Research Article

Psychological Science
2018, Vol. 29(8) 1207–1220

DOI: 10.1177/0956797618762400
www.psychologicalscience.org/PS

Statistical Learning Creates Novel Object Associations via Transitive Relations

Yu Luo[1] and Jiaying Zhao[1,2]
[1]Department of Psychology, The University of British Columbia, and [2]Institute for Resources, Environment and Sustainability, The University of British Columbia

Abstract

A remarkable ability of the cognitive system is to make novel inferences on the basis of prior experiences. What mechanism supports such inferences? We propose that statistical learning is a process through which 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 in which new transitive associations between objects are implicitly inferred.

Keywords

statistical learning, transitive inference, implicit associations;, regularities;, categorical hierarchy, open data, open materials

Received 10/22/16; Revision accepted 1/23/18

A remarkable ability of the human cognitive system is that it is able to make new inferences on the basis of prior experiences. What cognitive mechanisms support such inferences? One mechanism is logical reasoning. For example, if Ann is taller than Beth, and Beth is taller than Cath, then people can infer that Ann is taller than Cath through transitive relations (Goodwin & Johnson-Laird, 2005). Such transitive inference often requires deliberate and conscious reasoning. However, other studies have demonstrated that transitive relations can also be formed without awareness. For example, people can successfully learn the hierarchical order of objects (e.g., if A < B and B < C, then A < C) without awareness using trial-by-trial feedback (Greene, Spellman, Dusek, Eichenbaum, & Levy, 2001; Kumaran & Ludwig, 2013). A recent study shows that the hippocampus supports the transfer of values across objects that were previously associated, enabling people to make decisions between options that were never directly rewarded (Wimmer & Shohamy, 2012). Given the possibility of reactivating previous connections between objects, we propose that statistical learning is a process that establishes connections between object representations, which allows transitive inferences to form on the basis of prior experiences.

Statistical learning is the extraction of regularities between individual objects in terms of how they co-occur over space or time (Fiser & Aslin, 2001; Saffran, Aslin, & Newport, 1996). 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 2 min (Saffran et al., 1996). In addition to the auditory domain, statistical learning can operate in multiple sensory modalities and feature dimensions (Conway & Christiansen, 2005; Fiser & Aslin, 2001; Saffran et al., 1996; Turk-Browne, Isola,

Corresponding Author:
Jiaying Zhao, Department of Psychology and Institute for Resources, Environment and Sustainability, The University of British Columbia, Vancouver, British Columbia, V6T 1Z4, Canada
E-mail: jiayingz@psych.ubc.ca

Scholl, & Treat, 2008), spontaneously draw attention to the co-occurring objects (Yu & Zhao, 2015; Zhao, Al-Aidroos, & Turk-Browne, 2013; Zhao & Luo, 2017), and facilitate the compression of redundant information in the environment (Brady, Konkle, & Alvarez, 2009; Zhao & Yu, 2016).

A distinct feature of statistical learning is that this process occurs incidentally and automatically, without conscious intent or explicit awareness, because 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 be associated only via transitive relations.

The goal of our current study 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 our experiments, we asked a series of questions. First, given pairs A–B and B–C, can people automatically infer a new pair A–C (Experiments 1–3)? Second, given exemplar pairs at one categorical level (e.g., New York–London), can people infer new pairs at the subordinate level (e.g., Central Park–Hyde Park) and the superordinate level (e.g., USA–UK; Experiments 4–6)? And finally, 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)?

Experiment 1

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

Method

Participants. Forty undergraduates (28 female; age: M = 20.4 years, SD = 2.3) from The University of British Columbia (UBC) participated for course credit.[1] 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 software (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 (α = .05) to reveal the effect in our experiments.

Apparatus. Participants in all experiments were seated 60 cm from a computer monitor (refresh rate = 60 Hz). Stimuli were presented using MATLAB (The MathWorks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

Stimuli. The stimuli consisted of nine circles in nine distinct colors (red: RGB = 255, 0, 0; green: RGB = 0, 255, 0; blue: RGB = 0, 0, 255; yellow: RGB = 255, 255, 0; magenta: RGB = 255, 0, 255; cyan: RGB = 0, 255, 255; gray: RGB = 185, 185, 185; brown: RGB = 103, 29, 0; black: RGB = 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; Fig. 1a). This allowed us to test whether 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 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 nonadjacent dependencies.

