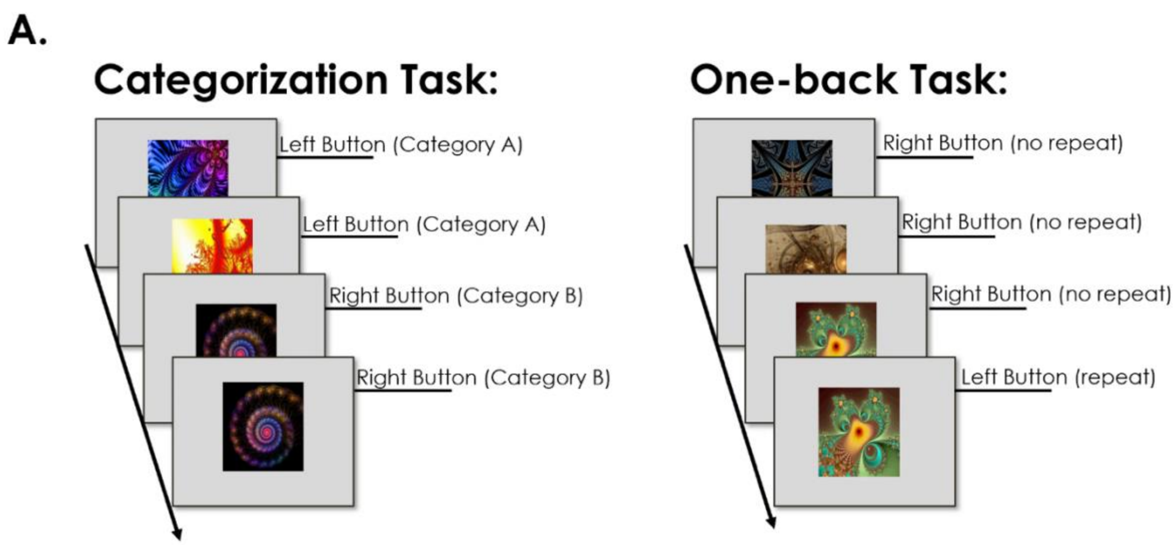

Figure 26. Example sequence of stimuli presentation and required response for the categorization task (left) and the one-back task (right).

Participants completed a total of 6 runs of familiarization where they viewed all of the images a total of 18 times across both tasks. Each run consisted of interwoven blocks of each task such that each task flipped an equal number of times between runs (e.g., ABABAB). Whether the participant began with the categorization task (A) or the one-back task (B) was counterbalanced (i.e., half of participants experienced runs of BABABA). Each run contained a single repetition of all images assigned to that task.

The categorization task was similar to that of Experiment 1; participants learned through trial and error which category each image belongs to. In order to equate responses between tasks, participants used the same keys to make responses in the one-back task. If the image was not a repeated image, they pressed one key. If the image was an immediately repeated image, they pressed the other key. Repeated images appeared in both tasks, such that each of the 96 images repeated once per run, divided amongst the 6 blocks.

During scanning, participants passively viewed all 96 images one at a time in the scanner throughout each run, with a total of 4 runs. Pairs were maintained, but all pairs and singletons from both tasks were intermixed in a single run. Each image was displayed for 1 second with a variable 4-7 second inter-trial-interval consisting of a black screen and a small fixation circle. Participants were instructed to lay as still as possible and fixate on each image as it passed by, and an MRI compatible eye tracker was used to ensure participants do so. After scanning, participants were led to another testing room where they were given a two alternative forced choice recognition test, the same used in Experiments 1 and 2.

4.1.3 fMRI Data Acquisition

Data were collected using a 3T Siemens Prisma system using a 64-channel head/neck coil. High resolution (0.7 mm isometric voxel) structural data used for registration was collected as a T1-weighted MPRAGE structural image. Functional scans were collected using T2 weighted Siemens Multiband sequences that acquired 80 interleaved slices at an oblique axial orientation at 25° from the anterior commissure/posterior commissure line with a resolution of 2.0 mm x 2.0 mm x 2.0 mm (TR= 1s, TE = 32 ms, flip angle 61°). Each run collected 628 volumes and lasted approximately 10 minutes.

4.1.4 Behavioral Results

During training, participants learned the categories for each stimulus by the final block($t(19) = 38.7$, $p < .001$, Cohen's $d = 5.26$) and correctly detected the majority of one-back events (for proportion of hits, $M = 0.92$ $SD = 0.08$) while making relatively few false alarms (for proportion of false alarms, $M = 0.01$, $SD = 0.01$). During the test phase, pairs learned during categorization were recognized significantly above chance ($t(19) = 5.35$, $p < .001$, Cohen's $d = 1.2$), but pairs learned during the one-back task were not ($t(29) = 1.29$, $p = .289$, Cohen's $d = 0.29$). Overall, pairs learned during categorization were learned better than pairs learned during the one-back task ($t(19) = 2.43$, $p = .025$, Cohen's $d = 0.542$, Figure 27). Pair learning under categorization did not differ by function of pair combination ($t(19) = -0.06$, $p = .954$, Cohen's $d = -0.01$), with learning for same-category pairs not differing from learning for different-category pairs (Figure 28).

