STATISTICAL LEARNING CREATES NOVEL OBJECT ASSOCIATIONS VIA TRANSITIVE RELATIONS

by

Yu Luo

B.Sc., The University of British Columbia, 2015

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ARTS

in

The Faculty of Graduate and Postdoctoral Studies

(Psychology)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August 2018

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, a thesis/dissertation entitled:

Statistical learning creates novel object associations via transitive relations

submitted by	Yu Luo	in partial fulfillment of the requirements for
the degree of	Master of Arts	
in	Psychology	

Examining Committee:

Jiaying Zhao, Psychology
Supervisor

Ronald A. Rensink, Psychology
Supervisory Committee Member

Lawrence Ward, Psychology
Supervisory Committee Member

Additional Examiner

Additional Supervisory Committee Members:

Supervisory Committee Member

Supervisory Committee Member

Abstract

A remarkable ability of the cognitive system is to make novel inferences based on prior experiences. What mechanism supports such inference? We propose that statistical learning is a process where transitive inferences of new associations are made between objects that have never been directly associated. After viewing a continuous sequence containing two base pairs (e.g., A-B, B-C), participants automatically inferred a transitive pair (e.g., A-C) where the two objects had never co-occurred before (Experiment 1). This transitive inference occurred in the absence of explicit awareness of the base pairs. However, participants failed to infer the transitive pair from three base pairs (Experiment 2), showing the limits of the transitive inference (Experiment 3). We further demonstrated that this transitive inference can operate across the categorical hierarchy (Experiments 4-7). The findings revealed a novel consequence of statistical learning where new transitive associations between objects are implicitly inferred.

Lay Summary

How the mind generates new knowledge has fascinated psychologists and philosophers for centuries. The mechanism supporting this ability is still not well understood. Here we explore how statistical learning, a basic form of visual learning, allows the creation of new knowledge through transitive inference. We demonstrate that people are able to not only learn relationships between objects that have appeared together in the past, but also successfully infer new relationships between objects that have never been directly associated with each other before. By revealing how statistical learning forms new associations between objects via transitive inference and the limits of this inference, the study offers new insights into how the mind creates new knowledge based on prior experiences.

Preface

Chapter 2, 3, and 4 are based on work conducted in the University of British Columbia's Behavioral Sustainability Lab by Yu Luo and Dr. Jiaying Zhao. Dr. Jiaying Zhao and I developed the study concept and experimental design. I collected the data and analyzed the data under the supervision of Dr. Jiaying Zhao. Experimental preparation and data collection were helped by the members at the Behavioral Sustainability Lab.

Chapter 2, 3, and 4 of the thesis have recently been accepted in Psychological Science (Luo & Zhao, in press). I conducted all the data collection and drafted the manuscript.

All experiments reported here were approved by the UBC Behavioral Research Ethics Board under the project entitled: Attentional control driven by statistical learning (BREB number: H13-02684).

Table of Contents

Abstract ..... **iii**

Lay Summary ..... **iv**

Preface ..... **v**

Table of Contents ..... **vi**

List of Figures ..... **viii**

Acknowledgements ..... **ix**

1 Introduction **1**

2 Statistical learning creates novel associations **8**

2.1 Experiment 1 8

2.1.1 Participants 8

2.1.2 Apparatus 8

2.1.3 Stimuli 8

2.1.4 Procedure 9

2.1.5 Results and Discussion 11

2.2 Experiment 2 12

2.2.1 Participants 12

2.2.2 Stimuli and Procedure 12

2.2.3 Results and Discussion 13

2.3 Experiment 3 14

2.3.1 Participants 14

2.3.2 Stimuli and Procedure 14

2.3.3 Results and Discussion 15

3 Novel association across categorical levels **17**

3.1 Experiment 4a 17

3.1.1 Participants 17

3.1.2 Stimuli 17

3.1.3 Procedure 18

3.1.4 Results and Discussion 20

3.2 Experiment 4b 21

3.2.1 Participants 21

3.2.2 Stimuli and Procedure 22

3.2.3 Results and Discussion 22

3.3 Experiment 5 23

3.3.1 Participants.. 24
3.3.2 Stimuli and Procedure 24
3.3.3 Results and Discussion.. 24
3.4 Experiment 6 25
3.4.1 Participants.. 25
3.4.2 Stimuli and Procedure 26
3.4.3 Results and Discussion.. 26
4 Transitive inference within and across categorical levels 28
4.1 Experiment 7a 28
4.1.1 Participants.. 28
4.1.2 Stimuli.. 28
4.1.3 Procedure 29
4.1.4 Results and Discussion.. 31
4.2 Experiment 7b 32
4.2.1 Participants.. 32
4.2.2 Stimuli and Procedure 32
4.2.3 Results and Discussion.. 34
5 General Discussion.. 37
Bibliography 43

List of Figures

Figure 1 Stimuli and results of Experiment 1 11
Figure 2 Stimuli and results of Experiment 2 13
Figure 3 Stimuli and results of Experiment 3 15
Figure 4 Stimuli and results of Experiment 4 19
Figure 5 Stimuli and results of Experiment 4b 23
Figure 6 Stimuli and results of Experiment 5 24
Figure 7 Stimuli and results of Experiment 6 26
Figure 8 Stimuli and results of Experiment 7a 30
Figure 9 Stimuli and results of Experiment 7b 33

Acknowledgements

I would like to thank my lab mates from Behavioral Sustainability Lab who helped with data collection and gave helpful suggestions. I would like to thank faculty members and fellow graduate students in the Department of Psychology who provided thoughtful feedback and comments. I really appreciate having Dr. Ronald A. Rensink and Dr. Lawrence Ward on my committee. I would also like to thank an anonymous reviewer from *Psychological Science* for suggesting the experiment in Chapter 4. I sincerely appreciate the generous funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Department of Psychology. Most of all, I would like to thank my supervisor, Dr. Jiaying Zhao, for her support and outstanding mentorship.

1 Introduction

A remarkable ability of the human cognitive system is its ability to make new inferences based on prior experiences. For instance, if we know that A is related to B and B is related to C, we might automatically think that A is sometimes related to C. What cognitive mechanisms support such transitive inference? One mechanism is logical reasoning. If Ann is taller than Beth, and Beth is taller than Cath, then people can infer that Ann is taller than Cath using logical reasoning (Goodwin & Johnson-Laird, 2005). Although such inference often requires deliberate and conscious reasoning, there is no doubt that healthy adult humans are able to perform such transitive judgments. However, whether young children can successfully infer such transitive relationships using logical reasoning remains controversial. Piaget (1928) claimed that children approximately below 7-year-old are incapable to perform such logical reasoning. Contrary to Piaget's view, a study has demonstrated that when 4-year-old children were trained to learn the descending order of rods' size without visual feedback (e.g., blue was larger than red and red was larger than green), they could successfully infer all the possible transitive comparisons (e.g., blue was larger than green; Bryant & Trabasso, 1971). These findings suggest that young children can make transitive judgments, as long as the order of the overlapping premise pairs is remembered.

More impressively, this ability is not only limited to humans, but it's also observed in other species. In one study, squirrel monkeys were trained to learn the relations between five different colored containers following a paradigm similar to Bryant and Trabasso's method (1971). For example, when choosing a yellow container against a blue container resulted in a reward and choosing a blue container against a green container resulted in a reward, the squirrel monkeys successfully chose the yellow container against a green container. This suggests that they made transitive judgments to a certain extent to solve the problems (McGonigle & Chalmers, 1977). In other studies, researchers have demonstrated similar results in chimpanzees (Gillan, 1981), pigeons (Von Fersen, Wynne, Delius, & Staddon, 1991), and rats (Davis, 1992).

These studies have shown that human and nonhuman species can successfully make transitive inference after extensive training on the order between objects, meaning that at least in human subjects, participants are explicitly aware of the order during training. For example, in Bryant & Trabasso's study (1971), children were explicitly told to select the "taller" or the "shorter" rod during the training. However, other studies have demonstrated that transitive relations can also be formed without awareness of the embedding order in the premise pairs during training. In one study, participants were not informed of the order of the objects at the beginning of the experiment and were trained to learn the embedding order using trial-by-trial feedback task. In this task, one object in the premise pair is always correct and the other is always incorrect (e.g., in A>B premise pair, A is always correct and B is always incorrect). In the test phase, participants can successfully infer the transitive pairs (e.g., if A>B and B>C, then A>C) based on the premise pairs that were learned without awareness of the embedding order (Greene, Spellman, Dusek, Eichenbaum, & Levy 2001). In addition, participants' performance in the transitive task does not correlate with their task awareness. In another study, researchers have also shown that without providing explicit feedback on the embedding order between objects in the training phase, participants can infer the transitive relations, although their performance is lower than participants who received explicit feedback, such as "older" rather than "correct", in a trial-by-trial feedback task (Kumaran & Ludwig, 2013). Thus, these studies suggest that people can learn the embedding order between objects and make further transitive judgments without knowing the embedding order during training.

Based on the studies mentioned above, transitive inference depends not only on the ability to learn the embedding order between objects but also on the ability to retrieve previous knowledge. Several studies have shown that the hippocampus plays an important role in retrieving past knowledge to support inferential reasoning (see review Zeithamova, Schlichting, & Preston, 2012). For example, the hippocampus is activated to retrieve that A is paired with B and B is paired with C when the AC pair was tested. In a recent study, the hippocampus was found to support the transfer of values across objects that were previously associated (Wimmer & Shohamy, 2012). Specifically, participants learned pairs of objects (e.g., S_1-S_2) through associative learning and then learned that half of S_2 is followed by a monetary reward (i.e., $S_{2+)}$ and the other half of S_2 is followed by a neutral outcome (i.e., S_{2-}). In the test phase, participants were biased toward choosing S_1 that was associated with S_{2+} rather than S_1 that was associated with S_{2-}, because the value of S_1 that was associated with S_{2+} was increased by the monetary reward linked to S_{2+}. With further analysis, hippocampal activation during reward learning was found to be greater for the stimulus that led in greater choice bias, and additionally, reactivation of previously associated stimuli (e.g., S_1) with S_{2+} via hippocampus during reward learning is necessary to transfer the positive value of reward to S_1. Given that we can reactivate previously learned relations between objects when making new inferences, we propose a second mechanism that supports transitive inferences: statistical learning, a process that establishes connections between object representations, which can allow transitive inferences to form based on prior experiences.