Procedure. The experiment contained two phases: the exposure phase and the test phase. During the exposure phase, one colored circle appeared at the center of the screen for 500 ms, followed by a 500-ms interstimulus interval (ISI) in each trial. Participants performed a 1-back task in which 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 was 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 a 20% chance of repeating itself (e.g., AAB or ABB). This 1-back task served as a cover task irrelevant to learning in order to conceal the true purpose of the experiment, 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 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 500 ms followed by a

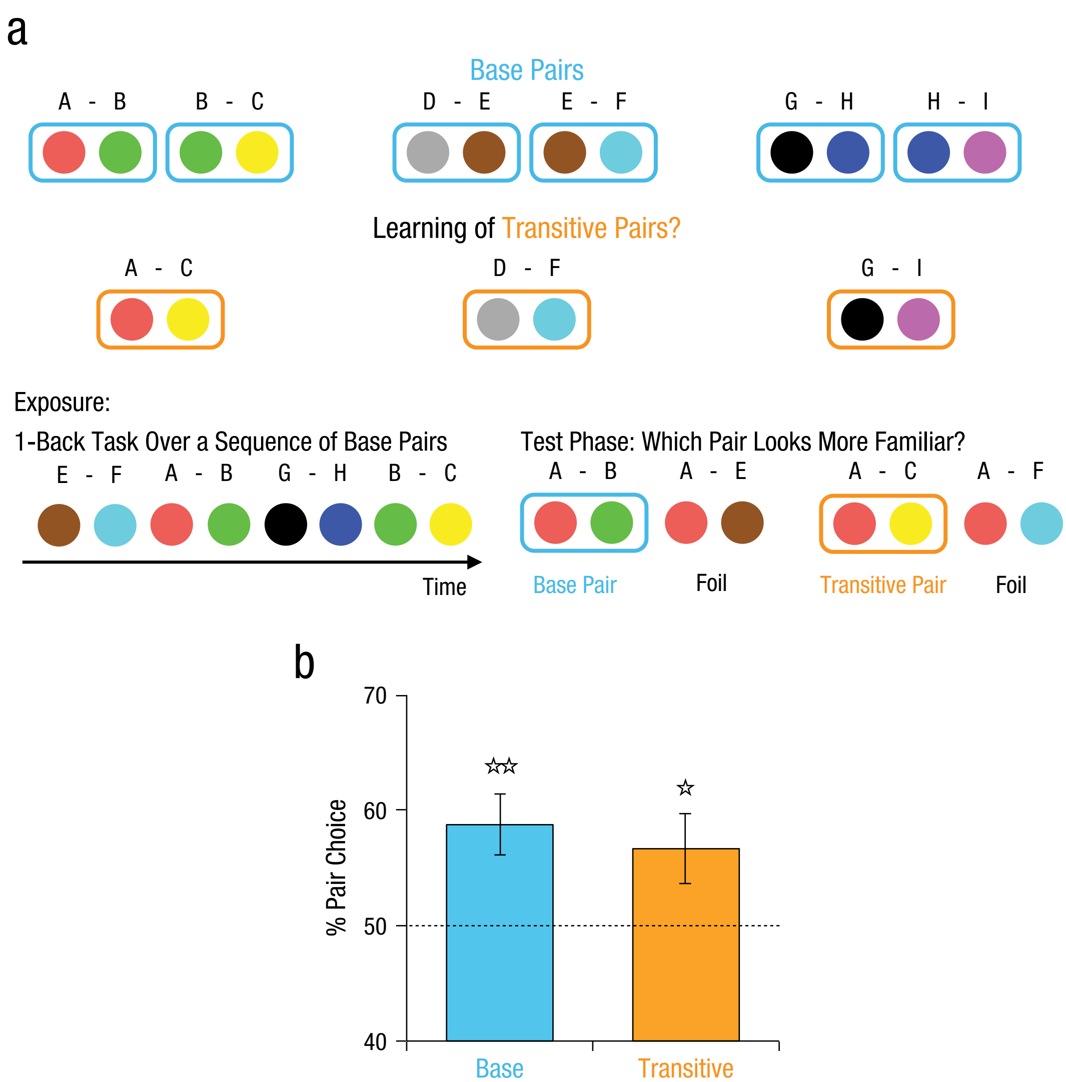

Fig. 1. Paradigm and results of Experiment 1. Nine colors were paired into six base pairs, and three transitive pairs were generated from the base pairs; examples are shown in (a). To ensure incidental encoding of the pairs, we asked participants to perform a 1-back task over the sequence during exposure. In the test phase, participants viewed a base pair against a foil or a transitive pair against a foil and chose which pair looked more familiar. The graph (b) shows the mean percentage of times the base and the transitive pairs were chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance (*$p < .05$, **$p < .01$). The dotted line indicates the level of chance performance.

500-ms ISI, and each sequence was separated by a 1,000-ms pause. Participants judged whether the first or second sequence looked more familiar from 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 another 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 is 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 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 pair (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 be inferred given only 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, we first asked them to answer "yes" or "no" to the question whether 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 report at least one correct pair.

Results

At the test phase, base pairs were chosen as more familiar than foils 58.8% (SD = 16.7%) of the time, which was reliably above chance (50%), $t(39) = 3.32$, $p = .002$, $d = 0.53$ (Fig. 1b). This indicates that participants successfully learned the temporal co-occurrences between the two colors in the base pairs. Moreover, transitive pairs were chosen as more familiar than foils 56.7% (SD = 19.1%) of the time, which was again reliably above chance (50%), $t(39) = 2.21$, $p = .03$, $d = 0.35$ (Fig. 1b), suggesting that participants 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 for 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 6 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 be associated only via transitive relations on the basis of prior experiences.