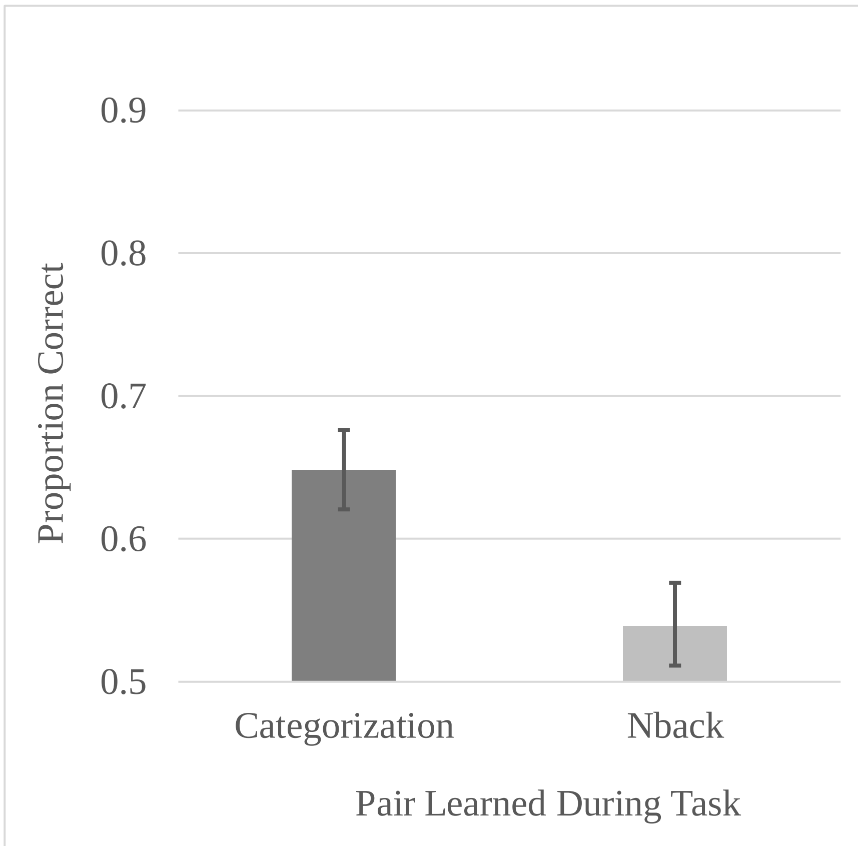

Figure 27. Experiment 6 test phase results obtained after scanning.

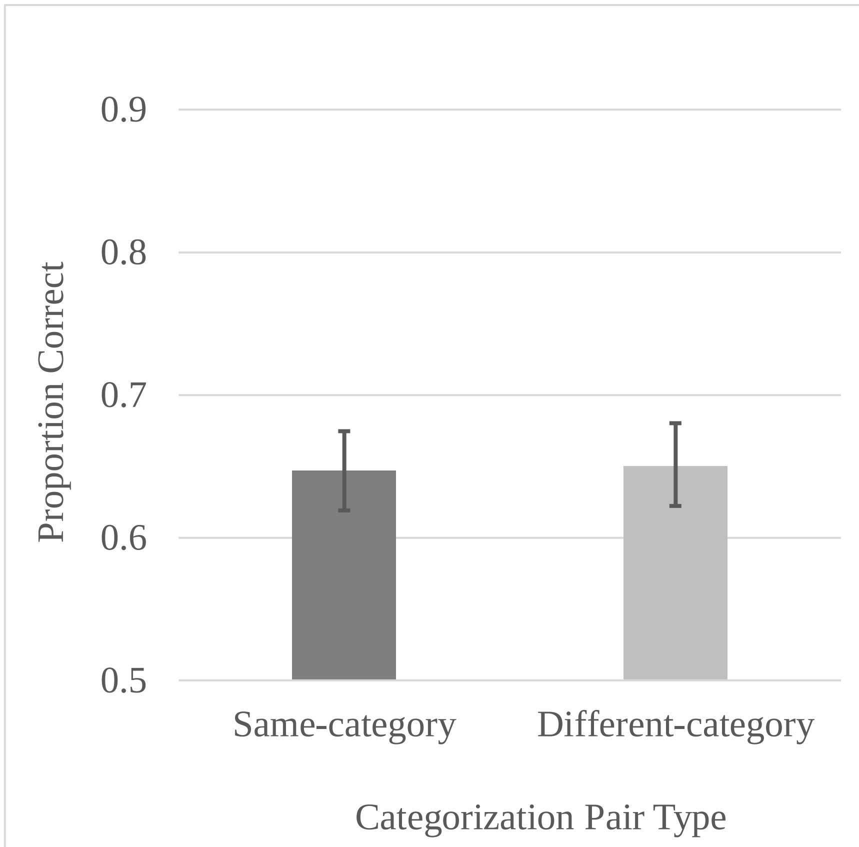

Figure 28. Experiment 6 test phase results divided by category membership.

4.1.5 Neuroimaging Results

Data analyses was conducted using fMRIB Software Library (Jenkinson et al., 2012) version 5.0.9, fMRI Expert Analysis Tool (FEAT) version 6.0 (get citation), and the AFNI software package (Cox, 1996). The high-resolution structural scan was skull stripped using BET (Smith, 2002) before being registered to a standard MNI152 2mm template using a nonlinear 12 degrees of freedom normal search. Functional runs were first be de-obliqued using ANFI's 3dWarp and then re-oriented to match the standard template using FSL's fslreorient2std. MCFLIRT motion correction was applied along with a 5mm FWHM Gaussian kernel for smoothing. A high pass temporal filter was applied to remove low frequency artifacts.

For the first-level analysis, all runs from all participants were modeled using a standard GLM approach. A total of 10 explanatory variables (EVs) was used. For the following explanation of the EV's used, the abbreviations stand as follows: One = first image of a pair, Two = second image of a pair, Cat = image learned during the categorization task, Nback = image learned during the 1-back task, Same = same category as it's respective paired image, Dif = different category as it's respective paired image, Sing = singleton. The 10 EVs used are as follows: One_Cat_Same, One_Cat_Dif, One_Nback_Same, One_Nback_Dif, Two_Cat_Same, Two_Cat_Dif, Two_Nback_Same, Two_Nback_Dif, Sing_Cat, Sing_Nback. All EVs were convolved using a double-gamma standard hemodynamic response function. At the first level, several contrasts were computed, and these contrasts were passed forward to higher level analyses. For the second level analysis, a fixed effects analysis was used to combine across the 4 passive viewing runs for each participant. The third level analysis combined runs across participants using a mixed effects analysis. The whole-brain statistical maps resulting from this analysis were cluster corrected for multiple comparisons using FSL's Randomise (Jenkinson et al., 2012), a nonparametric testing tool, using 5000 permutations and threshold free cluster enhancement (TFCE).