Statistical learning is a mechanism that supports the extraction of regularities between individual objects in terms of how they co-occur over time (Saffran, Aslin, & Newport, 1996) or space (Fiser & Aslin, 2001). A pioneering study in statistical learning demonstrated that 8-

month-old infants were able to distinguish the specific order of temporally co-occurring syllables (e.g., bi, da, ku) from the same syllables presented in a random order, after being exposed to the regularities for only two minutes (Saffran et al., 1996). In addition to temporal relationships between objects, people can visually extract spatial relationships between objects. In one study, participants were exposed to arrays in which certain objects reliably co-occurred in spatial position (e.g., A always on top of B), and participants were able to discriminate the spatial pair against a foil (e.g., A on top of B against A on top of another shape; Fiser & Aslin, 2001). In addition to the auditory and visual domain, statistical learning can also operate in the tactile domain (Conway & Christiansen, 2005).

These pioneering studies have shown our remarkable ability to extract regularities in different sensory domains, but the stimuli in these studies were limited to a single feature. Given that real-world objects are more complex and have multiple features, such as color, shape, or texture, one study has demonstrated that statistical learning can perform at multiple feature dimensions (Turk-Browne, Isola, Scholl, & Treat, 2008). In this study, participants viewed a sequence of triplets consisting of colored shapes (e.g., blue diamond-yellow circle-red triangle). After exposure, one group was tested on the triplets with combined features, and another group was tested separately on the color features (e.g., blue-yellow-red) and on the shape features (e.g., diamond-circle-triangle). Participants showed much higher learning for triplets with combined features than separated features, suggesting that statistical learning operates over multidimensional objects in which low-level visual features are bonded together.

Statistical learning is not only limited to the extraction of regularities in the environment but also has multiple consequences on our cognitive system. One major cognitive consequence is that it can spontaneously draw attention to the co-occurring objects (Zhao, Al-Aidroos, & Turk-

Browne, 2013). In this study, participants were exposed to streams of shapes that were shown at the top, bottom, left, and right of the central fixation. Importantly, one of the four streams contained temporal regularities (i.e., structured stream) and the other three stream has no regularities (i.e., random stream). During exposure, participants occasionally performed a visual search task in which they searched for a *T* target among three *L* distractors in the four locations. The results showed that the response time was faster when the *T* target was in the structured stream than when the *T* target was in the random stream, suggesting an attentional bias for regularities during statistical learning. In a follow-up study, this bias was found to persist when the regularities were removed or when new regularities were embedded in one of the previous random streams. However, this bias was weakened when the initial regularities in the structured stream were removed and new regularities were introduced in one of the previous random streams, or when novel shapes were introduced in the random streams (Yu & Zhao, 2015).

Given that attention is biased toward statistical regularities, another study has examined whether statistical regularities can prioritize local and global processing. Participants were exposed to arrays of Navon-like figures consisting of small objects making up a global form (Navon, 1977), and their task was to identify the local object or the global form. Unbeknownst to the participants, the Navon figures contain local regularities (e.g., A always appeared next to B), or global regularities (e.g., A and B co-occurred globally). As a result, local object identification was faster when local regularities were embedded, and global form identification was faster when global regularities were embedded. These findings suggest that attention is biased to the local scale when local regularities are presented, prioritizing the processing of individual objects, and attention is biased to the global scale when global regularities are presented, prioritizing the processing of the global form (Zhao & Luo, 2017).

In addition to its consequence on attention, statistical learning can facilitate the compression of redundant information in the environment. One study has shown that participants can remember more colors on a display when the colors are frequently paired together than when the colors are presented randomly on the display, suggesting that statistical regularities can be used to facilitate memory encoding by compressing information (Brady, Konkle, & Alvarez, 2009). In another study, participants viewed arrays of colored dots and estimated the number of dots in the array. In one condition the colored dots were grouped into pairs (i.e., structured condition), and in another condition, the colored dots were randomly placed in the array (i.e., random condition). The results showed that participants in the structured condition provided lower estimates than those in the random condition, suggesting that statistical regularities reduce perceived numerosity by grouping information (Zhao & Yu, 2016). Thus, both studies suggest that object co-occurrences can be used as grouping cues to compress information.

Among all the characteristics of statistical learning, a distinct feature of statistical learning is that this process occurs incidentally and automatically, without conscious intent or explicit awareness, since observers tend not to be explicitly aware of object co-occurrences (Turk-Browne, Jungé, & Scholl, 2005; Turk-Browne, Scholl, Chun & Johnson, 2009). Given the automatic and implicit nature of statistical learning, an unexplored question is whether statistical learning forms new associations between objects that have never co-occurred before and can only be associated via transitive relations.

The goal of this thesis was to examine whether statistical learning is a process in which new transitive associations are created among objects that have never been directly associated in the following three chapters:

1) After being exposed to temporal relationships between objects, can new associations be formed between objects that had never appeared together? For example, given the pairs, A-B and B-C, can people automatically infer a new pair A-C (Experiment 1)? If statistical learning supports such transitive inference, are there any limits in the transitive inference (Experiments 2-3)?

2) Can people form new association based on another type of transitivity from set theory (Ciesielski, 1997) using hierarchical relations between objects? Specifically, based on exposures to the association between two objects at one categorical level (e.g., New York-London), can people infer the same association at the subordinate level (e.g., Central Park-Hyde Park) and the superordinate level (e.g., USA-UK) (Experiment 4)? Are there any limits in such transitive inference across hierarchical categories (Experiments 5-6)?

3) Can people perform more demanding inference such as operating two types of transitivity simultaneously? Specifically, given the pairs A-B and B-C at one categorical level (e.g., New York-London and London-Vancouver), can people infer a new pair A-C at the subordinate level (e.g., Central Park-Stanley Park) and the superordinate level (e.g., USA-Canada, Experiment 7)?

2 Statistical learning creates novel associations

2.1 Experiment 1

This experiment examined whether new associations could be formed between objects that had never appeared together.

2.1.1 Participants

Forty undergraduates (28 female; mean age=20.4 years, SD=2.3) from University of British Columbia (UBC) participated for course credit. Participants in all experiments reported normal or corrected-to-normal vision and provided informed consent. All experiments reported here were approved by the UBC Behavioral Research Ethics Board. A power analysis was conducted using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007). Given a Cohen's *d* of 0.54 based on a prior study (Yu & Zhao, 2015), a minimum of 39 participants was required to have 90% power (alpha=0.05) to reveal the effect in our experiments.

2.1.2 Apparatus

Participants in all experiments were seated 60cm from a computer monitor (refresh rate=60Hz). Stimuli were presented using MATLAB (Mathworks) and Psychophysics Toolbox (http://psychtoolbox.org).

2.1.3 Stimuli

The stimuli consisted of nine circles in nine distinct colors (color name=R/G/B values: red=255/0/0; green=0/255/0; blue=0/0/255; yellow=255/255/0; magenta=255/0/255; cyan=0/255/255; gray=185/185/185; brown=103/29/0; black=0/0/0). Each circle subtended 1.6° of visual angle. The colored circles were randomly assigned into six base pairs for each participant and remained constant throughout the experiment. The six base pairs contained three sets of two base pairs. In each group, the second color in the first pair was the same as the first color in the second pair (e.g., A-B, B-C; Figure 1a). This allowed us to test if people could automatically infer a transitive pair (A-C) given the two base pairs (A-B and B-C). There were three transitive pairs from the six base pairs. Importantly, the two colors in the transitive pair had never directly followed each other. Each base pair was repeated 50 times to form a single continuous temporal sequence of colored circles presented in a pseudorandom order with two constraints: no single base pair could repeat back-to-back and no two base pairs with a shared color (e.g., A-B, B-C) could be presented consecutively. Thus, the inference of the transitive pair cannot be driven by non-adjacent dependencies.

2.1.4 Procedure

The experiment contained two phases: exposure phase and test phase. During the exposure phase, one colored circle appeared at the center of the screen for 500ms followed by a 500ms inter-stimulus interval (ISI) in each trial. Participants performed a 1-back task where they judged as quickly and accurately as possible whether the current circle had the same color as the previous circle (by pressing the "/" or "z" key for same or different, respectively, key assignment counterbalanced). For the 1-back task, each color had a 20% chance of repeating the previous color, producing a total of 720 trials. Specifically, each member in the pair had 20% chance of repeating itself (e.g., A A B or A B B). This 1-back task served as a cover task irrelevant to learning, in order to conceal the true purpose of the study, ensuring that learning of the color pairs was incidental. Participants were not told anything about the color pairs.

After exposure, participants completed a surprise two-alternative forced choice (2AFC) test phase to examine whether they had successfully learned the base pairs. In each trial, participants viewed two sequences of circles at fixation. Each circle appeared for 500ms followed by a 500ms ISI, and each sequence was separated by a 1000ms pause. Participants judged whether the first or second sequence looked more familiar based on the exposure phase. If they did not respond during the sequence presentation or ISI, the screen remained blank until response. One sequence was a base pair, and the other was a "foil" (e.g., A-E) composed of one color from a base pair (e.g., A-B), and the other from a different base pair (e.g., D-E) while preserving the temporal positions in the pairs. Thus, the two colors in the foil had never directly followed each other during exposure. Each base pair was tested against a foil, which was then repeated, resulting in 12 trials in total. It's important to note that each base pair and foil were presented the same number of times at test. Thus, to discriminate the base pair from the foil, participants needed to know which two particular colors followed each other during exposure. The order of the trials was randomized, and whether the base pair or foil appeared first was counterbalanced across trials.

At the test phase, we also examined whether participants inferred the transitive pair from base pairs (e.g., A-C from A-B and B-C). The foil (e.g., A-F) was constructed by selecting one color from one transitive pairs (e.g., A-C), and the other from a different transitive pair (e.g., D-F from D-E and E-F) while maintaining the temporal positions in the pairs. Each transitive pair was tested against a foil. Each transitive pair and the foil were presented the same number of times at test. If participants chose the transitive pair as more familiar, this would suggest that they had automatically inferred a new association between two objects that had never directly followed each other and could only be inferred given the exposure to the two base pairs.

A debriefing session was conducted after the test phase. To assess whether or not participants explicitly noticed the pairs, they first had to answer "yes" or "no" to the question on whether or not they noticed the pairs during the first part of the experiment. If they answered "yes", they were then asked to report which objects appeared together. To count as being aware, they had to at least report one correct pair.