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).

Method

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

Stimuli and procedure. The stimuli and the procedure were identical to those in Experiment 1, except that we added one more base pair, so 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; Fig. 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 1-back task as a cover to ensure incidental encoding of the pairs during exposure. Afterward, participants completed the surprise test phase, in which they chose whether the pair or the foil looked more familiar.

Results

At the test phase, base pairs were chosen as more familiar than foils 62.1% (SD = 16.9%) of the time, which was reliably above chance (50%), $t(39) = 4.53$, $p < .001$, $d = 0.72$ (Fig. 2b), suggesting that participants 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 51.3% (SD = 24.6%) of the time, which was not reliably different from chance (50%), $t(39) = 0.32$, $p = .75$, $d = 0.05$ (Fig. 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 2 participants reported noticing color pairs, but the participants 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.

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

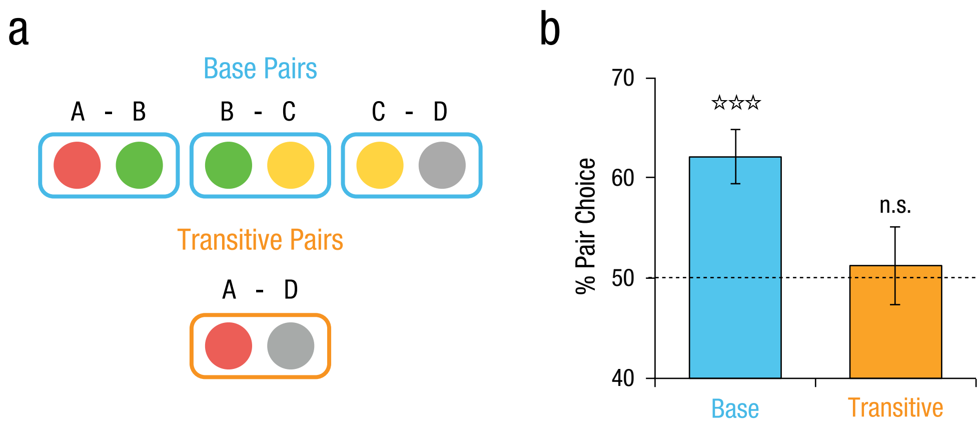

Fig. 2. Stimuli and results of Experiment 2. Examples of three base pairs and a transitive pair are shown in (a). The graph (b) shows the mean percentage of times the base and the transitive pairs were chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance ($p < .001$). The dotted line indicates the level of chance performance.

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 (four) as in Experiment 2 (Bays, Turk-Browne, & Seitz, 2015); however, we increased the extent of overlap (e.g., A–B–C and B–C–D) and examined whether participants could infer the transitive pair (A–D).

Method

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

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; Fig. 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 needed to make only one inference (from A–B–C to B–C–D). We will address this confound in the General Discussion.

As before, participants performed a 1-back task as a cover to ensure incidental encoding of the triplets during exposure. At the test phase, participants completed the surprise test phase, where they chose whether the subpairs (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 subpairs and transitive pairs were intermixed in the test trials (order was randomized).

Results

At the test phase, the subpairs were chosen as more familiar than foils 66.4% (SD = 17.2%) of the time, which was reliably above chance (50%), $t(39) = 6.04$, $p < .001$, $d = 0.95$ (Fig. 3b), suggesting that participants successfully learned the co-occurrences between the two colors in the subpairs.[2] More importantly, the transitive pairs were chosen as more familiar than foils 59.6% (SD = 28.0%) of the time, which was again reliably above chance (50%), $t(39) = 2.17$, $p = .04$, $d = 0.34$ (Fig. 3b), suggesting that participants also successfully inferred the transitive pairs, although the two colors in the transitive pair never directly followed each other during exposure. However, there was no reliable difference between the results for subpairs and transitive pairs, $t(39) = 1.62$, $p = .11$, $d = 0.29$. We compared the results of Experiments 2 and 3 in a mixed-effects analysis of variance and found that there was no interaction effect between experiments and conditions, $F(1, 78) = 0.38$, $p = .53$, $\eta_p^2 = .002$. Moreover, there was a moderate correlation, $r(38) = .38$, $p = .02$, between the learning of subpairs and the inference of transitive pairs. During debriefing, 4 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 subpairs or the transitive pairs.