As stated previously, we were interested in understanding if attention to categories evoke similar or distinct mechanisms of statistical learning. To probe for this, we compared structured items to singletons for both tasks. This includeded One_Cat > Sing_Cat and Two_Cat > Sing_Cat for categorization and One_Nback > Sing_Nback and Two_Nback > Sing_Nback for the 1-back task. If we had observed similar patterns of activity in these contrasts, it may have indicated general hippocampal activity for structured items (Turk-Browne et al., 2009). Unfortunately, none of our comparisons yielded significant results from these comparisons.

Alternatively, within these contrasts, we speculated about how we may approach findings that suggests distinct mechanisms of statistical learning between tasks. For instance, a distinct pattern of activity between tasks could have manifested in the form of increased activity within medial temporal lobe and/or striatum for images learned during categorization compared to images learned during the 1-back task. This pattern may suggest that information shared between items of a pair from the categorization task engendered activity from areas relating to associative learning in a way that items learned in the 1-back task did not (i.e., a possible link to the behavioral effects observed from Chapter 2 of the proposal). In our design, this explanation was explorable with additional contrasts scrutinizing activity arising from same-category pairs or different-category pairs. These contrasts included One_Cat_Same > One_Cat_Different and Two_Cat_Same > Two_Cat_Different. In the case that increased medial temporal lobe/striatal activity was observed for same-category images compared to different-category images, it may have suggested that these associative learning areas predict the strength of statistical learning as a function of the sum total information shared between items of a pair (again, relating back to an interpretation of behavioral evidence from Chapter 2). However, the additional contrasts we included here also did not yield any significant differences.

4.1.6 Discussion

Behavioral results suggest that visual statistical learning was evident, but only for the categorization task. Although blocks of the categorization task and blocks of the one-back task were intermixed within and between runs, this pattern of results may still be due to the use of multiple tasks in a within-subjects design. In the context of the present experiment, participants were exposed to one task where they could only perform at chance levels before slowly learning categories over time, and a one-back task where performance easily reached ceiling within the first few trials. In fact, many participants expressed an unprompted verbal preference for the 1-back task (occasionally referring to it as the easy task).

In order to actively learn the category membership of each image in the categorization task, participants were required to engage with stimuli in a way that encouraged studying the visual details of the image while trying to remember the category it belonged to. In contrast with the one-back task, no such effort was required; participants were only required to look for whether each image on screen was an immediate repetition of the previous image or not. We made a similar argument between Experiment 3A and 3B, although the major difference in

Experiment 6 is that visual statistical regularities were presented within both the categorization and the one-back task. While these visual statistical learning tasks were both used in the same within-subjects design, this may explain why statistical learning was evident for pairs learned during the categorization task but not pairs learned during the one-back task (and why participants spontaneously referred to them as the "hard task" and "easy task", respectively). Along with the influence of task on statistical learning (Vickery et al., 2018) this supports the idea that visual statistical learning is not an entirely passive process (Baker et al., 2004).

Additionally, within the category-learning task, there were no observable differences between learning for same-category pairs and different-category pairs. Unlike Experiment 1and 2 , the present experiment involved switching between tasks multiple times within a run, used 96 total images as opposed to the 32 images used in Experiment 1, and included the use of unpaired singleton images. In addition, four runs of the passive viewing task intervened between categorization exposure and the pair testing phase, which may have muted differences.

In sum, neuroimaging results from Experiment 6 were not as revealing as we had hoped, at least beyond what we have learned about the methodological considerations needed to compare tasks in a within-subjects design. Despite extensive familiarization with visual statistical regularities, and despite successfully observing some behavioral evidence of visual statistical learning (from the category learning task), our experiment may have been underpowered at the stage of the passive viewing task. There are several methodological decisions made that could have contributed to this scenario (e.g., the use of singletons, the passive viewing that shuffled all pairs and singletons together, and an ambitiously large stimulus set).

Chapter 5

SUMMARY AND CONCLUSION

Incidental learning is reliably elicited in laboratory settings, even in very short-term training episodes, implying that it is a ubiquitous phenomenon. How statistical learning – specifically *visual* statistical learning - operates in real-world scenarios remains open to scientific study. Some have argued that visual statistical learning provides an automatic and incidental means of constructing representations (Schapiro et al., 2012) , or even representations wherein statistical regularities are processed the same way as objects (Lengyel et al., 2021). However, most prior work as used unfamiliar stimuli about which nothing else is learnable except statistical regularity in the context of the experiment. Our work aimed to explore key features that we believe are important to how statistical learning might apply in naturalistic settings: new learning and prior knowledge for categories, perceptual similarities shared between stimuli, and the impact of task demands.

5.1 Category Information Predicts Visual Statistical Learning

Chapter 1 investigated the role of new category learning, old category memory, and simultaneous explicit learning on visual statistical learning. Experiment 1 assigned arbitrary categories to fractal images that did not possess any prior category information and were not easily categorizable by visual features (i.e., the random assignment of fractal images to each group roughly equated average within-category similarity with between-category similarity). Additionally, participants were exposed to statistical regularities while simultaneously learning the category information, which to our knowledge, is the first attempt to investigate whether visual statistical learning can occur simultaneously with explicit learning. Participants both incidentally learned the statistical contingencies within the task as well as the explicit category information. Critically, the novel category information influenced statistical learning such that same-category pairs were learned better than different category pairs.

Reintroducing memory for category information (as well as similarity between stimuli), Experiment 2 used the same paradigm but replaced fractal images with faces and scenes. If category information had to be explicitly attended to while being exposed to statistical regularities, we would have only replicated Experiment 1 insofar as expecting a difference in learning for arbitrary category pairs. Instead, the prior knowledge for the categories (e.g., faces and scenes) continued to predict learning such that same-category information, whether it be newly learned arbitrary category information or memory for existing categories, predicted greater visual statistical learning. Although we could not disentangle visual similarity in Experiment 2 (e.g., two faces have a great deal of visual similarity in common regardless of category information), evidence from Experiment 1 suggests that category information can still drive these differences with visual similarity roughly average between newly learned categories.

