

Statistical learning creates implicit subadditive predictions

Yu Luo (yuluo@psych.ubc.ca)

Department of Psychology, University of British Columbia

Jiaying Zhao (jiayingz@psych.ubc.ca)

Department of Psychology, and Institute for Resources, Environment and Sustainability, University of British Columbia

Abstract

The cognitive system readily learns when multiple cues jointly predict a specific outcome. What is less known is how the mind generates predictions when only a single cue is present. In four experiments, participants were first exposed to two objects followed by a circle with a specific size or a specific numeric value. Afterwards, participants viewed a single object and estimated the associated size or value. Finally, participants recalled the size or value that followed the initial two objects. We found that the estimated size associated with the single object was significantly smaller than 100% but significantly larger than 50% of the recalled size associated with the two objects. No participants were consciously aware of the associations. The results reveal a new consequence of statistical learning on automatic inferences: When multiple objects were previously associated with an outcome, the single object is implicitly expected to predict a subadditive outcome.

Keywords: Implicit learning; support theory; subadditive inferences; regularities; predictions

Introduction

A remarkable capacity of the cognitive system is to extract the relationships among objects in the environment. Statistical learning is one mechanism that detects the statistical relationships between individual objects in terms of co-occurrences over space or time (Fiser & Aslin, 2001; Saffran, Aslin, & Newport, 1996). In contrast to other forms of associative learning, statistical learning occurs incidentally, without conscious intent or explicit awareness, and thus observers are often not explicitly aware of object co-occurrences (Turk-Browne, Jungé, & Scholl, 2005; Turk-Browne, Scholl, Chun & Johnson, 2009).

The ability to extract statistical regularities from the environment has a series of cognitive consequences. For example, statistical learning encodes the co-occurring objects more efficiently in working memory (Brady, Konkle, & Alvarez, 2009; Zhao & Yu, 2016), draws attention spontaneously and persistently to the co-occurring objects (Yu & Zhao, 2015; Zhao, Al-Aidroos, & Turk-Browne, 2013; Zhao & Luo, 2017), forms new transitive inferences based on prior associations (Luo & Zhao, 2018), enhances memory representation of individual objects (Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014; Otsuka & Saiki, 2016), and induces false memories of co-occurring objects (Luo & Zhao, 2017).

Past research on statistical learning has predominately focused on associations between individual objects that co-occur in space or time (e.g., A appears next to or before B).

Moreover, most studies in associative learning focused on how the relationship between the cue and the outcome is learned, how learning modulates subsequent processes, and how predictive cues are selectively prioritized (e.g., Mackintosh, 1975; Le Pelley et al., 2016).

In the daily visual environment, multiple objects sometimes co-occur to jointly predict a specific outcome. For example, two co-authors often publish a paper together, or two co-founders start a company. What is less known is how the mind generates predictions when only a single cue is present, after learning that two cues were previously jointly associated with an outcome. For example, when author A and author B have been publishing high-quality papers together, what's the automatic inference when you see a paper by only author A?

Here we examine three possible hypotheses: (1) the complete inheritance hypothesis that suggests that the single cue predicts 100% of the outcome previously associated with the two cues, (2) the proportional inheritance hypothesis that suggests that the single cue predicts 50% of the outcome, and (3) the subadditive hypothesis that suggests that the single cue predicts more than 50% but less than 100% of the outcome previously associated with the two cues. The subadditive hypothesis is consistent with support theory (Tversky & Koehler, 1994), that suggests that when people unpack an event (e.g., the probability of death due to natural causes) into disjoint components (e.g., the probability of death due to heart attack, cancer, or other natural causes), they tend to increase the evidentiary support for the event. In other words, people tend to provide a higher probability of death due to natural causes when they are asked to estimate the probability of death due to each component of natural causes separately, compare to reporting the probability of death due to natural causes as one category.

To test these hypotheses, we conducted a series of four experiments to examine how the mind makes predictions when a single cue is present after learning that multiple cues previously jointly predicted an outcome.

Experiment 1

In this experiment, participants were first exposed to two cues (e.g., red and blue squares) that were immediately followed by an outcome (e.g., a circle with a specific size). We examined how they generated predictions of the outcome when only a single cue was present (e.g., a red square).