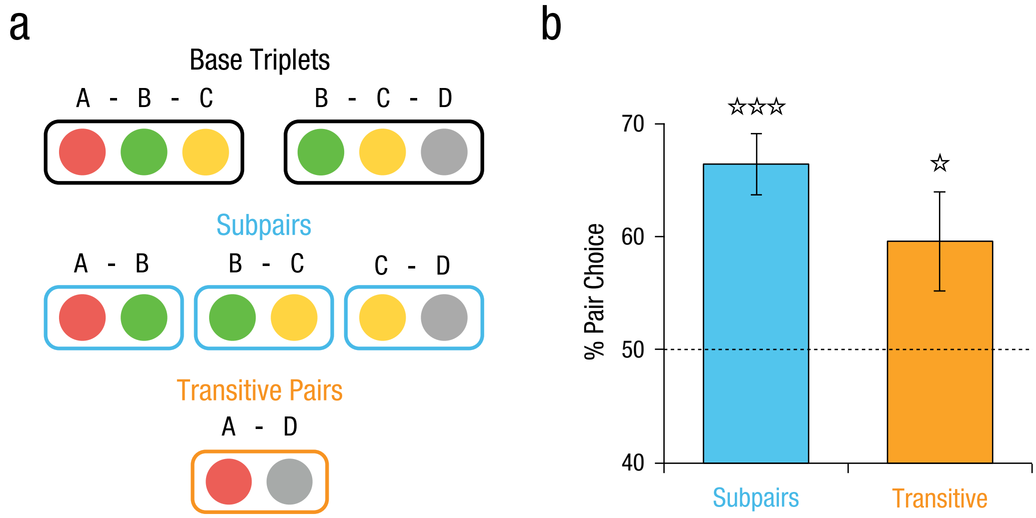

Fig. 3. Stimuli and results of Experiment 3. Examples of two base triplets, three subpairs, and a transitive pair are shown in (a). The graph (b) shows the mean percentage of times the subpairs and the transitive pairs were chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance (**p* < .05, ****p* < .001). The dotted line indicates the level of chance performance.

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.

Experiment 4

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). Specifically, given 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)?

Method

Participants. A new group of 80 undergraduates (58 female; age: M = 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 software (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 (α = .05) to reveal the effect in this experiment.

Stimuli. The stimuli consisted of eight city names (New York, London, Vancouver, Paris, Tokyo, Beijing, Barcelona, and Bangkok), 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 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; Fig. 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.

Procedure. Because 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 because 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.

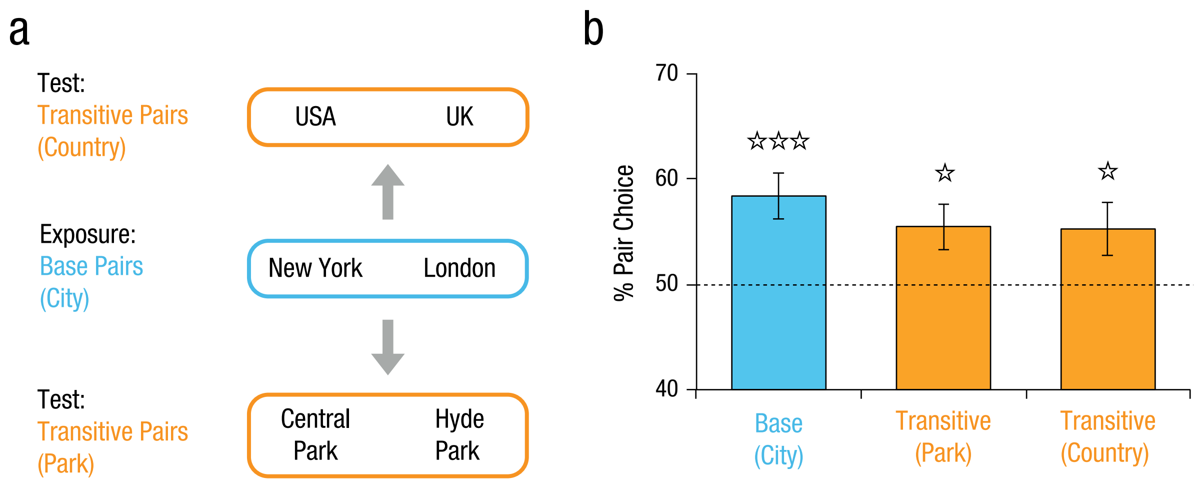

Fig. 4. Stimuli and results of Experiment 4. Examples of a city base pair, the corresponding transitive park pair, and the corresponding transitive country pair are shown in (a). The graph (b) shows the mean percentage of times each pair type was chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance (*$p < .05$, ***$p < .001$). The dotted line indicates the level of chance performance.

As before, the sequence contained the city base pairs unbeknownst to the participants. One city name was presented for 500 ms followed by a 500-ms ISI in each trial. Each city name had a 20% chance of repeating the city name on the previous trial, producing a total of 480 trials. The 1-back task served as a cover task to conceal the true purpose of the experiment, ensuring incidental learning of the city pairs.

After exposure, participants completed the surprise test phase as before, to see whether they had learned the city pairs and, more importantly, to see whether 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 given 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, in which 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.