Finally, to ensure that visual statistical learning was being driven by category information disassociated from similarity, and not the immediate task demands associated with responding to arbitrary category information, Experiment 3A sought to separate novel category learning from visual statistical learning. We attempted to equate Experiment 3A with Experiments 1 and 2 by maintaining the overall time spent in the experiment through cutting training trials and pairs in half. We used a category learning task that presented the stimuli randomly (i.e., not in pairs) followed by a jitter-detection task that covertly familiarized participants with the pairs. To our surprise, we failed to observe any evidence of visual statistical learning. We conducted Experiment 3B using the same number of pairs but almost twice the number of exposures in both phases of the experiment and, critically, we used a one-back task in place the jitter-detection task to expose participants to statistical regularities.

We again found evidence of visual statistical learning in Experiment 3B, even though the category training was separated from pair exposure. We reasoned that participants were demotivated after the category learning task and only engaged with the jitter-detection task in Experiment 3A insofar as to detect and respond to motion. We compensated for this by increasing the trial count and using a one-back task in Experiment 3B. This speculation was, in part, motivated by prior work that had no trouble observing evidence of visual statistical learning using a jitter-detection task that did not follow a category learning task (Vickery et al., 2018), although the overall degree of learning observed in jitter-detection tasks is also relatively weak to begin with. Importantly, the evidence we found in Experiment 3B highlighted better learning for same-category pairs as compared to different-category pairs, which did not reach above-chance levels of learning.

In light of the findings presented in Chapter 1, it is worthwhile considering the mechanisms underlying the influence of category information on visual statistical learning. Experiment 3B required participants to compare the previous stimulus with the stimulus viewed on the current trial, and we attribute this to our ability to successfully detect evidence of visual statistical learning after our jitter-detection task in Experiment 3A failed (along with an increase in trials). Prior work has recognized that individual working memory capacity may predict visual statistical learning (McCarter, 2021) and that working memory capacity may facilitate the processing of statistical regularities (Cashdollar et al., 2017). By this count, category information may be influencing visual statistical learning in several ways.

Based upon the work from Chapter 1, we hypothesized that the sum total information shared between items constituting a statistical representation predicts the strength of that representation. If visual statistical learning is relying on working memory, it could be the case that shared category information is contributing to the successful extraction of statistical regularity information. Additionally, category switching could be inducing a sort of event boundary, and memory for regularity information has been shown to be better for those items contained within such boundaries (DuBrow & Davachi, 2013, 2016). On the other hand, if working memory is maintaining some item on the current trial, as well as some temporally decayed item from the previous trial, it could be the case that category switching intensifies the decay of the previous item. Thus, rather than same-category information facilitating working memory in the extraction of statistical regularities (or rather than considering this explanation alone in the context of our findings) different-category information could burden working memory which may lead to the creation of weaker statistical representations.

Finally, we have previously discussed the importance of the hippocampus to visual statistical learning (Schapiro et al., 2012) and how patterns of activity within the hippocampus become increasingly correlated as strong representations are built. Additional work has modeled the hippocampus to explore the ways in which different inputs (including regularity information) are extracted from the environment (Schapiro et al., 2017). One interesting route for future work would be to apply different category inputs to such a model and test our hypothesis about how shared information (or similarity between stimuli) drive increased learning between statistical representations. In this way, similarity could include conceptual similarity (like category information) or perceptual similarity (as explored in Chapter 3).

Together, these findings provide critical insights into how statistical learning may operate in the real world. Category information alone (albeit newly learned), as observed in Experiment 1, can influence statistical learning, as can well-established category learning that may be task-irrelevant, as in Experiment 2. Finally, even when newly learned category information immediately precedes exposure to statistical regularities, it can have a profound impact of what statistical representations are formed. Rarely would visual statistical learning in real-world contexts be separated from our wealth of prior knowledge or our constantly changing goals on a moment-to-moment basis, and the set of experiments from Chapter 1 highlight the incredible importance tasks and prior knowledge may have on visual statistical learning in a naturalistic setting.

5.2 Perceptual Similarity Interacts with Task Demands but is Not Altered by Visual Statistical Learning

Chapter 3 focused on investigating the relationship between visual statistical learning and perceptual similarity. Experiment 4 was designed to investigate the impact visual similarity has on visual statistical learning while controlling for the high-level visual information that accompanies complex visual stimuli such as faces and scenes. We used a continuum of basic shape stimuli (Li et al., 2020) that we divided into two categories. Using our arbitrary category learning paradigm as a cover task, we exposed participants to perceptually similar and perpetually dissimilar pairs of stimuli that were covertly structured into same-category and different-category pairs. Although we hypothesized a pattern of visual statistical learning similar to that found in Chapter 2, we instead uncovered a unique pattern that we speculate was engendered, in part, by task demands.

Different-category pairs appeared to be learned better than same-category pairs, particularly for different-category perceptually similar pairs. Why might this be the case? Our decision to define arbitrary categories by dividing a similarity space in half fundamentally changed the task; participants could now use perceptual similarity to guide their learning of the categories. Thus, determination of category membership was hardest for items lying near the borders; different-category items necessarily lay closest to the boundaries, and different-category perceptually similar pairs presented the greatest challenge. This may have led to awareness of the contingency through noticing the difficulty and the similarity of two back-to-back items with different category identities.

Momentarily turning back to Experiment 3A, we explained our inability to observe evidence of visual statistical learning as being due to, in part, a lack of engagement with the jitter-detection task after being exhausted by the categorization task. Experiment 4, by contrast, demands the most effortful engagement for those pairs that require response switching, particularly for the different-category perceptual similar pairs where such subtle differences between stimuli introduced a great deal of ambiguity. If task demands from Experiment 3B (e.g., the switch from a jitter-detection tasks to a one-back task) required participants to engage with stimuli on a deeper level (e.g., compare the current image with the previous image in a one-back task, as opposed to responding to any kind of motion from a jitter-detection task), similar demands for engagement can be found in Experiment 4 for those specific conditions where we observed the most learning.