Participants

A total of 42 undergraduates (31 female; mean age=19.6 years, SD=1.5) from University of British Columbia (UBC) participated in the experiment for course credit. Participants reported normal or corrected-to-normal visual acuity and provided informed consent. The protocol was approved by the UBC Behavioral Research Ethics Board.

Stimuli

The stimuli consisted of eight squares in eight 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; orange = 255/158/0; brown = 103/29/0). Each square subtended 2.7° of visual angle. The colored squares were randomly assigned into four pairs for each participant and remained constant throughout the experiment. Each color pair was randomly associated with a gray circle (R/G/B = 128/128/128) with a specific diameter. The circle diameter subtended 3.0° (or 100 pixels), 6.0° (or 200 pixels), 9.0° (or 300 pixels), or 12.0° (or 400 pixels) of visual angle (Fig.1a). Thus, each color pair was associated with a circle of a specific size.

Apparatus

Participants in all experiments were seated 50cm from a computer monitor (refresh rate = 60 Hz). Stimuli were presented using MATLAB and PsychophysicsToolbox (<http://psychtoolbox.org>).


Procedure

The experiment consisted of three phases: exposure phase, inference phase, and recall phase. During exposure, two colored squares (e.g., red and blue squares) appeared in a horizontal configuration at the center of the screen for 500ms, followed by a 500ms inter-stimulus interval (ISI), and then the circle with a rotated T in the middle appeared at the center of the screen for 500ms in each trial (Fig.1b). Each color-size pair was repeated 80 times to form a single continuous temporal sequence of color-size pairs in a pseudorandom order with a constraint where no single color-size pair could repeat back-to-back. In total, there were 320 trials. Participants performed a cover task where they judged as quickly and accurately as possible whether the rotated T in the circle was pointing to the left or right (by pressing the “1” or “0” key for left or right, respectively). The cover task was irrelevant to learning the color-size pairs, in order to conceal the true purpose of the study. This also ensured that statistical learning of the color-size pairs was incidental. Participants were not told anything about the color-size pairs.


After exposure, participants performed an inference phase (Fig.1c). In each trial, participants viewed a single color square for 500ms followed by a 3000ms blank screen. Afterwards, a probe circle with a diameter subtending 0.6° (or 20 pixels) was presented on the screen. Participants were asked to estimate the size of the circle that was associated

with the color square by adjusting the size of the probe circle using their mouse. The diameter of the adjustable circle was restricted to a range from 20 pixels to 420 pixels. The adjustable circle remained on the screen until the “a” key was pressed to register participant’s estimate. Each member of a color pair was tested four times, resulting in 32 trials in total (the order of the trials was randomized).


a) 2 colors - size pairings



b) Exposure phase: (cover task) is the rotated T pointing left or right?



c) Inference phase: estimate the size of the dot that follows the color



d) Recall phase: recall the size of the dot that follows the colors




Figure 1. Experiment 1 paradigm. (a) Four color-size pairs were presented (e.g., red and blue squares followed by a circle with a diameter of 100 pixels). (b) Exposure phase using a cover task to expose the color-size pairs to participants. (c) Inference phase where participants estimated the size of the circle that was associated with the color. (d) Recall phase where participants recalled the size of the circle that followed the two color squares.

To examine whether participants had successfully learned the color-size pairs (i.e., the association between the two color squares and the size of the circle), participants completed a size recall task following the inference phase (Fig.1d). In each trial, participants viewed the original color pair (e.g., red and blue squares) that they viewed during exposure for 500ms followed by a 3000ms blank screen. Afterwards, a probe circle with a diameter subtending 0.6° (or 20 pixels) was presented on the screen. Participants were asked to recall the size of the circle that was associated with the original two colors during exposure by adjusting the size of the probe using their mouse. The diameter of the

adjustable circle was restricted to a range from 20 pixels to 420 pixels. The adjustable circle remained on the screen until the “a” key was pressed to register participant’s estimate. Each color pair was tested four times, producing 16 trials in total (the order of the trials was randomized).