Results

At the test phase, the city base pairs were chosen as more familiar than foils 58.4% (SD = 19.6%) of the time, which was reliably above chance (50%), $t(79) = 3.82$, $p < .001$, $d = 0.43$ (Fig. 4b). This indicates that participants 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 55.5% (SD = 19.6%) of the time, which was again reliably above chance (50%), $t(79) = 2.49$, $p = .01$, $d = 0.28$ (Fig. 4b), suggesting that participants have successfully inferred the park pairs, although no parks were presented during exposure. Likewise, country pairs were also chosen as more familiar than foils 55.2% (SD = 22.4%) of the time, which was again reliably above chance (50%), $t(79) = 2.09$, $p = .04$, $d = 0.23$ (Fig. 4b), suggesting that participants 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, $F(2, 158) = 0.95$, $p = .39$, $\eta_p^2 = .01$. 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, $r(78) = .50$, $p < .001$; this further supports that participants successfully inferred the park pairs and the country pairs.

During debriefing, 7 participants reported noticing city pairs, but none correctly reported which specific names followed each other. This suggests that participants had

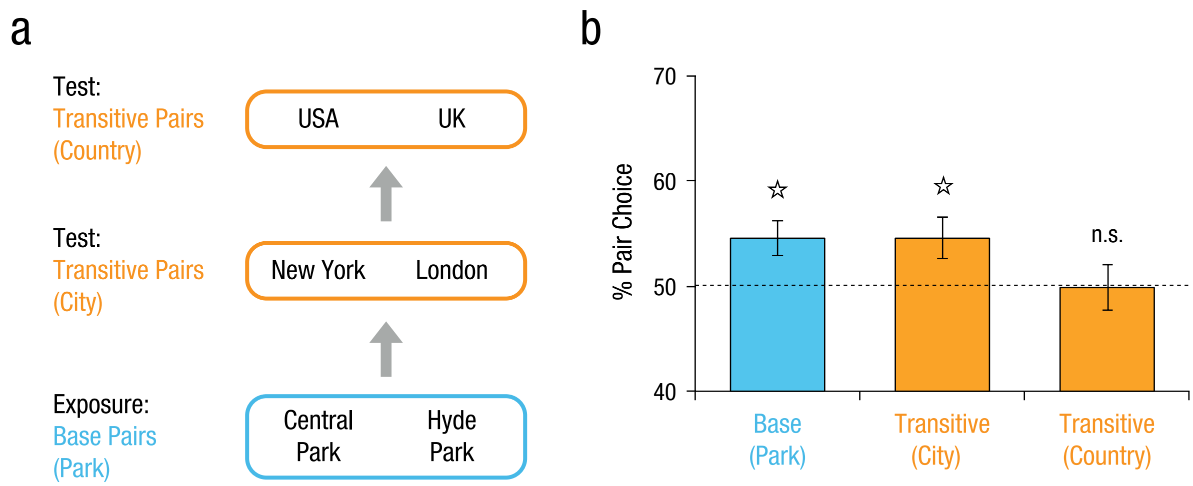

Fig. 5. Stimuli and results of Experiment 5. Park pairs were presented as base pairs during exposure, whereas city and country pairs were transitive pairs at the test phase. Examples are shown in (a). The graph (b) shows the mean percentage of times each pair type was chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance ($p < .05$). The dotted line indicates the level of chance performance.

no explicit awareness of the pairs. Because 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. To address the priming issue, we conducted a follow-up study ($N = 481$) on Amazon Mechanical Turk (MTurk; see the Supplemental Material available online), in which participants did not go through training prior to the experiment but were instead tested on the park–city associations after the experiment. Importantly, we also randomized the block order at the test phase. We found that for participants who knew which park was in which city and successfully learned the base city pairs, they automatically inferred park pairs and country pairs.

These results suggest that participants spontaneously inferred new associations at both the subordinate and the superordinate levels on the basis of 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.

Experiment 5

The goal of this experiment was to examine whether the transitive inference could be made only on the basis of 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 (Fig. 5a).

Method

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

Stimuli and procedure. The stimuli and the procedure were identical to those in Experiment 4, except that park pairs served as base pairs during exposure, and city pairs and country pairs served as transitive pairs at the test phase (Fig. 5a).

Results

At the test phase, the park base pairs were chosen as more familiar than foils 54.5% ($SD = 15.3\%$) of the time, which was reliably above chance (50%), $t(79) = 2.65$, $p = .01$, $d = 0.30$ (Fig. 5b), indicating that participants 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 54.5% ($SD = 17.3\%$) of the time, which was again reliably above chance (50%), $t(79) = 2.35$, $p = .02$, $d = 0.26$ (Fig. 5b), suggesting that participants successfully inferred the city pairs, even though no cities were presented during exposure. However, country pairs were chosen as more familiar than foils 49.8% ($SD = 19.2\%$) of the time, which was not different from chance (50%), $t(79) = 0.07$, $p = .94$, $d = 0.01$ (Fig. 5b). There was a marginal difference among the three conditions, $F(2, 158) = 2.79$, $p = .06$, $\eta_p^2 = .03$, in that the performance in the country condition was marginally weaker than that in the city or park conditions (ps = .1). There