These findings from Experiment 4 do not provide insight into how perceptual similarity alone may influence visual statistical learning, but they do provide very clear evidence that an interaction between perceptual similarity and task demands can predict which types of statistical regularities are extracted. Although we chose to use basic shape stimuli to avoid visual complexity that was less controllable, the present experiment highlights how challenges relating to perceptual discriminability can drive visual statistical learning. Future work should examine less constrained feature spaces – even the most dissimilar shapes employed in Experiment 4 were still highly similar-looking, compared to a large, diverse, multi-dimensional space that defines natural categories such as faces. The spectrum of basic shape stimuli was arguably less useful in Experiment 5, in which we considered the possibility of visual statistical learning altering perceptual similarity judgements.

Using the basic shape stimuli from Experiment 4, Experiment 5 first exposed participants to statistical regularities in a one-back task. Half of the stimuli were covertly organized into pairs, while the other half appeared randomly as singletons. After exposure, singletons were organized into pairs that were matched in terms of perceptual similarity with pairs that appeared during familiarization, and participants were asked to rate the two items within each pair based on how similar or dissimilar they appear to be. No differences were found between ratings of pairs that possessed statistical regularities (from the familiarization phase) and those that did not, which may suggest that visual statistical learning does not impact judgements of perceptual differences. However, there are a few limitations worth considering in this case.

First, Experiment 5 did not have a visual statistical learning test phase as the rest of the experiments did. This is because we could not reliably test for memory for pairs after revealing the presence and organization of pairs during our perceptual differences judgement task, as participants would have been directly shown which pairs appeared during familiarization when completing their ratings between pairs after learning. It remains a possibility that participants did not learn the visual regularities during familiarization, which lead to no observable differences between pairs that appeared previously and pairs that did not. However, given the fact that we used a one-back task, which has been shown to lead to higher rates of visual statistical learning in prior work (Himberger et al., 2019) as well as Experiment 3B, this possibility seems unlikely. Rather, future work would benefit from reintroducing higher-levels of visual complexity that provide more context upon which perception and similarity judgements may operate, rather than the relatively anemic and feature-constrained basic shape stimuli used in Chapter 3.

Together, the experiments from Chapter 3 provide important insight into how perceptual differences can interact with task demands to influence visual statistical learning (Experiment 4), while also guiding future investigations that may scrutinize how visual statistical learning alters perceptual similarity judgements (Experiment 5). As mentioned previously, we rarely encounter sets of stimuli for which we have little to no knowledge about. However, these efforts to parse the unique impact of perceptual similarity from category related information on visual statistical learning are critical to understanding how perceptual differences may interact with visual statistical learning.

5.3 Methodological Insights for Future Work

The separating of category learning from visual statistical learning in Experiment 3B replicated the category-related effects observed in Experiments 1 and 2, but different-category pairs were not learned at above-chance levels. This may be due to the separation of category learning from visual statistical learning, or it may be due to a task-related influence on participants that had to complete a category learning task before the one-back task that contained statistical regularities. We had speculated that the failure to observe evidence of visual statistical learning in Experiment 3A was due to participants engaging in a relatively easy jitter-detection task after being exhausted by a difficult category-learning task. This was in contrast to prior work which had no issue eliciting evidence of visual statistical learning using a jitter-detection task (Vickery et al., 2018) or even no task at all other than passive viewing (Fiser & Aslin, 2002).

Likewise, Experiment 6 in Chapter 4 was designed to investigate the neural underpinnings of visual statistical learning across tasks. Initially, a within-subjects design wherein participants were exposed to separate statistical regularities across two different tasks seemed like a direct way to examine similarities or differences across tasks that may or may not be consistent with the literature. However, while the present dissertation was interested in task-related influences on visual statistical learning, we did not expect to discover that such influences would be powerful enough to impact learning across independent tasks. Along with the evidence of cross-task influences on participants' ability to extract statistical regularities from Experiments 3A and 3B, Experiment 6 provides critical insight for future work considering the use of multiple statistical learning tasks in a within-subjects design.

Despite the multiple independent opportunities to extract regularities from the tasks, participants were only able to reliably display evidence of visual statistical learning from the categorization task. We had hoped to create a scenario wherein both tasks would have elicited learning so that we may compare the neural correlates of such learning between tasks, but the unexpected discovery of these cross-task influences had made this impossible. Additionally, in retrospect, our design may have also suffered from including too many pairs for subjects to learn effectively during a single learning session. Recalling Experiment 3B, we failed to observe evidence of visual statistical learning for different-category pairs despite using half the number of pairs as in Experiment 1 (16 pairs down to 8), while maintaining nearly just as many presentations. Experiment 6 contained 32 pairs of stimuli, across two different tasks, and appeared within streams that contained non-structured singletons. Along with the unexpected powerful impact of cross-task influences on statistical learning, this design choice may also have contributed to our inability to detect meaningful neural activity between our critical contrasts.

Experiment 6 may have also suffered from a number of other issues that are difficult to specify or remediate. First, there were only four runs of the main passive-viewing task in the scanner, which may have constituted too little within-subject power for detecting neural differences. The passive-viewing task, itself, may have interfered with previous learning, serving to nullify differences between same- and different-category pairs. We chose to intermingle images from both the category learning and one-back learning tasks, and preserved pairwise contingencies. Any or all of these factors may have played against us, but it is difficult to say which ones are most important. Another important consideration is that it is possible that previous detection of differences were false alarms or otherwise rare successes, and that in general, fMRI studies of visual statistical learning are underpowered for the size of effects that should be expected.