A debriefing session was conducted at the end of the experiment, where participants were asked if they had noticed any pairings of squares and circles that appeared one after another. For those who responded yes, we further asked them to write in sentences which type of circle followed which colors.

Results and Discussion

We first analyzed whether the inferred circle size associated with one single object in the pair (e.g., red square) was different from the inferred circle size associated with the other member of the pair (e.g., blue square) to rule out any spatial positioning bias. We found that the inferred circle size associated with one object was not different from the inferred circle size associated with the other member in the pair for all four types of circle diameter (p 's > .19). Thus, we combined the inferred size of either member in the pair.

We also found that in the recall phase, participants overestimated the size of the small circle (mean recalled circle diameter of a circle diameter of 100 pixels was 176.1, $SD=84.5$), and they underestimated the size of the large circle (mean recalled circle diameter of a circle diameter of 400 pixels was 225.8, $SD=100.6$). Given these recall biases, we compared the inferred size with the recalled size, not with the objective size in the following analyses.

The purpose of this experiment was to examine how the mind predicts the outcome given a single predictor, after learning that two predictors were associated with a specific outcome. We compared the inferred size associated with the single object during inference phase to the recalled size associated with the two objects to test the complete inheritance hypothesis. We also compared the inferred size associated with the single object during inference phase to the 50% of the recalled size to test the proportional inheritance hypothesis (Fig.2a).

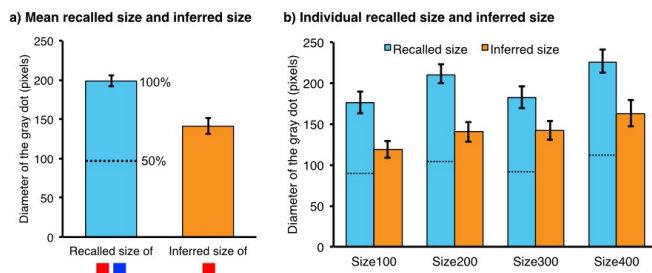


Figure 2. Experiment 1 results. (a) The mean recalled size of the circle associated with two objects and the mean inferred size of the circle associated with a single object. (b) The recalled size of the circle associated with two objects and the inferred size of the circle associated with a single object for each color-size pairing (error bars reflect ± 1 SEM; dashed line represents 50% of the recalled size).

We found that the inferred size associated with the single object (mean inferred diameter=141.1, $SD=63.6$) was significantly smaller than the recalled size associated with the two objects (mean recalled diameter=198.6, $SD=46.8$) [$t(41)=6.90$, $p<.001$, $d=1.03$], but significantly larger than 50% of the recalled size [$t(41)=5.01$, $p<.001$, $d=0.87$] (corrected for multiple comparisons). Additionally, the same results were consistently found for each color-size pairing (Fig.2b). The results support the subadditive hypothesis. During debriefing, three participants reported noticing the color-size pairs, but none could correctly report which circle size followed which specific color pair. This suggests that participants had no explicit awareness of the color-size pairs.

These findings suggest that people implicitly predict a subadditive outcome from a single predictor after learning that two predictors previously jointly predicted a specific outcome.

Experiment 2

This experiment aimed to replicate and extend the findings in Experiment 1 by increasing the number of predictors from two to three.

Participants

A new group of 40 undergraduates (34 female, mean age=19.7 years, $SD=2.2$) from UBC participated in the experiment for course credit.

Stimuli

The stimuli were identical to those in Experiment 1, except that we added a black color (R/G/B=0/0/0) to the color set. There were nine color squares in total, randomly assigned into three triplets for each participant. Each triplet was randomly associated to a gray circle with a specific diameter. The circle diameter subtended 3.0° (or 100 pixels), 7.5° (or 250 pixels), or 12.0° (or 400 pixels) of visual angle (Fig.3a).




Figure 3. Experiment 2. (a) Three color-size pairs (e.g., red, blue, and green squares–circle with a diameter of 100 pixels). (b) The mean recalled size of the circle associated with three objects and the mean inferred size of the circle associated with a single object (error bars reflect ± 1 SEM; dashed line represents 33% of the recalled size).