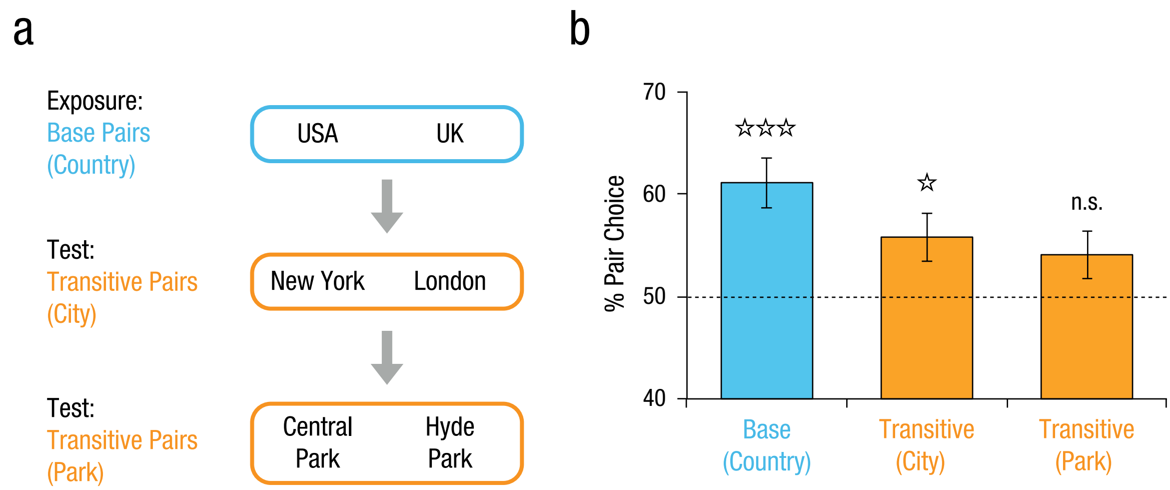

Fig. 6. Stimuli and results of Experiment 6. Country pairs were presented as base pairs during exposure, whereas city and park pairs were transitive pairs at the test phase. Examples are shown in (a). The graph (b) shows the mean percentage of times each pair type was chosen as more familiar than the foil at the test phase. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance (*$p < .05$, ***$p < .001$). The dotted line indicates the level of chance performance.

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, $r(78) = .16$, $p = .16$; this supports that participants successfully inferred city pairs but not country pairs. During debriefing, 3 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 4, demonstrating that participants could successfully infer new associations at the superordinate level above the original categorical level at which regularities were learned. However, there was a limit in how far the inference could be made beyond the level at which objects were originally associated.

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 (Fig. 6a).

Method

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

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 (Fig. 6a).

Results

At the test phase, the country base pairs were chosen as more familiar than foils 61.1% ($SD = 22.0\%$) of the time, which was reliably above chance (50%), $t(79) = 4.52$, $p < .001$, $d = 0.51$ (Fig. 6b), indicating that participants 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 55.8% ($SD = 21.0\%$) of the time, which was again reliably above chance (50%), $t(79) = 2.46$, $p = .02$, $d = 0.27$ (Fig. 6b), suggesting that participants successfully inferred the city pairs, even though no cities were presented during exposure. However, park pairs were chosen as more familiar than foils 54.1% ($SD = 20.8\%$) of the time, which was not reliably above chance (50%), $t(79) = 1.75$, $p = .08$, $d = 0.20$ (Fig. 6b). Moreover, there was a significant difference among the three conditions, $F(2, 158) = 4.05$, $p = .02$, $\eta_p^2 = .05$, in that 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, $r(78) = .23$, $p = .04$; this supports that participants successfully inferred city pairs but that the inference of park pairs was weaker. During debriefing, 8 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 4, showing that participants could successfully infer new associations at the subordinate level below the original categorical level at which regularities were learned. However, there was again a limit in how far the inference could be made beyond the level at which objects were originally associated.

Experiment 7

Given the findings in previous experiments, we next examined whether seeing A–B and B–C 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., USA–Canada) and the subordinate park level (e.g., Central Park–Stanley Park).[3]

Method

Participants. A new group of 200 undergraduates (149 female; age: $M = 20.1$ years, $SD = 2.5$) 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 ($\alpha = .05$). 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 the two conditions in this experiment).

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 the same as the first city in the second pair (Fig. 7a). 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 (Fig. 7a). This allowed us to replicate the findings in Experiment 4 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 on the basis of exposure to the base A–B and B–C city pairs (Fig. 7b).

Procedure. The exposure phase was similar to that in Experiment 4, 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 nonadjacent dependencies. The city A–C pairs and all transitive park and country pairs were never presented in the exposure phase. Because of the demanding transitive inferences to be tested in this experiment, we made two changes in the exposure in order to enhance learning of the base city pairs. First, the presentation time of each city name was increased from 500 ms to 1,000 ms 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.