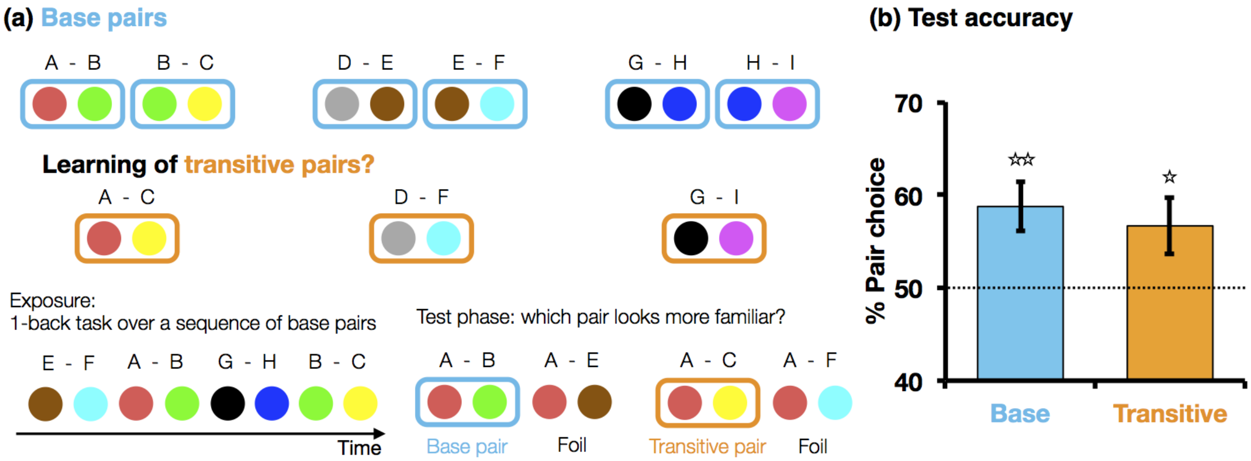

Figure 1 Stimuli and results of Experiment 1
*(a) Nine colors were paired into six base pairs (e.g., A-B, B-C), and three transitive pairs (e.g., A-C) were generated from the base pairs. To ensure incidental encoding of the pairs, participants performed a cover 1-back task over the sequence during exposure. In the test phase, participants viewed a base pair against a foil, or a transitive pair against a foil, and chose which looked more familiar. (b) Percent of times where the pair was chosen as more familiar over the foil at the test phase (error bars reflect ±1 SEM; *p<.05, **p<.01; dotted line=chance performance of 50%).*

2.1.5 Results and Discussion

At the test phase, base pairs were chosen as more familiar than foils for 58.8% (SD=16.7%) of the time, which was reliably above chance (50%) [*t*(39)=3.=32, *p*=.002, *d*=0.53] (Figure 1b). This indicates that participants have successfully learned the temporal co-occurrences between the two colors in the base pairs. Moreover, transitive pairs were chosen as more familiar than foils for 56.7% (SD=19.1%) of the time, which was again reliably above chance (50%) [*t*(39)=2.21, *p*=.03, *d*=0.35] (Figure 1b), suggesting that participants have also successfully inferred the transitive pairs, although the two colors in the transitive pair never directly followed each other during the exposure phase. There was no difference between the results of base pairs and transitive pairs [*t*(39)=0.53, *p*=.60, *d*=0.12]. However, there was no correlation [*r*(38)=.05, *p*=.75] between learning of base pairs and the inference of transitive pairs.

At debriefing, only six participants reported noticing color pairs, but none could correctly report which specific colors followed each other. This suggests that participants had no explicit awareness of the base pairs or the transitive pairs. These findings demonstrate that statistical learning automatically and implicitly forms novel associations between objects that have never appeared together and can only be associated via transitive relations based on prior experiences.

2.2 Experiment 2

This experiment aimed to examine the limits of the transitive inference by increasing the chain of object associations. Specifically, we added one more base pair (e.g., C-D) and tested whether people could infer the transitive pair (e.g., A-D).

2.2.1 Participants

A new group of 40 undergraduates (28 female, mean age=20.6 years, SD=2.5) from UBC participated in the experiment for course credit.

2.2.2 Stimuli and Procedure

The stimuli and the procedure were identical to those in Experiment 1, except that we added one more base pair such that three base pairs formed a transitive pair. As before, six base pairs were created for each participant. For every three base pairs, the second color in the first pair was the same as the first color in the second pair, and the second color in the second pair was the same as the first color in the third pair (e.g., A-B, B-C, C-D; Figure 2a). The transitive pair (e.g., A-D) consisted of the first color in the first pair and the second color in the third pair. As before, participants performed a cover 1-back task to ensure incidental encoding of the pairs during exposure. Afterwards, participants completed the surprise test phase where they chose whether the pair or the foil looked more familiar.

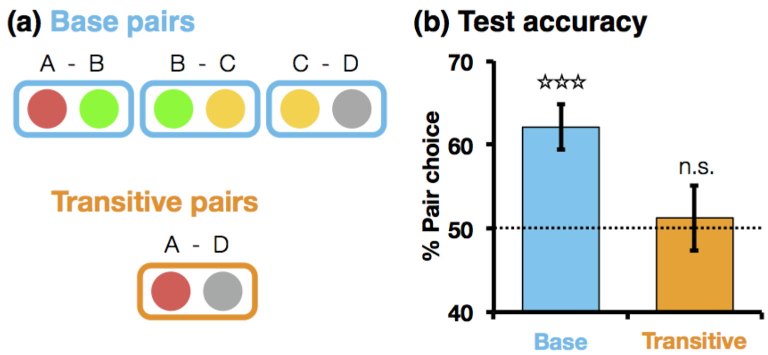

Figure 2 Stimuli and results of Experiment 2
*(a) Three example base pairs (A-B, B-C, C-D) and a transitive pair (A-D). (b) Percent of times the pair was chosen as more familiar than the foil at test (error bars reflect ±1 SEM; **p<.01; dotted line=chance performance of 50%).*

2.2.3 Results and Discussion

At the test phase, base pairs were chosen as more familiar than foils for 62.1% (SD=16.9%) of the time, which was reliably above chance (50%) [$t(39)=4.53$, $p<.001$, $d=0.72$] (Figure 2b), suggesting that participants have successfully learned the co-occurrences between the two colors in the base pairs. However, this time the transitive pairs were chosen as more familiar than foils for 51.3% (SD=24.6%) of the time, which was not reliably different from chance (50%) [$t(39)=0.32$, $p=.75$, $d=0.05$] (Figure 2b). There was a reliable difference between the results of base pairs and transitive pairs [$t(39)=2.18$, $p=.03$, $d=0.51$]. This suggests that people failed to infer the transitive pairs even though they successfully learned the base pairs. Moreover, there was no correlation [$r(38)=.12$, $p=.46$] between learning of base pairs and the inference of transitive pairs. During debriefing, only two participants reported noticing color pairs, but the participant could not correctly report which specific colors temporally followed each other. This again suggests that participants had no explicit awareness of the base pairs or the transitive pairs. This result reveals a limit in the novel associations that can be formed transitively across the base pairs. This limit may reflect processing constraints as the number of associations increases.

2.3 Experiment 3

In Experiment 2, the failure of transitive inference from A to D could be driven by the weak overlap between two base pairs (having only one shared object between pairs). Thus, this experiment aimed to overcome this failure by strengthening the extent of overlap as defined by the number of shared objects. Specifically, we maintained the same number of pair associations (A-B, B-C, and C-D) and the same number of objects (4) as in Experiment 2 (Bays, Turk-Browne, & Seitz, 2016), but increased the extent of overlap (e.g., A-B-C and B-C-D), and examined whether people could infer the transitive pair (A-D).

2.3.1 Participants

A new group of 40 undergraduates (32 female, mean age=21.6 years, SD=4.2) from UBC participated in the experiment for course credit.

2.3.2 Stimuli and Procedure

The stimuli and the procedure were identical to those in Experiment 2, except that the colored circles formed four base triplets, where the three colors within the triplet temporally appeared one after another (e.g., A-B-C). For every two base triplets, the second and the third colors in the first triplet were the same as the first and the second colors in the second triplet (e.g., A-B-C, B-C-D; Figure 3a). The transitive pair (e.g., A-D) consisted of the first color in the first triplet and the third color in the second triplet. We recognize that the number of transitive inferences was different between Experiments 2 and 3. In Experiment 2, participants needed to make two inferences (one from A-B to B-C, and another from B-C to C-D), whereas in this experiment, participants only needed to make one inference (from A-B-C to B-C-D). We will discuss this confound in the General Discussion.

As before, participants performed a cover 1-back task to ensure incidental encoding of the triplets during exposure. At the test phase, participants completed the surprise test phase where they chose whether the sub-pairs (e.g., A-B, B-C, C-D) or the foil looked more familiar, and chose whether the transitive pair or the foil looked more familiar. The sub-pairs and transitive pairs were intermixed in the test trials (order randomized).

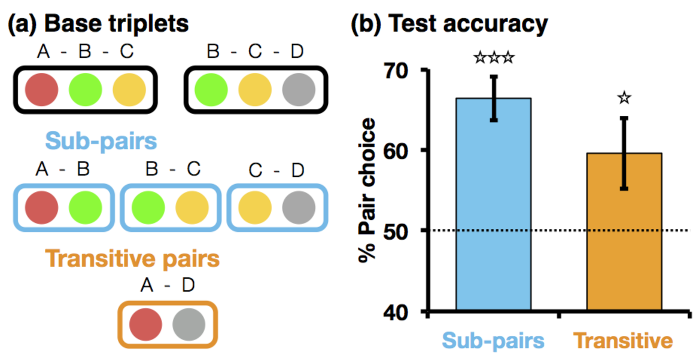

Figure 3 Stimuli and results of Experiment 3
*(a) Two example base triplets (A-B-C, B-C-D), three sub-pairs (A-B, B-C, C-D), and a transitive pair (A-D). (b) Percent of times the pair was chosen as more familiar than the foil at test (error bars reflect ±1 SEM; *p<.05, ***p<.001; dotted line=chance performance of 50%).*

2.3.3 Results and Discussion

At the test phase, the sub-pairs were chosen as more familiar than foils for 66.4% (SD=17.2%) of the time, which was reliably above chance (50%) [$t(39)=6.04$, $p<.001$, $d=0.95$] (Figure 3b), suggesting that participants have successfully learned the co-occurrences between the two colors in the sub-pairs.[1] More importantly, the transitive pairs were chosen as more familiar than foils for 59.6% (SD=28.0%) of the time, which was again reliably above chance

[1] In a separate experiment (N=30), participants chose the base triplets as more familiar than foils for 65.8% (SD=18.7%) of the time, which was significantly above chance (50%) [$t(29)=4.63$, $p<.001$, $d=0.84$], suggesting that they have successfully learned the base triplets.