Procedure

The procedure was identical to that in Experiment 1, except that the three color squares were followed by a circle of a given size in the exposure phase, and participants recalled the circle size associated with the three squares in the recall phase.

Results and Discussion

In a one-way repeated-measures ANOVA, we found no difference between the inferred circle size associated with each object in the triplet for all three types of circle size (p 's > .55). Thus, we combined the inferred size of each member in the triplet. We also found that participants overestimated the size of the small circle (mean recalled circle diameter of a circle diameter of 100 pixels was 188.4, $SD=78.1$) and underestimated the size of the larger circle (mean recalled circle diameter of a circle diameter of 400 pixels was 197.3, $SD=94.8$). Given these biases, we compared the inferred size with the recalled size, not with the objective size in the following analyses.

We found that the inferred size associated with the single object (mean diameter = 124.2, $SD=59.5$) was significantly smaller than the recalled size associated with the three objects (mean diameter = 198.3, $SD=53.1$) [$t(41)=7.87$, $p<.001$, $d=1.31$], but significantly larger than 33% of the recalled size [$t(41)=6.90$, $p<.001$, $d=1.32$] (corrected for multiple comparisons; Fig. 3b). The results again support the subadditive hypothesis.

During debriefing, two participants reported noticing the color-size pairs, but none could correctly report which circle size followed the specific color triplet. This suggests that participants had no explicit awareness of the color-size pairs.

These findings successfully replicated the findings in Experiment 1, showing that people implicitly predict a subadditive outcome from a single predictor after learning that three predictors previously jointly predicted a specific outcome.

Experiment 3

Experiment 3 aimed to generalize the findings to other types of outcomes from circle sizes to numeric values. Specifically, after learning that two objects (e.g., red and blue squares) were associated with a specific numeric value, we examined how people made predictions of value from a single predictor (e.g., red square).

Participants

A new group of 45 undergraduates (41 female, mean age = 20.38 years, $SD=2.8$) from UBC participated in the experiment for course credit.

Stimuli

The stimuli were identical to those in Experiment 1, except that each color pair was associated with a specific three-digit number. There were four three-digit numbers: 150,

400, 650, and 900. Each number was associated with a color pair (Fig. 4a).

Procedure

As in Experiment 1, there were three phases (exposure, inference, and recognition). The exposure phase was identical to Experiment 1, except that in the cover task, participants viewed a three-digit number above the rotated T in the circle (Fig. 4b). Since that a specific number may be easier to learn than the size of a circle, we reduced the number of repetitions for each color-number pair to 40 times, resulting in 160 trials in total (the order of trials was randomized).

a) 2 colors - number pairings



b) Exposure phase: (cover task) is the rotated T pointing left or right?




Figure 4. Experiment 3 paradigm. (a) Four color-number pairs (e.g., red and blue squares-150). (b) Exposure phase using a cover task to expose the color-number pairs to participants.

In the inference phase, participants viewed a single colored square and were asked to estimate the number that was associated with the color square by typing a number on the keyboard. The estimated number was restricted to a range from 0 to 1050. Participants had the option to delete and revise their estimated number until the “a” key was pressed to register their estimate.

In the recognition phase, participants viewed a pair of color squares that was presented in exposure and were asked to recall the number that was associated with the color pair by typing the number on the keyboard. The recalled number was restricted to a range from 0 to 1050. Participants had the option to delete and revise their recalled number until the “a” key was pressed to register their estimate. A debriefing session was conducted at the end as before.

Results and Discussion

We found that the inferred number associated with a single object was not different from the inferred number associated with the other member in the pair (p 's > .32). Thus, we combined the inferred number of each member in the pair. We also found that participants overestimated the small number (mean recalled number of 150 was 489.7,

SD=227.3) and underestimated the large number (mean recalled number of 900 was 557.0, SD=251.7). Given these biases, we compared the inferred number with the recalled number, not with the objective number in the following analyses.