There were two between-subjects conditions in the test phase to which participants were randomly assigned. In one condition ($n = 100$), participants performed a 2AFC test in which 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 (Fig. 7a). The order of the three blocks was randomized at test. In the other condition ($n = 100$), participants performed a 2AFC test in which 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 (Fig. 7b). 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 a pilot study (see the Supplemental Material) we combined all the testing blocks in the same experiment, so each participant completed the six blocks in a random order in the test phase. This more than doubled the length of the test phase and thus reduced performance in the 2AFC test overall. To keep the test phase short, we

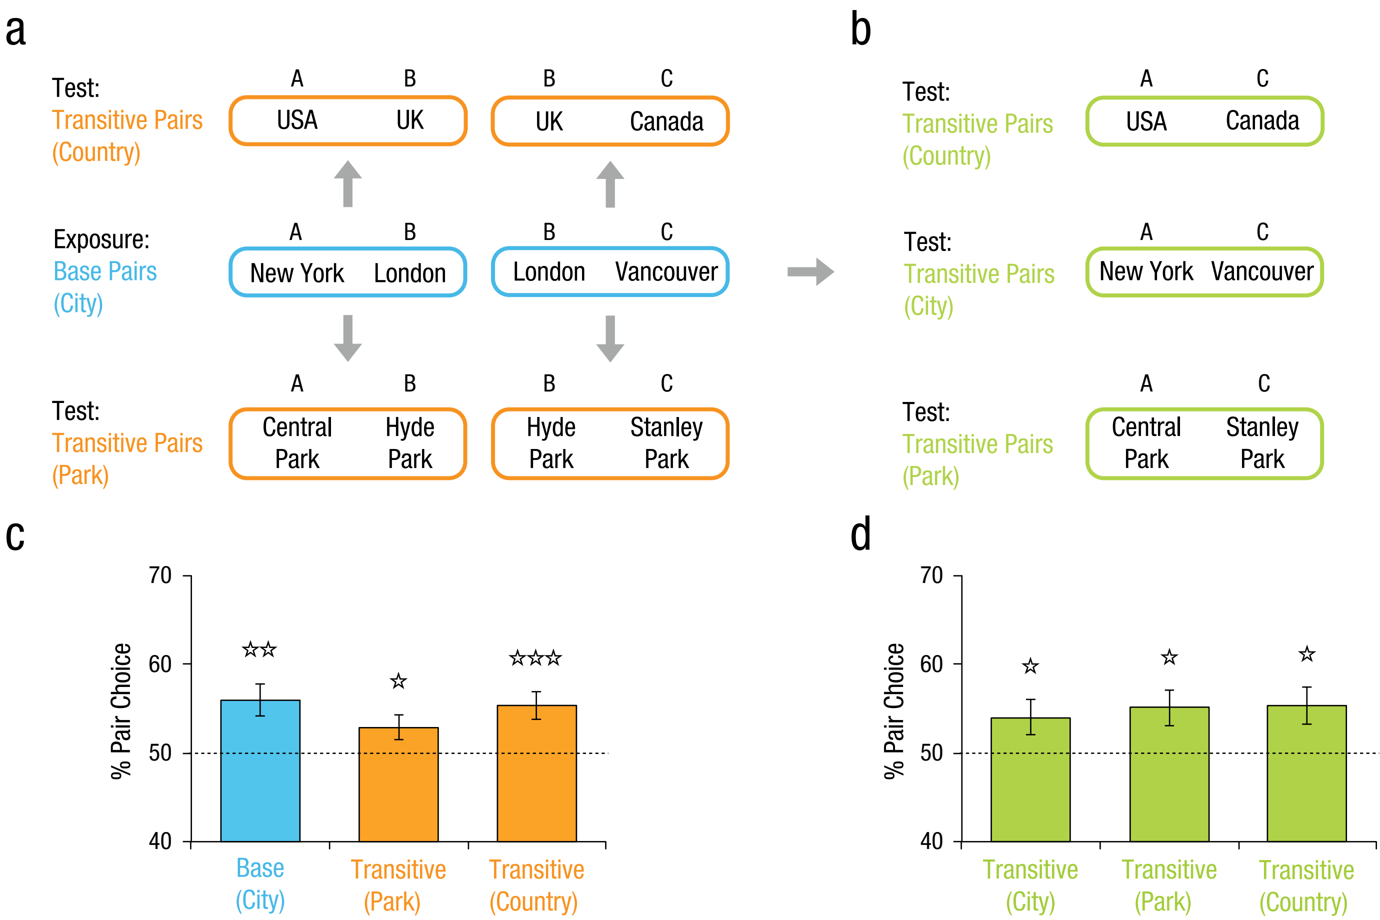

Fig. 7. Stimuli and results of Experiment 7. During exposure, participants viewed a continuous sequence of base city pairs and performed a 1-back cover task. Each base city pair corresponded to a subordinate park pair and a superordinate country pair. Examples are shown in (a). There were three types of A–C transitive pairs: city, park, and country; examples are shown in (b). In one condition (c; n = 100), participants chose whether the A–B and B–C pairs or a foil looked more familiar. The test contained base city pairs and transitive park or country pairs. The graph shows the mean percentage of times each pair type was chosen as more familiar than the foil at the test phase. In another condition (d; n = 100), participants chose whether the A–C pair or a foil looked more familiar at the city, park, or country level. The graph shows the mean percentage of times the transitive A–C pair was chosen as more familiar than the foil. Error bars reflect ±1 *SEM*. Asterisks indicate results significantly different from chance (*$p < .05$, **$p < .01$, ***$p < .001$). The dotted line indicates the level of chance performance.

chose to use the current design, in which all participants completed the same exposure phase and half performed each testing condition.

Results

At test, participants chose the base city pairs as more familiar than foils 55.9% (SD = 17.6%) of the time, which was reliably above chance (50%), $t(99) = 3.36$, $p = .001$, $d = 0.33$ (Fig. 7c), 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 52.8% (SD = 14.1%) of the time, which was again reliably above chance (50%), $t(99) = 2.01$, $p = .048$, $d = 0.20$ (Fig. 7c), 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 55.3% (SD = 15.6%) of the time, which was again reliably above chance (50%), $t(99) = 3.42$, $p < .001$, $d = 0.34$ (Fig. 7c), 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, $F(2, 198) = 1.71$, $p = .18$, $\eta_p^2 = .02$. 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, 9 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 54.0% (SD = 19.6%) of the time, which was 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 55.1% (SD = 20.4%) of the time, which was 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 55.3% (SD = 21.4%) of the time, again reliably above chance (50%), $t(99)$ = 2.47, p = .02, d = 0.25 (Fig. 7d). This suggests that participants successfully inferred the transitive A–C pairs at the base, subordinate, and superordinate levels, even though they saw only 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 = .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, and 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, in which participants successfully inferred A–C pairs after seeing only 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 exposed only 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 conducted 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 marginally 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.

General Discussion

The goal of the current study 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, 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.