(50%) [$t(39)=2.17$, $p=.04$, $d=0.34$] (Figure 3b), suggesting that participants have also successfully inferred the transitive pairs, although the two colors in the transitive pair have never directly followed each other during exposure. However, there was no reliable difference between the results of sub-pairs and transitive pairs [$t(39)=1.62$, $p=.11$, $d=0.29$]. We compared the results in Experiments 2 and 3 in a mixed-effects ANOVA, and found that there was no interaction effect between experiments and conditions [$F(1, 78)=0.38$, $p=.53$, $\eta_p^2=0.002$]. Moreover, there was a moderate correlation [$r(38)=.38$, $p=.02$] between the learning of sub-pairs and the inference of transitive pairs. During debriefing, four participants reported noticing color pairs, but no participants could correctly report which specific colors temporally followed each other. This again suggests that participants had no explicit awareness of the sub-pairs or the transitive pairs.

The findings demonstrate that the transitive inference from A to D was slightly stronger in Experiment 3, suggesting that increasing the overlap between pairs and reducing the number of transitive inferences did not completely remove the limit observed in Experiment 2.

3 Novel association across categorical levels

3.1 Experiment 4a

As shown in previous experiments, novel associations can be formed transitively between objects that have never co-occurred before. In this experiment, we aimed to extend our findings by examining another type of transitivity from set theory (Ciesielski, 1997) using hierarchical relations between objects in an ordered set. Specifically, based on exposures to the association between two objects at one categorical level (e.g., New York-London), can people infer the same association at the subordinate level (e.g., Central Park-Hyde Park) and the superordinate level (e.g., USA-UK)?

3.1.1 Participants

A new group of 80 undergraduates (58 female; mean age =20.6 years, SD=3.2) from UBC participated in the experiment for course credit. The sample size was determined by a power analysis using G*Power (Faul, et al., 2007). Given a Cohen's *d* of 0.35 which was based on the transitive inference result in Experiment 1, a minimum of 72 participants was required to have 90% power (alpha=0.05) to reveal the effect in this experiment.

3.1.2 Stimuli

The stimuli consisted of eight city names (New York, London, Vancouver, Paris, Tokyo, Beijing, Barcelona, and Bangkok), the eight corresponding park names (Central Park, Hyde Park, Stanley Park, Champ de Mars Park, Yoyogi Park, Bei Hai Park, Güell Park, and Lumpini Park), and the eight corresponding country names (USA, UK, Canada, France, Japan, China, Spain, and Thailand). The eight cities were randomly grouped into four base pairs for each participant (e.g., New York-London, Figure 4a). The city base pairs produced four park pairs at the subordinate level (e.g., Central Park-Hyde Park), and four country pairs at the superordinate level (e.g., USA-UK). The park pairs and the country pairs served as transitive pairs to be tested at the test phase, and were never presented in the exposure phase. Each city base pair was repeated 50 times to form a single continuous sequence of cities in a pseudorandom order with the constraint that no city pair could repeat back-to-back.

3.1.3 Procedure

Since participants may not know which park is in which city, they were first trained on a separate task to associate each park with a given city prior to the start of the experiment. In this task, participants viewed a park and selected which city contained the park (by pressing a key from “1” to “8”), and received feedback on each trial. They had to achieve 100% accuracy on this task before starting the experiment. We did not test city-country association, since we assumed that participants should know which city is in which country. There was no mention of any country names before starting the experiment.

The experiment consisted of an exposure phase and a test phase, as in Experiment 1. During exposure, participants performed a 1-back task over a continuous sequence of city names where they judged whether the current city name was the same as the previous one. As before, the sequence contained the city base pairs unbeknownst to the participants. One city name was presented for 500ms followed by a 500ms ISI in each trial. Each city name had a 20% chance of repeating the previous name, producing a total of 480 trials. The 1-back task served as a cover task to conceal the true purpose of the study, ensuring incidental learning of the city pairs.

After exposure, participants completed the surprise test phase as before, to see if they had learned the city pairs, and more importantly, to see if they could infer the corresponding park pairs or country pairs which were never presented during exposure. In each test trial, participants judged whether the pair or the foil looked more familiar based on what they saw in the exposure phase. There were three blocks of trials. In the first block, each city pair was tested against a foil where two cities never followed each other during exposure. The foil was constructed by selecting one city from one base pair and another city from a different base pair while maintaining the temporal positions in the pairs. In the second block, each park pair corresponding to its city pair was tested against a foil that contained the two parks corresponding to the two cities in the foil in the first block. In the third block, each country pair corresponding to its city pair was tested against a foil that contained the two countries corresponding to the two cities in the foil in the first block. The order of the last two blocks was randomized. It is important to note that in each block, the base pair or the transitive pair was presented the same number of times as the foils.

A debriefing session was conducted after the test phase, where participants were asked if they had noticed any names that appeared one after another. For those who responded yes, we further asked them to specify which names followed each other.

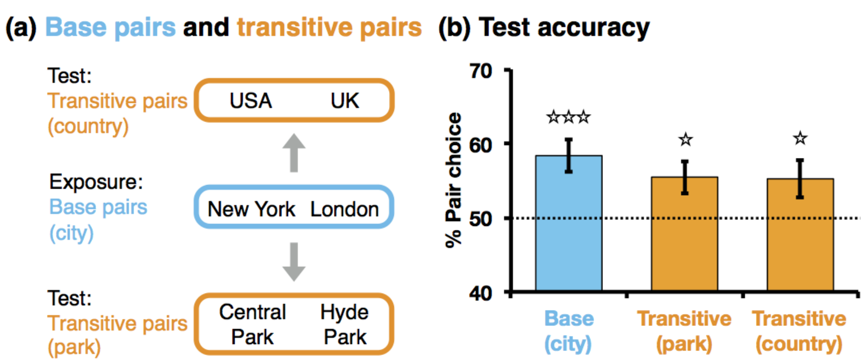

Figure 4 Stimuli and results of Experiment 4a
*(a) An example city base pairs (e.g., New York-London), the corresponding transitive park pair (e.g., Central Park-Hyde Park), and the corresponding transitive country pair (e.g., USA-UK). (b) Percent of times the pair was chosen as more familiar than the foil at test (error bars reflect ±1 SEM; *p<.05, ***p<.001; dotted line=chance performance of 50%).*

3.1.4 Results and Discussion

At the test phase, the city base pairs were chosen as more familiar than foils for 58.4% (SD=19.6%) of the time, which was reliably above chance (50%) [$t(79)=3.82$, $p<.001$, $d=0.43$] (Figure 4b). This indicates that participants have successfully learned the temporal co-occurrences between the two cities in a base pair during exposure. More importantly, park pairs were chosen as more familiar than foils for 55.5% (SD=19.6%) of the time, which was again reliably above chance (50%) [$t(79)=2.49$, $p=.01$, $d=0.28$] (Figure 4b), suggesting that participants have successfully inferred the park pairs, although no parks were ever presented during exposure. Likewise, country pairs were also chosen as more familiar than foils for 55.2% (SD=22.4%) of the time, which was again reliably above chance (50%) [$t(79)=2.09$, $p=.04$, $d=0.23$] (Figure 4b), suggesting that participants have also successfully inferred the country pairs, although no countries were presented during exposure. Moreover, there was no reliable difference in the results among the three conditions [$F(2, 158)=0.95$, $p=.39$, $\eta_p^2=0.01$]. There was a moderate correlation between the learning of city pairs and the inference of park pairs [$r(78)=.35$, $p=.002$], between city pairs and country pairs [$r(78)=.34$, $p=.002$], and between park pairs and country pairs [$r(78)=.50$, $p<.001$], further supporting that participants have successfully inferred the park pairs and the country pairs. During debriefing, seven participants reported noticing city pairs, but none correctly reported which specific names followed each other. This suggests that participants had no explicit awareness of the pairs.

These results suggest that participants spontaneously inferred new associations at both the subordinate and the superordinate levels, based on the regularities extracted at one categorical level. This provides further evidence that statistical learning forms novel associations between objects at different levels along the categorical hierarchy, even if these objects are never directly experienced or associated with each other.

3.2 Experiment 4b

Since in Experiment 4a participants were trained on the park-city associations prior to the experiment, the training could have facilitated the transitive inference from city pairs to park pairs. Moreover, the order of three blocks was not randomized in the test phase. To address the priming issue and to avoid possible order effect in the test phase, we conducted a follow-up experiment on Amazon Mechanical Turk. The goal of the current experiment was two-fold: to examine whether participants could infer the park pairs without prior training of the park-city associations, and to randomize the block order during the test phase to demonstrate that the 2AFC test performance was not driven by the fixed block order in Experiment 4a.

3.2.1 Participants

A total of 481 participants was recruited on Amazon Mechanical Turk, and received $1 for participation. Since we had no control over the testing conditions of the online participants, we screened for eligible participants based on two criteria. First, participants had to know in advance which park was in which city. This resulted in 156 participants who scored a 100% accuracy in identifying which park was in which city. Second, participants had to show successful learning of the base city pairs. This further resulted in 79 participants (38 females, mean age=33.7 years, SD=10.7) who scored above 50% for the base city pairs at the test phase. Thus, 79 participants were included in the analysis. This sample size was comparable to that in Experiment 4a (N=80).

3.2.2 Stimuli and Procedure

The stimuli and the procedure were identical to those in Experiment 4a (Figure 5a), except for two important changes. First, there was no training on the park-city association prior the experiment. To examine participants' knowledge of the location of the parks, we tested their knowledge at the end of the experiment using multiple-choice questions, where participants matched each of the eight parks to its corresponding city. Second, the order of the blocks at the test phase (i.e., one block for base city pairs and two blocks for transitive park and country pairs) was randomized in order to avoid possible order effects in Experiment 4a.