We found that the inferred number associated with the single object (mean inferred number=476.5, SD=150.8) was marginally smaller than the recalled number associated with the two objects (mean recalled number=513.2, SD=97.3) [$t(44)=1.79$, $p=.08$, $d=0.29$], but significantly larger than 50% of the recalled number [$t(44)=10.86$, $p<.001$, $d=1.96$] (corrected for multiple comparisons; Fig.5). The results again support the subadditive hypothesis.




Figure 5. Experiment 3 results. The mean recalled number associated with two objects and the mean inferred number associated with a single object (error bars reflect ± 1 SEM; dashed line represents 50% of the recalled number).

During debriefing, two participants reported noticing the color-number pairs, but none could correctly report which number followed which specific colors. This suggests that participants had no explicit awareness of the color-number pairs.

These findings again replicated the findings in Experiment 1, showing that people implicitly predict a subadditive outcome from a single predictor after learning that two predictors previously jointly predicted a specific outcome.

Experiment 4

This experiment aimed to extend the findings in Experiment 3 by increasing the number of predictors from two to three.

Participants

A new group of 33 undergraduates (28 female, mean age=20.2 years, SD=1.7) from UBC participated in the experiment for course credit.

Stimuli and Procedure

The stimuli and the procedure were identical to Experiment 3, except that there were three color triplets and each triplet was associated with 150, 525, or 900 (Fig.6a).

Results and Discussion

In a one-way repeated-measures ANOVA, we found no difference between the inferred number associated with each object in the triplet for all three types of numbers ($p>.35$). Thus, we combined the inferred number of each member in the triplet. We also found that participants overestimated the small number (mean recalled number of 150 was 485.7, SD=267.0) and underestimated the large number (mean recalled number of 900 was 592.7, SD=246.8). Given these biases, we compared the inferred number with the recalled number, not with the objective number in the following analyses.

We found that the inferred number associated with the single object (mean inferred number=401.1, SD=180.3) was significantly smaller than the recalled number associated with the three objects (mean recalled number=501.3, SD=106.8), [$t(32)=3.03$, $p=.005$, $d=0.68$], but significantly larger than 33% of the recalled number [$t(32)=7.61$, $p<.001$, $d=1.80$] (corrected for multiple comparisons; Fig.6b). The results again support the subadditive hypothesis.

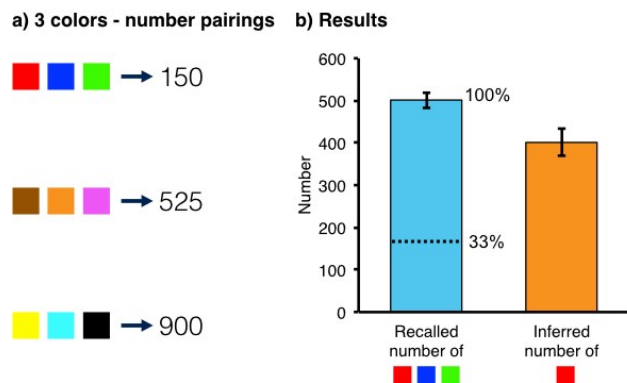


Figure 6. Experiment 4. (a) Three color-number triplets (e.g., red, blue, and green squares-150). (b) The mean recalled number associated with three objects and the mean inferred number associated with a single object (error bars reflect ± 1 SEM; dashed line represents 33% of the recalled number).

During debriefing, one participant reported noticing the color-number pairs, but the participant could not correctly report which number followed which specific colors. This suggests that participants had no explicit awareness of the color-number pairs.

These findings replicated the findings in Experiment 3, showing that people implicitly predict a subadditive outcome from a single predictor after learning that three predictors previously jointly predicted a specific outcome.