Taken together with Experiment 6, future work must use caution considering designs that use multiple statistical learning tasks within-subjects. Even with a reduction of the number of images covertly presented in pairs, visual statistical learning may be interrupted, as potentially evidenced in Experiment 3B. The use of a task that presents statistically structured pairs simultaneously with singletons should be considered carefully; the present dissertation does not provide insight into the impact singleton presence can have on learning, but no evidence to date suggests that it does not have an impact. A future iteration of Experiment 6 may benefit from the use of fewer pairs, elimination of singletons to control for any additional unknown influences, and the use of tasks that are better equated on difficulty and engagement (e.g., a categorization task and a two-back task, rather than a one-back task).

5.4 Conclusion

To our knowledge, we have provided the first evidence that the explicit learning of category information and the incidental learning of visual statistical information can co-occur, and both novel category information and prior well-learned category information can impact visual statistical learning. Even when novel category learning is entirely separated from (i.e., immediately precedes) exposure to statistical regularities, such category information predicts what visual statistical information is learned. If visual statistical learning is a highly automatic form of incidental learning that is constantly extracting regularities from our environment, these findings shed light on how newly acquired as well as old, well-learned category information can predict what is learned.

We also found evidence that task demands and perceptual similarity may interact in such a way that predicts relatively better learning for certain regularities. We used our category-learning task while exposing participants to pairs of stimuli that were either visually similar or visually dissimilar and found a pattern of results that harshly contrasted with our earlier experiments; we argue that the visually ambiguous category boundaries drove participants to more effortfully engage with different-category pairs than same-category pairs, which led to higher rates of learning for different-category pairs. If visual similarity drives visual statistical learning, our evidence suggests that subtle but important (i.e., category-defining) differences for which task demands force participants to prioritize predict greater learning for statistical contingencies. Additional work use may use this understanding investigate other ways perceptual similarity drives statistical learning, especially when considering the use of more complex stimuli, as our investigation into how visual statistical learning may inversely alter perceptual similarity ratings yielded no differences.

Finally, along with the numerous influences of task demands on visual statistical learning, we found that exposure to statistical regularities across multiple, different contexts revealed uniquely different patterns of learning within subjects. Although our use of a one-back task was effective in eliciting statistical learning in earlier experiments with fewer instances of statistical regularities, simultaneous exposure to differing tasks (e.g., a category-learning task and a one-back task), each with their own set of statistical regularities, can introduce cross-task influences in such a way that we only observed evidence of learning in the category-learning task and not the one-back task. What is more clear than ever, as a result of the work put forth by the present dissertation, is that task influences are even more powerful and pervasive than previously believed.

Altogether, the present dissertation has provided evidence that visual statistical learning may be influenced by a number of factors that are present in day-to-day functioning. Newly learned category information can differentially predict which regularities are extracted from our environment, as well as long-standing prior knowledge about categories. As we navigate our world with whatever our present goals may be, perceptual similarity also interacts with visual statistical learning in a way that predicts what may be learned. Finally, how statistical regularities may or may not be extracted across multiple tasks, and what ultimately predicts successful learning, remains a potentially fruitful avenue for future research. The present dissertation has revealed multiple ways prior knowledge, perceptual similarity, and task demands can operate on visual statistical learning. From this, many more paths may be open for investigation into how this unique form of incidental learning serves us in real world scenarios.

REFERENCES

Akyürek, E. G., Hommel, B., & Jolicoeur, P. (2007). Direct evidence for a role of working memory in the attentional blink. *Memory & Cognition*, *35*(4), 621–627.

Ashby, F. G., & Maddox, W. T. (2005). Human Category Learning. *Annual Review of Psychology*, *56*(1), 149–178. https://doi.org/10.1146/annurev.psych.56.091103.070217

Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*(4), 321–324. https://doi.org/10.1111/1467-9280.00063

Baker, C. I., Olson, C. R., & Behrmann, M. (2004). Role of attention and perceptual grouping in visual statistical learning. *Psychological Science*, *15*(7), 460–466. https://doi.org/10.1111/j.0956-7976.2004.00702.x

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*(3), 211–227. https://doi.org/10.3758/BF03196968

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive psychology*, *14*(2), 143-177.

Bomba, P. C., & Siqueland, E. R. (1983). The nature and structure of infant form categories. *Journal of Experimental Child Psychology*, *35*(2), 294–328. https://doi.org/10.1016/0022-0965(83)90085-1

Brady, T. F., & Oliva, A. (2008). Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent. *Psychological science*, *19*(7), 678-685. https://doi.org/10.1111/j.1467-9280.2008.02142.x

Campbell, K. L., Zimerman, S., Healey, M. K., Lee, M., & Hasher, L. (2012). Age differences in visual statistical learning. *Psychology and Aging*, *27*(3), 650. https://doi.org/10.1037/a0026780

Cashdollar, N., Ruhnau, P., Weisz, N., & Hasson, U. (2017). The Role of Working Memory in the Probabilistic Inference of Future Sensory Events. *Cerebral Cortex*, *27*(5), 2955–2969. https://doi.org/10.1093/cercor/bhw138

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive psychology*, *36*(1), 28-71. https://doi.org/10.1006/cogp.1998.0681

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*(3), 162–173. https://doi.org/10.1006/cbmr.1996.0014

DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*, *142*(4), 1277–1286. https://doi.org/10.1037/a0034024

DuBrow, S., & Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of Learning and Memory*, *134*, 107–114. https://doi.org/10.1016/j.nlm.2016.07.011

Etcoff, N. L., & Magee, J. J. (1992). Categorical perception of facial expressions. *Cognition*, *44*(3), 227–240. https://doi.org/10.1016/0010-0277(92)90002-Y

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499–504.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467.