3.2.3 Results and Discussion

The 2AFC performance at the test phase was presented in Figure 5b. The base city pairs were chosen as more familiar than foils for 65.7% (SD=10.9%) of the time. This of course is not surprising because we selected participant who scored above 50% for the base city pairs. Importantly, park pairs were chosen as more familiar than foils for 53.8% (SD=14.3%) of the time, which was reliably above chance (50%) [$t(78)=2.36$, $p=.02$, $d=0.27$], suggesting that participants successfully inferred the transitive park pairs, although no park was ever presented during exposure. Likewise, country pairs were also chosen as more familiar than foils for 55.1% (SD=17.7%) of the time, which was again reliably above chance (50%) [$t(78)=2.55$, $p=.01$, $d=0.29$], suggesting that participants also successfully inferred the transitive country pairs, although no country was presented during exposure. There was no difference in performance between the park and the country conditions [$t(78)=0.54$, $p=.59$, $d=0.08$].

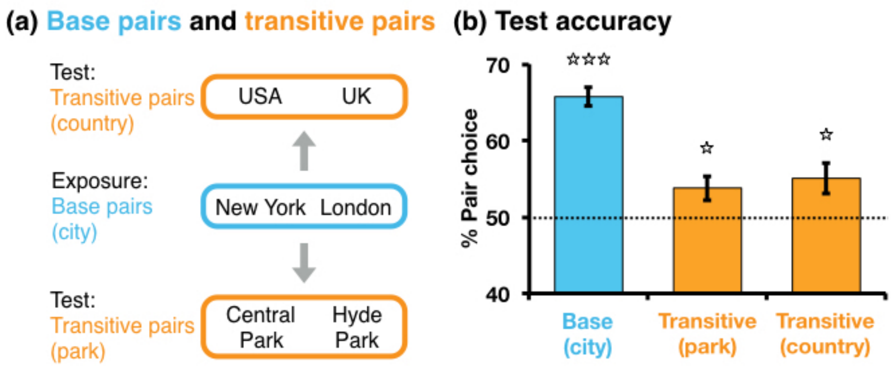

Figure 5 Stimuli and results of Experiment 4b
*(a) An example base city pair (New York-London) during exposure, the corresponding transitive park pair (Central Park-Hyde Park), and the corresponding transitive country pair (USA-UK). (b) Percent of times the pair was chosen as more familiar than the foil at test (error bars reflect ±1 SEM; *p<.05, ***p<.001; dotted line=chance performance of 50%).*

There was a weak correlation between the learning of city pairs and the inference of country pairs [$r(77)=.22$, $p<.05$], a marginal correlation between city pairs and park pairs [$r(77)=.21$, $p=.06$], but no correlation between park pairs and country pairs [$r(77)=.16$, $p=.17$]. During debriefing, three participants reported noticing park pairs, but none of them correctly reported which specific parks followed each other, suggesting that participants had no explicit awareness of the pairs. These results replicated those in Experiment 4a, suggesting that participants spontaneously inferred new associations at both the subordinate and the superordinate level, based on the regularities extracted at one categorical level. In addition, the fixed block order at the test phase in Experiment 4a did not fully account for the test performance.

3.3 Experiment 5

The goal of this experiment was to examine whether the transitive inference could only be made based on city pairs, and whether there were limits in forming the novel transitive associations across the categorical hierarchy. Specifically, park pairs served as base pairs during exposure, and city pairs and country pairs served as transitive pairs at the test phase (Figure 6a).

3.3.1 Participants

A new group of 80 undergraduates (65 female; mean age=20.3 years, SD=2.1) from UBC participated in the experiment for course credit.

3.3.2 Stimuli and Procedure

The stimuli and the procedure were identical to those in Experiment 4a, except that park pairs served as base pairs during exposure, and city pairs and country pairs served as transitive pairs at the test phase (Figure 6a).

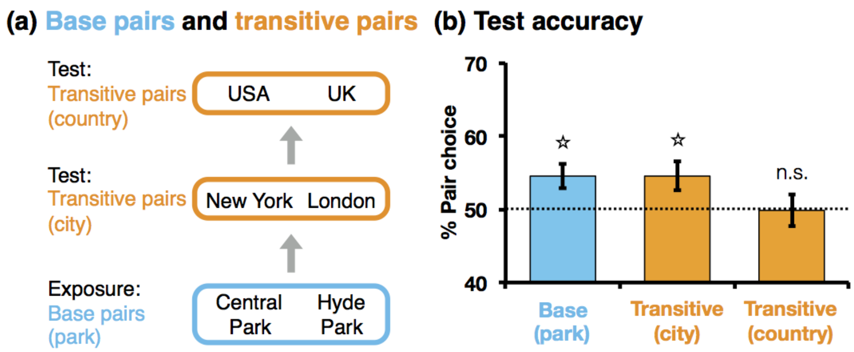

Figure 6 Stimuli and results of Experiment 5
*(a) Park pairs were presented as base pairs during exposure. City and country pairs were transitive pairs at the test phase. (b) Percent of times the pair was chosen as more familiar (error bars reflect ±1 SEM; *p<.05; dotted line=chance performance of 50%).*

3.3.3 Results and Discussion

At the test phase, the park base pairs were chosen as more familiar than foils for 54.5% (SD=15.3%) of the time, which was reliably above chance (50%) [$t(79)=2.65$, $p=.01$, $d=0.30$] (Figure 6b), indicating that participants have successfully learned the temporal co-occurrences between the two parks in a base pair during exposure. More importantly, city pairs were chosen as more familiar than foils for 54.5% (SD=17.3%) of the time, which was again reliably above chance (50%) [$t(79)=2.35$, $p=.02$, $d=0.26$] (Figure 6b), suggesting that participants have successfully inferred the city pairs, even though no cities were ever presented during exposure.

However, country pairs were chosen as more familiar than foils for 49.8% (SD=19.2%) of the time, which was not different from chance (50%) [$t(79)=0.07$, $p=.94$, $d=0.01$] (Figure 6b). There was a marginal difference among the three conditions [$F(2, 158)=2.79$, $p=.06$, $\eta_p^2=0.03$] in that the performance in the country condition was marginally weaker than that of the city or park conditions (p's=.1). There was a weak correlation between the learning of park pairs and the inference of city pairs [$r(78)=.24$, $p=.03$], and a moderate correlation between city pairs and country pairs [$r(78)=.48$, $p<.001$], but no correlation between park pairs and country pairs [$r(78)=.16$, $p=.16$], supporting that participants have successfully inferred city pairs, but not country pairs. During debriefing, three participants reported noticing park pairs, but none correctly reported which specific parks followed each other, suggesting that they had no explicit awareness of the pairs.

These results replicated those in Experiment 4a, demonstrating that participants could successfully infer new associations at the superordinate level above the original categorical level where regularities were learned. However, there was a limit in how far the inference can be made beyond the level where objects were originally associated.

3.4 Experiment 6

This experiment aimed to examine whether the limit in the transitive inference was specific to superordinate levels. Specifically, country pairs served as base pairs during exposure, and city pairs and park pairs served as transitive pairs at test (Figure 7a).

3.4.1 Participants

A new group of 80 undergraduates (61 female; mean age =20.8 years, SD=5.2) from UBC participated in the experiment for course credit.

3.4.2 Stimuli and Procedure

The stimuli and the procedure were identical to those in Experiment 4, except that country pairs served as base pairs during exposure, and city pairs and park pairs served as transitive pairs at test (Figure 7a).

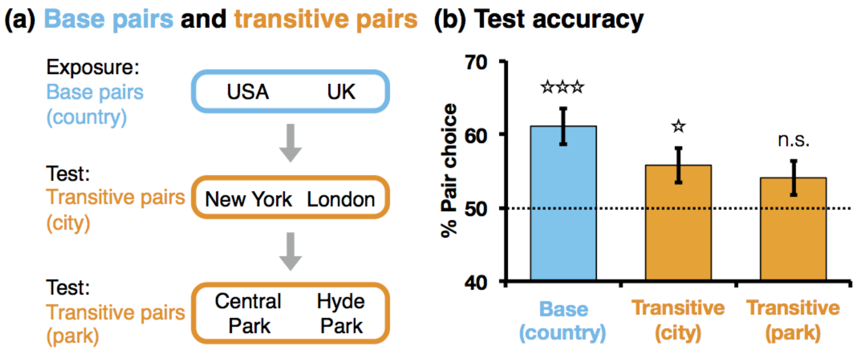

Figure 7 Stimuli and results of Experiment 6
*(a) Country pairs were presented as base pairs during exposure. City and park pairs were transitive pairs at the test phase. (b) Percent of times the pair was chosen as more familiar (error bars reflect ±1 SEM; *p<.05, ***p<.001; dotted line=chance performance of 50%).*

3.4.3 Results and Discussion

At the test phase, the country base pairs were chosen as more familiar than foils for 61.1% (SD=22.0%) of the time, which was reliably above chance (50%) [$t(79)=4.52$, $p<.001$, $d=0.51$] (Figure 7b), indicating that participants have successfully learned the temporal co-occurrences between the two countries in a base pair during exposure. More importantly, city pairs were chosen as more familiar than foils for 55.8% (SD=21.0%) of the time, which was again reliably above chance (50%) [$t(79)=2.46$, $p=.02$, $d=0.27$] (Figure 7b), suggesting that participants have successfully inferred the city pairs, even though no cities were ever presented during exposure. However, park pairs were chosen as more familiar than foils for 54.1% (SD=20.8%) of the time, which was not reliably above chance (50%) [$t(79)=1.75$, $p=.08$, $d=0.20$] (Figure 7b). Moreover, there was a significant difference among the three conditions [F(2, 158)=4.05, p=.02, η_p^2=0.05] in that the performance in the park condition was reliably weaker than that in the country condition (p=.02) but only marginally weaker than that in the city condition (p=.1). There was a moderate correlation between the learning of country pairs and the inference of city pairs, [r(78)=.30, p=.007], a strong correlation between city pairs and park pairs [r(78)=.72, p<.001], and a weak correlation between park pairs and country pairs [r(78)=.23, p=.04], supporting that participants have successfully inferred city pairs, but the inference of park pairs was weaker. During debriefing, eight participants reported noticing country pairs, but none correctly reported which specific countries followed each other, suggesting that participants had no explicit awareness of the country pairs.