Focusing beyond temporal transitivity, we further examined whether 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 4). 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 the MTurk study, 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 driven entirely by the priming of the park names, although it was possible that in Experiments 4 to 7, the transitive inference from city to park could be facilitated by the training session.

Remarkably, in Experiment 7, we found that exposure to A–B and B–C base 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., USA–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 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 inferences to be made. In Experiment 3, we increased the overlap but also reduced the number of inferences and thus could not tease apart these two factors. However, follow-up studies can hold the overlap constant while increasing the number of inferences (e.g., A–B–C, B–C–D, 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 MTurk study, 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.

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).

The current study 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, the current study 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 notice 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.

Action Editor

Edward S. Awh served as action editor for this article.

Author Contributions

Y. Luo and J. Zhao developed the study concept. Both authors contributed to the study design. Testing and data collection were performed by Y. Luo. Y. Luo analyzed and interpreted the data under the supervision of J. Zhao. Both authors drafted the manuscript. Both authors approved the final version of the manuscript for submission.

ORCID iD

Jiaying Zhao https://orcid.org/0000-0003-0902-4343

Acknowledgments

We thank the Zhao Lab, Nick Turk-Browne, Brian Scholl, and Aaron Seitz for helpful comments.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This work was supported by a Natural Sciences and Engineering Research Council of Canada Discovery Grant (RGPIN-2014-05617 to J. Zhao), the Canada Research Chairs program (to J. Zhao), the Leaders Opportunity Fund from the Canadian Foundation for Innovation (F14-05370 to J. Zhao), and an Elizabeth Young Lacey Graduate Scholarship (to Y. Luo).

Supplemental Material

Additional supporting information can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797618762400

Open Practices

All data and materials have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/5yau4/. The design and analysis plans of these experiments were not preregistered. The complete Open Practices Disclosure for this article can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797618762400. This article has

received badges for Open Data and Open Materials. More information about the Open Practices badges can be found at http://www.psychologicalscience.org/publications/badges.

Notes

1. On the basis of the power analysis, we added 10 new participants from a previous round of review, for a total of 40 participants in Experiment 1, to ensure we had 90% power.
2. In a separate experiment ($N = 30$), participants chose the base triplets as more familiar than foils 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 successfully learned the base triplets.
3. We thank an anonymous reviewer for suggesting this experiment.

References

Bays, B. C., Turk-Browne, N. B., & Seitz, A. R. (2015). Dissociable behavioural outcomes of visual statistical learning. *Visual Cognition*, *23*, 1072–1097.

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.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

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

Goodwin, G. P., & Johnson-Laird, P. N. (2005). Reasoning about relations. *Psychological Review*, *112*, 468–493.

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.

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.

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.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

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

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.

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.

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

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

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