Goldstone, R. (1994). Influences of Categorization on Perceptual Discrimination. *Journal of Experimental Psychology: General*, *123*(2), 178–200. https://doi.org/10.1037/0096-3445.123.2.178

Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition*, *78*(1), 27–43. https://doi.org/10.1016/S0010-0277(00)00099-8

Himberger, K. D., Finn, A. S., & Honey, C. J. (2019). *Reconsidering the Automaticity of Visual Statistical Learning*. https://doi.org/https://doi.org/10.31234/osf.io/r659w

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *NeuroImage*, *62*(2), 782–790. https://doi.org/10.1016/j.neuroimage.2011.09.015

Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, *23*(2), 141–221. https://doi.org/10.1016/0010-0285(91)90009-D

Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive psychology*, *15*(4), 483-524. https://doi.org/10.1016/0010-0285(83)90017-8

Lengyel, G., Nagy, M., & Fiser, J. (2021). Statistically defined visual chunks engage object-based attention. *Nature Communications*, *12*(1). https://doi.org/10.1038/s41467-020-20589-z

Li, A. Y., Liang, J. C., Lee, A. C. H., & Barense, M. D. (2020). The validated circular shape space: Quantifying the visual similarity of shape. *Journal of Experimental Psychology: General*, *149*(5), 949. https://doi.org/10.1037/xge0000693

Liberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, *54*(5), 358–368. https://doi.org/10.1037/h0044417

McCarter, A. (2021). *The Role of Working Memory in Visual Statistical Learning*. University of Delaware.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419. https://doi.org/10.1037/0033-295X.102.3.419

Musz, E., Weber, M. J., & Thompson-Schill, S. L. (2015). Visual statistical learning is not reliably modulated by selective attention to isolated events. *Attention, Perception, & Psychophysics*, *77*(1), 78–96. https://doi.org/10.3758/s13414-014-0757-5

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*(1), 1–32. https://doi.org/10.1016/0010-0285(87)90002-8

Otsuka, S., Nishiyama, M., Nakahara, F., & Kawaguchi, J. (2013). Visual Statistical Learning Based on the Perceptual and Semantic Information of Objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(1), 196–207. https://doi.org/10.1037/a0028645

Palmer, S., Rock, I. (1994). Rethinking perceptual organization: The role of uniform connectedness. In *Psychonomic Bulletin & Review* 1**,** 29–55. https://doi.org/10.3758/BF03200760

Peterson, M. A. (1994). Object recognition processes can and do operate before figure–ground organization. *Current Directions in Psychological Science*, *3*(4), 105–111.

Quinn, P. C. (1987). The categorical representation of visual pattern information by young infants. *Cognition*, *27*(2), 145–179. https://doi.org/10.1016/0010-0277(87)90017-5

Richler, J. J., & Palmeri, T. J. (2014). Visual category learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(1), 75–94. https://doi.org/10.1002/wcs.1268

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*(11), 1019–1025. https://doi.org/10.1038/14819

Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. In *Wiley Interdisciplinary Reviews: Cognitive Science* (Vol. 1, Issue 6, pp. 906–914). https://doi.org/10.1002/wcs.78

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928. https://doi.org/10.1126/science.274.5294.1926

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*(1), 27–52. https://doi.org/10.1016/S0010-0277(98)00075-4

Schapiro, A. C., Kustner, L. v., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*(17), 1622–1627. https://doi.org/10.1016/j.cub.2012.06.056

Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711). https://doi.org/10.1098/rstb.2016.0049

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*(4820), 1317–1323. https://doi.org/10.1126/science.3629243

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, *17*(3), 143–155. https://doi.org/10.1002/hbm.10062

Spelke, E. S. (1990). Principles of object perception. *Cognitive science*, *14*(1), 29-56. https://doi.org/10.1016/0364-0213(90)90025-R

Strauss, M. S. (1979). Abstraction of prototypical information by adults and 10-month-old infants. *Journal of Experimental Psychology: Human Learning and Memory*, *5*(6), 618. https://doi.org/10.1037/0278-7393.5.6.618

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564. https://doi.org/10.1037/0096-3445.134.4.552

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*(10), 1934–1945. https://doi.org/10.1162/jocn.2009.21131

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit Perceptual Anticipation Triggered by Statistical Learning. *Journal of Neuroscience*, *30*(33), 11177–11187. https://doi.org/10.1523/JNEUROSCI.0858-10.2010

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*(4), 327. https://psycnet.apa.org/doi/10.1037/0033-295X.84.4.327

Vickery, T. J., Park, S. H., Gupta, J., & Berryhill, M. E. (2018). Tasks determine what is learned in visual statistical learning. *Psychonomic Bulletin & Review*, *25*(5), 1847–1854. https://doi.org/10.3758/s13423-017-1405-6

Younger, B., & Gotlieb, S. (1988). Development of categorization skills: Changes in the nature or structure of infant form categories? *Developmental Psychology*, *24*(5), 611. doi:10.1037/0012-1649.24.5.611

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention Is Spontaneously Biased Toward Regularities. *Psychological Science*, *24*(5), 667–677. https://doi.org/10.1177%2F0956797612460407

Zhao, J., Ngo, N., McKendrick, R., & Turk-Browne, N. B. (2011). Mutual interference between statistical summary perception and statistical learning. *Psychological Science*, *22*(9), 1212–1219. https://doi.org/10.1177/0956797611419304

Zhou, H., Friedman, H. S., & Diger Von Der Heydt, R. (2000). Coding of Border Ownership in Monkey Visual Cortex. *Journal of Neuroscience*, *20*(17), 6594-6611. https://doi.org/10.1523/JNEUROSCI.20-17-06594.2000

APPENDIX A

IRB APPROVAL

UNIVERSITY OF DELAWARE

Institutional Review Board
210H Hullihen Hall
Newark, DE 19716
Phone: 302-831-2137
Fax: 302-831-2828

DATE:	July 12, 2019
TO:	Timothy Vickery, PhD
FROM:	University of Delaware IRB
STUDY TITLE:	[630382-7] Production and Perception of Patterns
SUBMISSION TYPE:	Continuing Review/Progress Report
ACTION:	APPROVED
APPROVAL DATE:	July 12, 2019
EXPIRATION DATE:	July 14, 2020
REVIEW TYPE:	Expedited Review
REVIEW CATEGORY:	Expedited review category # *(7)*

Thank you for your Continuing Review/Progress Report submission to the University of Delaware Institutional Review Board (UD IRB). The UD IRB has reviewed and APPROVED the proposed research and submitted documents via Expedited Review in compliance with the pertinent federal regulations.