These results replicated those in Experiment 4a, showing that participants could successfully infer new associations at the subordinate level below the original categorical level where regularities were learned. However, there was again a limit in how far the inference can be made beyond the level where objects were originally associated.

4 Transitive inference within and across categorical levels

4.1 Experiment 7a

Based on the previous findings, in this experiment, we examine whether seeing the A-B and B-C based city pairs (e.g., New York-London, London-Vancouver) at exposure can induce the inference of the A-C pair at both the superordinate country level (e.g., US-Canada) and the subordinate park level (e.g., Central Park-Stanley Park). In the test phase, participants viewed a pair and a foil and chose which looked more familiar. The test included the base pairs and transitive pairs at the base, subordinate, and superordinate levels for A-B, B-C, and A-C pairs.[2]

4.1.1 Participants

A new group of 100 undergraduates (80 female; mean age =20.2 years, SD=2.2) from UBC participated in the experiment for course credit. Given the demanding transitive inferences in this experiment, we raised the power to 95% in the power analysis (alpha=0.05). Given a Cohen's *d* of 0.35 in the transitive inference in Experiment 1, a minimum of 90 participants was required in each condition (a minimum of 180 for two conditions in this experiment).

4.1.2 Stimuli

The stimuli consisted of nine city names (New York, London, Vancouver, Paris, Tokyo, Beijing, Barcelona, Bangkok, and São Paulo), nine corresponding park names (Central Park, Hyde Park, Stanley Park, Champ de Mars Park, Yoyogi Park, Bei Hai Park, Güell Park, Lumpini Park, and Ibirapuera Park), and nine corresponding country names (USA, UK, Canada, France, Japan, China, Spain, Thailand, and Brazil). As before, the nine city names were randomly grouped into six city base pairs for each participant. The six city pairs contained three groups. Each group contained two city pairs (A-B and B-C) where the second city in the first pair was

[2] We thank an anonymous reviewer for suggesting this experiment.

the same as the first city in the second pair (Figure 8a). This allowed us to replicate the findings in Experiment 1 by testing whether participants could automatically infer the transitive A-C pair. There were three A-C pairs from the six base pairs.

The six base city pairs corresponded to three groups of A-B and B-C park pairs at the subordinate level and three groups of country pairs at the superordinate level (Figure 8a). This allowed us to replicate the findings in Experiment 4a by testing whether participants could infer the transitive park pairs and country pairs. More importantly, this also allowed us to test whether participants could infer the transitive A-C park pairs and country pairs based on exposure to the base A-B and B-C city pairs (Figure 8b).

4.1.3 Procedure

The exposure phase was similar to that in Experiment 4a, where each base city pair was repeated 50 times to create a single continuous sequence in a pseudorandom order with two constraints: no single city pair could repeat back-to-back and no two base pairs with a shared city name (e.g., New York-London and London-Vancouver) could be presented consecutively. Thus, the inference of the transitive pair could not be driven by non-adjacent dependencies. The city A-C pairs and all transitive park and country pairs were never presented in the exposure phase. Due to the demanding transitive inferences to be tested in this experiment, we made two changes in the exposure in order to enhance the learning of the base city pairs. First, the presentation time of each city name was increased from 500ms to 1000ms during exposure. Second, for the 1-back task at exposure, the chance of repeating the previous city name was reduced from 20% to 10% to minimize disruptions to the pair.

In the test phase, participants performed a 2AFC test where an A-B or B-C pair was presented against a foil and they chose which looked more familiar to them. There were three blocks of test trials: one block for base A-B and B-C city pairs, two blocks for the transitive A-B and B-C park and country pairs (Figure 8a). In addition, participants performed the 2AFC test where a transitive A-C pair was presented against a foil and they chose which looked more familiar to them. There were again three blocks of test trials: one block for the A-C city pairs, one block for the A-C park pairs, and one block for the A-C country pairs (Figure 8b). The order of the six blocks was randomized at the test.

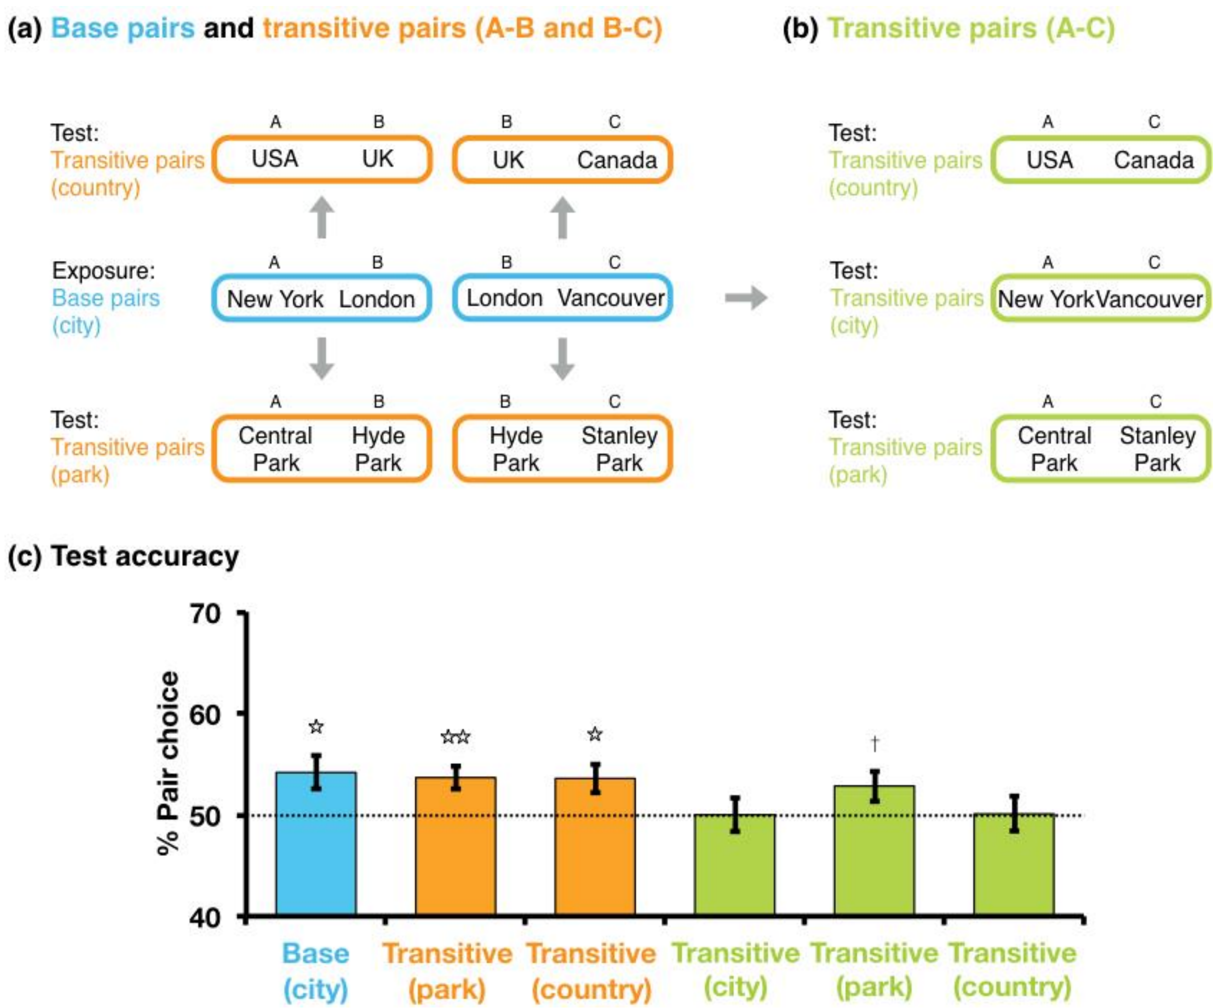

Figure 8 Stimuli and results of Experiment 7a

(a) During exposure, participants viewed a continuous sequence of base city pairs (e.g., A-B, B-C) and performed a 1-back cover task. Each base city pair corresponded to a subordinate park pair and a superordinate country pair. (b) There were three types of A-C transitive pairs: city, park, and country. (c) Participants chose whether A-B and B-C pairs or a foil looked more familiar at the city, park, or country level and chose whether A-C pair or a foil looked more familiar at the city, park, or country level.

*Percent of times the pair was chosen as more familiar than the foil at test (error bars reflect ±1 SEM; †p<.1, *p<.05, **p<.01; dotted line=chance performance of 50%).*

4.1.4 Results and Discussion

At test, participants chose the city A-B and B-C pairs as more familiar than foils for 54.2% (SD=16.0%) of the time, which was reliably above chance (50%) [$t(99)=2.60$, $p=.01$, $d=0.26$] (Figure 8c), suggesting that they successfully learned the base city pairs. They also chose the park A-B and B-C pairs as more familiar than foils for 53.7% (SD=11.1%) of the time, which was again reliably above chance (50%) [$t(99)=3.29$, $p=.001$, $d=0.33$] (Figure 8c), suggesting that they successfully inferred the transitive park pairs. Likewise, participants chose the country A-B and B-C pairs as more familiar than foils for 53.6% (SD=14.0%) of the time, which was again reliably above chance (50%) [$t(99)=2.55$, $p=.01$, $d=0.26$] (Figure 8c), suggesting that they successfully inferred the transitive country pairs. There was no correlation between the learning of city pairs and the inference of park pairs [$r(98)=.10$, $p=.31$], between city pairs and country pairs [$r(98)=.03$, $p=.79$], or between park pairs and country pairs [$r(98)=.06$, $p=.56$].

However, the city A-C pairs were chosen as more familiar than foils for 50.0% (SD=16.8%) of the time, which was not reliably different from chance (50%) [$t(99)=0.00$, $p=1$, $d=0.00$]; the park A-C pairs were chosen as more familiar than foils for 52.8% (SD=15.1%) of the time, which was marginally above chance (50%) [$t(99)=1.88$, $p=.06$, $d=0.19$], and the country A-C pairs were chosen as more familiar than foils for 50.1% (SD=16.8%) of the time, which was again not reliably different from chance (50%) [$t(99)=0.05$, $p=.96$, $d=0.005$] (Figure 8c). There was no difference in performance among all condition [$F(5, 495)=1.56$, $p=.17$, $\eta_p^2=0.02$]. This suggests that participants failed to infer the A-C pairs from the A-B and B-C pairs. This failure could be driven by the lengthy test phase which could hinder the inference of the A-C pairs.