General Discussion

The goal of this study was to examine how the mind automatically generates prediction when only a single cue is present, after learning that multiple cues were previously jointly associated with an outcome. We found that after learning that two co-occurring objects (e.g., red and blue squares) predicted a specific circle size, participants inferred the circle size associated with a single color (e.g., red

square) to be smaller than the original circle size associated with the color pair, but larger than 50% of the circle size associated with the color pair (Experiment 1). We further extended the number of predictors from two to three. After learning that three co-occurring objects predicted a specific circle size, participants inferred the circle size associated with a single color to be smaller than the circle size associated with the color triplet, but larger than 33% of the circle size associated with the color triplet (Experiment 2). We further replicated and extended the experiment from circle sizes to numeric values as outcomes for two predictors (Experiment 3) and three predictors (Experiment 4). Importantly, no participant was consciously aware of the association between the predictors and the outcome across all experiments, suggesting that the inference of the size or number associated with one single predictor was largely implicit. The current findings also suggest when people predict an outcome relying on a single cue from a set of cues, they do not inherently generate the prediction based on the outcome associated with the complete set of cues, nor do they proportionally inherit the outcome based on the number of cues. Instead, they make predictions in a subadditive manner, which is consistent with support theory (Tversky & Koehler, 1994).

One rationale behind support theory is that unpacking an event to its individual component may evoke other relevant elements that might have been missed. When participants were asked to infer the size associated with each individual color in the pair or triplet, they might have to think more extensively for each color, compared to recalling the outcome associated with the color pair or triplet. A second rationale behind support theory is that explicitly referring to an individual component of an event would increase its salience. When participants were asked to infer the size associated with a single color, their attention was drawn to the single color which may increase the weight of the single color in their prediction of the outcome.

Alternatively, previous studies have suggested that seeing one object in a pair may activate the unitized representation of the pair (e.g., Alvarez & Oliva, 2008). The co-occurring objects (e.g., red and blue squares) may be grouped in the mind during learning. When participants were asked to predict the outcome relying on a single object (e.g., a red square), the object may trigger the representation of the group but not fully activate the representation of the group. Therefore, participants may predict an outcome above 50% but less than 100% of the original outcome.

Another possible explanation is that participants could add the size predicted by each colored square in a sublinear fashion, creating a subadditive sum. A new experiment is needed to test this hypothesis to tease apart whether the subadditivity is driven by the sublinear representation of each size or the sublinear summation of the two sizes. Specifically, participants are first exposed to one unique color predicting a unique size during the exposure phase (e.g., a red square predicting a circle with a certain diameter, and a blue square predicting a circle with a certain

diameter). In the inference phase, participants see a red square with a blue square presented side by side simultaneously, and they will be asked to infer the circle size associated with the two squares. In the recall phase, participants simply recall the original size of the circle associated with the red square and the blue square. If the inferred size is equal to the sum of the two recalled sizes, then this would suggest that participants use an additive approach to predict the outcome. If the inferred size is smaller than the sum of the two recalled sizes but larger than each recalled size, then this would suggest that participants use a subadditive approach.

In summary, we found a new consequence of statistical learning on automatic inferences: When multiple objects jointly predict a specific outcome, the presence of a single object implicitly triggers a subadditive prediction.

Acknowledgments

We would like to thank Oleg Urminsky, Ru Qi Yu, Brandon Tomm, and two anonymous reviewers for their helpful comments. This work was supported by NSERC Discovery Grant (RGPIN-2014-05617 to JZ), the Canada Research Chairs program (to JZ), the Leaders Opportunity Fund from the Canadian Foundation for Innovation (F14-05370 to JZ), and Cordula and Gunter Paetzold Fellowship (to YL).

References

- Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of attention. *Psychological Science*, 19, 392–398.
- 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., Konkle, T., Alvarez, G. A. and Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences, USA*, 105, 14325–14329.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12, 499–504.
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences*, 111, 8997–9002.
- Luo, Y., & Zhao, J. (2017). Learning induced illusions: Statistical learning creates false memories. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, (pp. 774–779). Austin, TX: Cognitive Science Society.
- Luo, Y., & Zhao, J. (2018). Statistical Learning Creates Novel Object Associations via Transitive Relations. *Psychological Science*, 29, 1207–1220.

- Mackintosh, N. J. (1975). A theory of attention: variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276.
- Otsuka, S., & Saiki, J. (2016). Gift from statistical learning: Visual statistical learning enhances memory for sequence elements and impairs memory for items that disrupt regularities. *Cognition*, 147, 113-126.
- Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: an integrative review. *Psychological Bulletin*, 142, 1111.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926-1928.
- Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological review*, 101, 547.
- 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.
- 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.