As the Principal Investigator for this study, you are responsible for and agree that:

- All research must be conducted in accordance with the protocol and all other study forms as approved in this submission. Any revisions to the approved study procedures or documents must be reviewed and approved by the IRB prior to their implementation. Please use the UD amendment form to request the review of any changes to approved study procedures or documents.
- Informed consent is a process that must allow prospective participants sufficient opportunity to discuss and consider whether to participate. IRB-approved and stamped consent documents must be used when enrolling participants and a written copy shall be given to the person signing the informed consent form.
- Unanticipated problems, serious adverse events involving risk to participants, and all non-compliance issues must be reported to this office in a timely fashion according with the UD requirements for reportable events. All sponsor reporting requirements must also be followed.

Oversight of this study by the UD IRB REQUIRES the submission of a CONTINUING REVIEW seeking the renewal of this IRB approval, which will expire on July 14, 2020. A continuing review/progress report form and up-to-date copies of the protocol form and all other approved study materials must be submitted to the UD IRB at least 45 days prior to the expiration date to allow for the required IRB review of that report.

If you have any questions, please contact the UD IRB Office at (302) 831-2137 or via email at hsrb-research@udel.edu. Please include the study title and reference number in all correspondence with this office.

UNIVERSITY OF DELAWARE

Institutional Review Board
210H Hullihen Hall
Newark, DE 19716
Phone: 302-831-2137
Fax: 302-831-2828

DATE: February 25, 2020

TO: Timothy Vickery
FROM: University of Delaware IRB

STUDY TITLE: [1467811-2] Studies of Human Cognition
SUBMISSION TYPE: Amendment/Modification

ACTION: DETERMINATION OF EXEMPT STATUS
EFFECTIVE DATE: February 25, 2020

REVIEW CATEGORY: Exemption category # *(3)*

Thank you for your Amendment/Modification submission to the University of Delaware Institutional Review Board (UD IRB). According to the pertinent regulations, the UD IRB has determined this project is EXEMPT from most federal policy requirements for the protection of human subjects. The privacy of subjects and the confidentiality of participants must be safeguarded as prescribed in the reviewed protocol form.

This exempt determination is valid for the research study as described by the documents in this submission. Proposed revisions to previously approved procedures and documents that may affect this exempt determination must be reviewed and approved by this office prior to initiation. The UD amendment form must be used to request the review of changes that may substantially change the study design or data collected.

Unanticipated problems and serious adverse events involving risk to participants must be reported to this office in a timely fashion according with the UD requirements for reportable events.

A copy of this correspondence will be kept on file by our office. If you have any questions, please contact the UD IRB Office at (302) 831-2137 or via email at hsrb-research@udel.edu. Please include the study title and reference number in all correspondence with this office.

INSTITUTIONAL REVIEW BOARD

www.udel.edu

Institutional Review Board
210H Hullihen Hall
Newark, DE 19716
Phone: 302-831-2137
Fax: 302-831-2828

UNIVERSITY OF DELAWARE

DATE:	November 14, 2019
TO:	Timothy Vickery, PhD
FROM:	University of Delaware IRB
STUDY TITLE:	[889240-6] Reinforcement learning in the human brain: Dimensions, features, and contexts
SUBMISSION TYPE:	Amendment/Modification
ACTION:	APPROVED
APPROVAL DATE:	November 13, 2019
EXPIRATION DATE:	May 31, 2020
REVIEW TYPE:	Expedited Review
REVIEW CATEGORY:	Expedited review category # *(4,7)*

Thank you for your Amendment/Modification submission to the University of Delaware Institutional Review Board (UD IRB). The UD IRB has reviewed and APPROVED the proposed research and submitted documents via Expedited Review in compliance with the pertinent federal regulations.

As the Principal Investigator for this study, you are responsible for and agree that:

- All research must be conducted in accordance with the protocol and all other study forms as approved in this submission. Any revisions to the approved study procedures or documents must be reviewed and approved by the IRB prior to their implementation. Please use the UD amendment form to request the review of any changes to approved study procedures or documents.
- Informed consent is a process that must allow prospective participants sufficient opportunity to discuss and consider whether to participate. IRB-approved and stamped consent documents must be used when enrolling participants and a written copy shall be given to the person signing the informed consent form.
- Unanticipated problems, serious adverse events involving risk to participants, and all non-compliance issues must be reported to this office in a timely fashion according with the UD requirements for reportable events. All sponsor reporting requirements must also be followed.

Oversight of this study by the UD IRB REQUIRES the submission of a CONTINUING REVIEW seeking the renewal of this IRB approval, which will expire on May 31, 2020. A continuing review/progress report form and up-to-date copies of the protocol form and all other approved study materials must be submitted to the UD IRB at least 45 days prior to the expiration date to allow for the required IRB review of that report.

If you have any questions, please contact the UD IRB Office at (302) 831-2137 or via email at hsrb-research@udel.edu. Please include the study title and reference number in all correspondence with this office.

APPENDIX B

PERMISSIONS

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

Apr 13, 2022

This Agreement between University of Delaeware -- Leeland Rogers ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	5287130280628
License date	Apr 13, 2022
Licensed Content Publisher	Springer Nature
Licensed Content Publication	Psychonomic Bulletin & Review
Licensed Content Title	Visual statistical learning is modulated by arbitrary and natural categories
Licensed Content Author	Leeland L. Rogers et al
Licensed Content Date	Mar 31, 2021
Type of Use	Thesis/Dissertation
Requestor type	academic/university or research institute
Format	electronic
Portion	full article/chapter

Will you be translating?	no
Circulation/distribution	1 - 29
Author of this Springer Nature content	yes
Title	Visual statistical learning is modulated by arbitrary and natural categories
Institution name	University of Delaware
Expected presentation date	May 2022
Requestor Location	University of Delaeware 105 THE GRN RM 108 WOLF HALL NEWARK, DE 19716 United States Attn: University of Delaeware
Total	0.00 USD