4.2 Experiment 7b

To overcome the failure to infer A-C pairs from A-B and B-C pairs in Experiment 7a, we chose to shorten the length of the test phase by using two between-subjects conditions. In one condition participants were tested on A-B and B-C pairs, and in the other condition, participants were tested on A-C pairs.

4.2.1 Participants

A new group of 200 undergraduates (149 female; mean age=20.1 years, SD=2.5) from UBC participated in the experiment for course credit.

4.2.2 Stimuli and Procedure

The stimuli and the procedure were identical to Experiment 7a, except that in the test phase we used two between-subjects conditions to which participants were randomly assigned. In one condition (N=100), participants performed a 2AFC test where an A-B or B-C pair was presented against a foil and they chose which looked more familiar to them. There were three blocks of test trials: one block for base A-B and B-C city pairs, and two blocks for the transitive A-B and B-C park and country pairs (Figure 9a). The order of the three blocks was randomized at the test. In the other condition (N=100), participants performed the 2AFC test where a transitive A-C pair was presented against a foil and they chose which looked more familiar to them. There were again three blocks of test trials: one block for the A-C city pairs, one block for the A-C park pairs, and one block for the A-C country pairs (Figure 9b). The order of the three blocks was again randomized. The reason for having two separate testing conditions for the A-B and B-C pairs and the A-C pairs was because in Experiment 7a we combined all the testing blocks in the same experiment, where each participant completed the six blocks in a random order in the test phase. This is more than doubled the length of the test phase and thus reduced the performance at the 2AFC test overall. To keep the test phase short, we chose to use the current design where all participants completed the same exposure phase and half performed each testing condition.

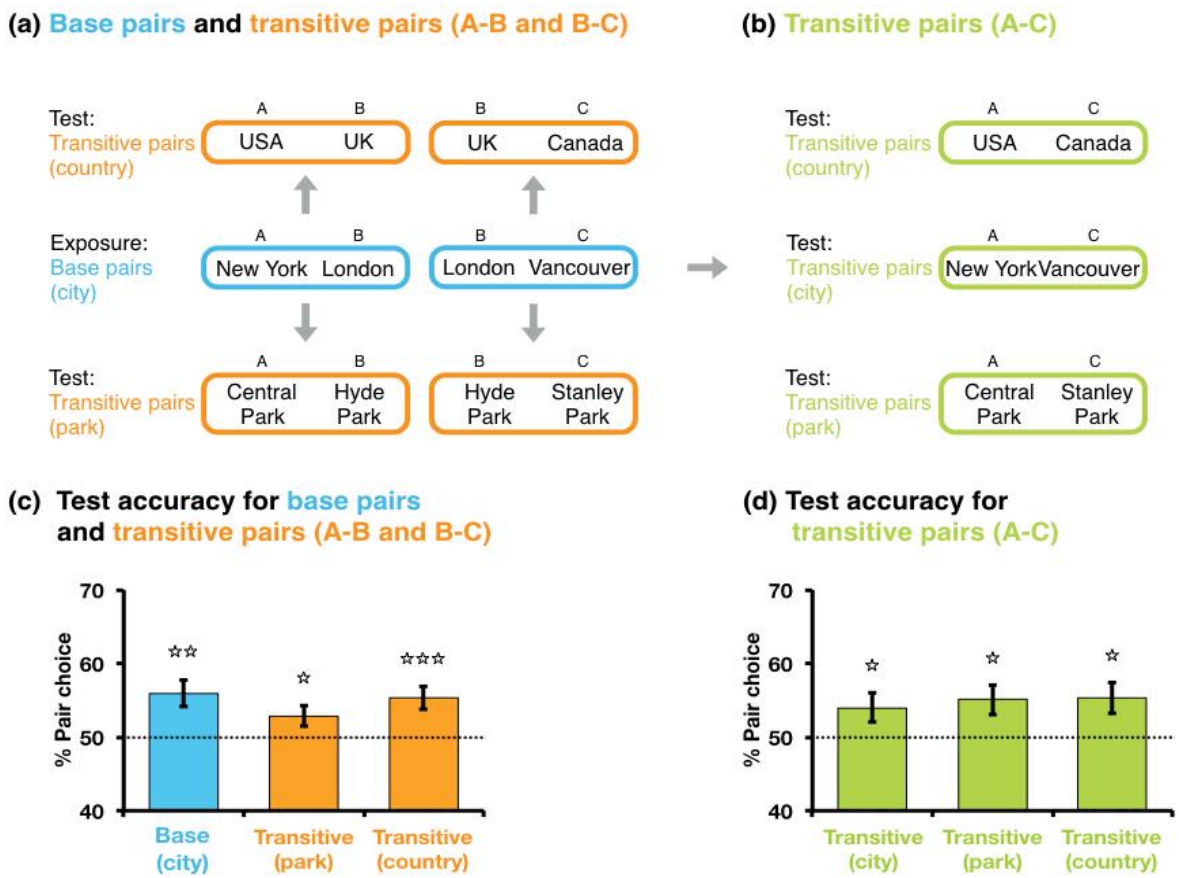

Figure 9 Stimuli and results of Experiment 7b
*(a) During exposure, participants viewed a continuous sequence of base city pairs (e.g., A-B, B-C) and performed a 1-back cover task. Each base city pair corresponded to a subordinate park pair and a superordinate country pair. (b) There were three types of A-C transitive pairs: city, park, and country. (c) In one condition (N=100), participants chose whether the A-B B-C pairs or a foil looked more familiar. The test contained base city pairs and transitive park or country pairs. The graph showed the percent of times the base or transitive pair was chosen as more familiar than the foil. (d) In another condition (N=100), participants chose whether the A-C pair or a foil looked more familiar at the city, park, or country level. The graph showed the percent of times the transitive A-C pair was chosen as more familiar than the foil (error bars reflect ±1 SEM, *p<.05, **p<.01; ***p<.001; dotted line=chance performance of 50%).*

4.2.3 Results and Discussion

At test, participants chose the base city pairs as more familiar than foils for 55.9% (SD=17.6%) of the time, which was reliably above chance (50%) [$t(99)=3.36$, $p=.001$, $d=0.33$] (Figure 9c), suggesting that they successfully learned the temporal co-occurrences between the two cities in a base pair during exposure. Participants also chose the A-B and B-C park pairs as more familiar than foils for 52.8% (SD=14.1%) of the time, which was again reliably above chance (50%) [$t(99)=2.01$, $p=.048$, $d=0.20$] (Figure 9c), suggesting that they successfully inferred the transitive park pairs, although no parks were ever presented during exposure. Likewise, participants chose the A-B and B-C country pairs as more familiar than foils for 55.3% (SD=15.6%) of the time, which was again reliably above chance (50%) [$t(99)=3.42$, $p<.001$, $d=0.34$] (Figure 9c), suggesting that they again successfully inferred the transitive country pairs, although no countries were presented during exposure. There was no difference in performance among the three conditions [$F(2, 198)=1.71$, $p=.18$, $\eta_p^2=0.02$]. There was a moderate correlation between the learning of base city pairs and the inference of transitive park pairs, [$r(98)=.40$, $p<.001$], between base city pairs and transitive country pairs [$r(98)=.36$, $p<.001$], and between transitive park pairs and country pairs [$r(98)=.39$, $p<.001$]. During debriefing, nine participants reported noticing city pairs, but none correctly reported which specific cities followed each other. This suggests that participants had no explicit awareness of the pairs.

Importantly, participants also chose the A-C city pairs as more familiar than foils for 54.0% (SD=19.6%) of the time reliably above chance (50%) [$t(99)=2.05$, $p=.04$, $d=0.20$]; they chose the A-C park pairs as more familiar than foils for 55.1% (SD=20.4%) of the time again reliably above chance (50%) [$t(99)=2.48$, $p=.01$, $d=0.25$]; and they chose the A-C country pairs as more familiar than foils for 55.3% (SD=21.4%) of the time again reliably above chance (50%)

[t(99)=2.47, p=.02, d=0.25] (Figure 9d). This suggests that participants successfully inferred the transitive A-C pairs at the base, subordinate, and superordinate levels, even though they only saw the A-B and B-C base pairs at exposure. There was no difference in performance among the three conditions [F(2, 198)=0.16, p=.85, η_p^2=0.002]. There was a moderate correlation between the inference of city A-C pairs and the inference of park A-C pairs, [r(98)=.32, p<.001], between city A-C pairs and country A-C pairs [r(98)=.32, p=.001], and a weak correlation between park A-C pairs and country A-C pairs [r(98)=.22, p=.03]. During debriefing, 13 participants reported noticing city pairs, but none correctly reported which specific cities followed each other. This suggests that participants had no explicit awareness of the pairs.

The experiment replicated the findings in previous experiments, where participants successfully inferred A-C pairs after only seeing A-B and B-C pairs (Experiment 1), and successfully inferred the same associations at subordinate and superordinate levels (Experiment 4). More interestingly, the experiment showed that participants successfully inferred A-C pairs at subordinate and superordinate levels, although they were only exposed to A-B and B-C pairs at the base level.

To further examine whether there were differences between participants who reported noticing the pairs and those who reported not noticing the pairs, we con-ducted an analysis by pooling the participants across the seven experiments, where 39 out of the total 460 participants met the criteria of explicitly noticing the pairs. We found that these participants showed margin-ally greater learning of the base pairs (63.8%) than those who reported not noticing the pairs (58.2%), t(458)=1.82, p=.07, d=0.31. However, there was no difference in the inference of transitive pairs between participants who reported noticing the base pairs (56.3%) and those who did not (54.4%), t(458)=0.58, p=.56, d=0.09. This suggests that although participants who noticed the base pairs showed marginally greater learning of the base pairs, they did not show a stronger inference of the transitive pairs. Thus, the transitive inference observed in the experiments was largely implicit

5 General Discussion

The goal of this thesis was to examine how statistical learning enables novel associations among objects that have never been directly associated before through transitive relations. We found that after learning that B followed A, and C followed B, participants automatically and implicitly inferred that C followed A, although C was never directly associated with A (Experiment 1). However, when there were three pairs (e.g., A-B, B-C, and C-D) this transitive inference (e.g., A-D) was not successful, revealing a limit in the extent of the transitive inference afforded by statistical learning (Experiment 2). This limit seemed to be partly driven by the extent of pair overlap or the number of transitive inferences required (Experiment 3). Moreover, the findings were largely supported by the correlation between learning of base pairs and the inference of transitive pairs in all experiments except Experiment 1.

Extending beyond temporal transitivity, we further examined if novel associations could be formed across the categorical hierarchy. We found that after learning a pair of objects at one categorical level (e.g., New York-London), participants implicitly inferred the same association at the subordinate level (e.g., Central Park-Hyde Park) and superordinate level (e.g., USA-UK), even if the subordinate or superordinate objects were never presented or associated with each other (Experiment 4a). Although participants were previously trained on the park-city associations before starting the experiment, they were never trained on the city and country associations. Moreover, in Experiment 4b, participants were not trained on the park-city association before the experiment. For those who already knew the park-city associations and learned the city base pairs, their transitive inference was successful at both the subordinate and superordinate levels. This suggests that the implicit inference at a different categorical level was not entirely driven by the priming of the park names, although it was possible that in

Experiments 4a, 5, 6, and 7 the transitive inference from city to park could be facilitated by the training session.

Remarkably, in Experiment 7b we found that exposure to A-B and B-C based city pairs (e.g., New York-London, London-Vancouver) induced the transitive inference of the A-C pairs at both the superordinate country level (e.g., US-Canada) and the subordinate park level (e.g., Central Park-Stanley Park).

The current findings suggest that statistical learning generates new associations beyond the statistical relationships between objects that are directly associated. Learning base pairs supports the inference of the same regularities across different objects and categorical levels. Interestingly, both the base pairs and the transitive pairs remained largely implicit, since no participant could accurately report which objects co-occurred in the experiments. This finding is consistent with previous work showing that people can infer relational information between objects without explicit awareness (Greene et al., 2001; Munnelly & Dymond, 2014).

It is important to understand the limits in this inference. The failure to infer the transitive pair (A-D) in Experiment 2 despite of successful learning of the base pairs (A-B, B-C, and C-D) could be explained by the weak overlap between pairs (i.e., A-B and B-C only shared one common object), or the greater number of transitive inference to be made. The current Experiment 3 increased the overlap but also reduced the number of inferences, and thus could not tease these two factors apart. However, follow-up studies can hold the overlap constant while increasing the number of inferences (e.g., A-B-C, B-C-D, and C-D-E). In Experiment 3, the limit was slightly alleviated but not completely removed.

The failure to infer transitive pairs at the subordinate or superordinate levels in Experiments 5 and 6 revealed a limit in transitive inference across the categorical hierarchy. This limit can be explained by several factors. The first factor was the weaker activation between countries and parks than between countries and cities. That is, a given country name may readily elicit its prominent city but not the park within the city. Second, the knowledge of which park is in which city may be weak to start with. Indeed, in the Experiment 4b only 32% of the participants knew the locations of the parks. Third, even after the initial training session participants may not fully retain the relationship between cities and parks in memory during exposure or at test, which could be a barrier to the inference from city pairs to park pairs.

A potential model that can be accounted for inferring transitive pairs after exposure is spreading activation model. This model suggests that information is encoded into units which form an interconnected network in our memory, and by activating one of the units during retrieval triggers a spread of activation in the network (Anderson, 1983). Moreover, past studies have shown that activation spreads not only a single step between directly associated objects, but also spreads in multiple steps beyond directly associated objects in the memory network (e.g., Balota & Lorch, 1986; Sharifian & Ramin, 1997). In Experiment 1, we found that after learning A-B and B-C, people can infer A-C. Given that A and B are directly associated and B-C are directly associated in the exposure phase, these two pairs are encoded into two units in the network by grouping individual information (Zhao & Yu, 2016). In the test phase, when A is presented, it automatically activates B, and then weakly activates C. When C is presented, it automatically activates B, and then weakly activates A. This bidirectional spreading of activation in the network results in choosing A-C pair as more familiar. However, when seeing a foil (e.g., A-F), A automatically activates B, and then weakly activates C, and F automatically activates E, and then weakly activates D. Given that there is no connection between A and F in the network, A-F pair was not chosen as more familiar. Thus, B is used as a mediated node to spread activation from A to C or C to A.

An unexplored question is the long-term retention of transitive inferences. In our current experiments, the test phase immediately followed the exposure phase, and therefore we do not know whether the inferences can be retained after a delay. Future studies should test the longevity of transitive inferences to further elucidate the memory strength of these newly inferred associations (Kim, Seitz, Feenstra, & Shams, 2009).

This thesis only explored how inferences are formed after learning temporal associations between objects, so future studies can train people on other associations and examine whether they can form novel inferences. Given that statistical regularities can be learned in a spatial context (Fiser & Aslin, 2001), we can test whether people can make new inferences spatially. For example, after learning that A is to the left of B and B is to the left of C, can people infer that A is likely to the left of C? Beyond the containment relations in Experiment 4, other hierarchies, such as inheritance relations, can be tested to generalize the findings in Experiment 4. For example, when the Mammal-Car pair is learned, can people infer the Chipmunk-Minivan pair at the subordinate level and the Animal-Vehicle pair at the superordinate level? To further explore the extent of our ability to make novel inferences and generalize findings in Experiment 4, cross-level pairs can be examined in future studies. Specifically, we can show one city base pair (New York-London) and one country base pair (UK-Canada) in the exposure phase and test whether people can infer the cross-level A-C transitive pair (New York-Canada).

Another interesting question to be explored is whether transitive inference occurs as early as in the exposure phase or occurs at the test phase. Participants' eye movement can be tracked during the experiment to test this question. The paradigm will be similar to Experiment 1, except that all nine colors will be shown randomly in a ring surrounding the central stimuli on the screen while participants' eye movements were tracked. For example, when A is presented at the center of the screen, we can measure whether participants tend to fixate and dwell more on A, B, and C in the ring, given that A is associated to B and B is associated to C. Also, when C is presented at the center of the screen, we can measure whether participants tend to fixate and dwell more on C, B, and A. If we see this looking bias in the exposure phase, this will suggest that transitive inference already occurs in the exposure phase which might facilitate learning of the regularities. If we only see this looking bias in the test phase, this will suggest that novel transitive inference only occurs when it is requested.

Although the majority of the participants reported that they did not notice the pairs and those who noticed cannot identify the pairs, the alternative choice task in the test phase requires explicit decisions and is considered as an explicit measure to a certain extent. To avoid using an explicit measure, future studies can use a visual search paradigm as an implicit measure to test learning and inference in the test phase. For example, we can show a red dot before the searching task as a cue. If participants learned the regularities (red co-occurs with blue and blue co-occurs with yellow) and could make novel inferences, they should be faster at finding a yellow car than a green car among other colored objects.

This thesis is significant in several ways. First, we found that people are able to automatically infer novel associations through transitive relations between objects that have never appeared together before. Second, this thesis extends beyond past work showing that people can learn categorical regularities from associations among individual exemplars (Brady & Oliva, 2008). We demonstrated that the regularities extracted at one categorical level can be inferred at the subordinate or superordinate level. This suggests that statistical learning not only operates at an abstract conceptual level, but also propagates object associations across the categorical hierarchy. Third, we revealed the limits in these transitive inferences. Fourth, in all experiments, participants reported that they did not noticed any pairs, suggesting that the transitive inference of the novel associations does not require conscious awareness of the regularities that are previously learned. Understanding the scope and the limits of the transitive associations in statistical learning can help reveal how the cognitive system generates inferences from prior experiences.

Bibliography

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of verbal learning and verbal behavior, 22*, 261-295.

Bays, B. C., Turk-Browne, N. B., & Seitz, A. R. (2016). Dissociable behavioural outcomes of visual statistical learning. *Visual Cognition,* 1-26.

Balota, D. A., & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *Journal of Experimental Psychology: Learning, memory, and cognition, 12*, 336.

Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General, 138*, 487-502.

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

Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature, 232, 456-458.*

Ciesielski, K. (1997*). Set theory for the working mathematician (Vol. 39).* Cambridge University Press.

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

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39,* 175-191.

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

Gillan, D. J. (1981). Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 150-164.

Goodwin, G. P., & Johnson-Laird, P. N. (2005). Reasoning about relations. P*sychological Review*, *112*, 468.

Greene A.J., Spellman B.A., Dusek J.A., Eichenbaum H.B., & Levy W.B. (2001). Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans. *Memory & Cognition*, *29,* 893–902.

Davis, H. (1992). Transitive inference in rats (rattus norvegicus). *Journal of Comparative Psychology*, *106*, 342-349.

Kim, R., Seitz, A., Feenstra, H., & Shams, L. (2009). Testing assumptions of statistical learning: Is it long-term and implicit?. *Neuroscience Letters, 461,* 145-149.

Kumaran, D., & Ludwig, H. (2013). Transitivity performance, relational hierarchy knowledge and awareness: Results of an instructional framing manipulation. *Hippocampus, 23,* 1259-1268.

McGonigle, B. O., & Chalmers, M. (1977). Are monkeys logical? *Nature*, *267,* 694-696.

Munnelly, A., & Dymond, S. (2014). Relational memory generalization and integration in a transitive inference task with and without instructed awareness. *Neurobiology of Learning and Memory, 109,* 169-177.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9,* 353–383.

Piaget, J. (1928). *Judgment and reasoning in the child.* Oxford, England: Harcourt, Brace.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926-1928.

Sharifian, F., & Samani, R. (1997). Hierarchical spreading of activation. *In Proc. of the Int'l Conference on Language, Cognition, and Interpretation* (pp. 1-10). IAU Press.

Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 399–407.

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General, 134,* 552–564.

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,* 1934–1945.

Von Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 17,* 334-341.

Wimmer, G., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science, 338,* 270-273.

Yu, R., & Zhao, J. (2015). The persistence of attentional bias to regularities in a changing environment. *Attention, Perception, & Psychophysics, 77,* 2217-2228.

Zeithamova, D., Schlichting, M. L., & Preston, A. R. (2012). The hippocampus and inferential reasoning: building memories to navigate future decisions. *Frontiers in human neuroscience, 6,* 70.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science, 24*, 667–677.

Zhao, J., & Yu, R. (2016). Statistical regularities reduce perceived numerosity. *Cognition, 146*, 217-222.

Zhao, J. & Luo, Y. (2017). Statistical regularities guide the spatial scale of attention. *Attention, Perception, & Psychophysics, 79*, 24-30.