Finding the pattern: online extraction of spatial structure during virtual navigation

Kathryn N. Graves

Department of Psychology

Yale University

New Haven, CT

`kathryn.graves@yale.edu`

James W. Antony

Princeton Neuroscience Institute

Princeton University

Princeton, NJ

Nicholas B. Turk-Browne

Department of Psychology

Yale University

New Haven, CT

Abstract

While navigating the world, we pick up on patterns of where things tend to appear. According to theories of memory and studies of animal behavior, knowledge of these patterns emerges gradually over days or weeks, via consolidation of individual navigation episodes. Here we discover that navigation patterns can also be extracted online, prior to the opportunity for offline consolidation, as a result of rapid statistical learning. Human participants navigated a virtual water maze in which platform locations were drawn from a spatial distribution. Within a single session, participants increasingly navigated through the mean of the distribution. This behavior was better simulated by random walks from a model with only an explicit representation of the current mean, compared to a model with only memory for the individual platform locations. These results suggest that participants rapidly summarized the underlying spatial distribution and used this statistical knowledge to guide future navigation.

Keywords: memory integration, statistical learning, episodic encoding, consolidation, complementary learning systems

Finding the pattern: online extraction of spatial structure during virtual navigation

Introduction

Once you find your way to a new restaurant for the first time, you are able to find it more easily the next time. This ability to use past episodes of navigation to achieve current goals is ubiquitous across species (Barry, Coogan, & Commins, 2016; Cornwell, Johnson, Holroyd, Carver, & Grillon, 2008; Judd & Collett, 1998). However, recollecting specific episodes is not the only way to draw upon past experiences when navigating. When seeking a new location, it may instead be more efficient to rely on general knowledge aggregated across multiple episodes. For example, if looking for a new restaurant, an effective strategy might be to head toward the neighborhood with other restaurants you have liked.

Aggregating across episodes is useful because it allows us to learn statistical structure, such as the spatial pattern of locations previously visited. Theories of memory (McClelland, McNaughton, & O'Reilly, 1995) typically describe this process as resulting from slow and gradual consolidation, whereby distinct episodes are integrated during offline sleep or rest over many days. This memory transformation has been demonstrated in rodents, who after 30 (but not 7) days navigated a familiar arena according to the spatial distribution of previously learned locations rather than individual past locations (Richards et al., 2014).

Recent work suggests that such statistical learning can also occur rapidly, within minutes, potentially in parallel with the encoding of episodes (Schapiro, Turk-Browne, Botvinick, & Norman, 2017). However, many of the studies showing statistical learning on this short timescale examine temporal structure in sequences (Saffran, Newport, & Aslin, 1996; Turk-Browne, Jungé, & Scholl, 2005), with the literature on spatial structure mostly limited to brief arrays of simple objects (Chun & Jiang, 1998; Fiser & Aslin, 2001). Spatial structure is inherently different in navigation, as patterns of locations exist in allocentric rather than egocentric space (i.e., in world-centered coordinates, relative to landmarks). How humans deploy statistical learning during spatial navigation remains an open question.

Here we test this question using a virtual analog of the delayed match-to-place water maze task (Morris, 1984). Across repeated attempts to locate hidden "platforms" within a single session, we quantify the extent to which humans navigate based on previous platform locations vs. the mean of their spatial distribution (itself not a platform location). We interpret this behavior by modeling it with artificial agents who take random walks with knowledge of either individual past platform locations or the mean location. Despite the lack of offline consolidation, we hypothesized that participants would shift within session from initially navigating to individual platforms (episode-based navigation) to navigating toward their mean (pattern-based navigation), after accruing enough experience to extract a pattern.

Methods

Participants

Thirty participants (16 female, 18-22 years) completed the study for course credit. The Yale University Human Subjects Committee approved the study protocol. We arrived at this sample size by conducting a power analysis on data from a pilot study, using a bootstrapping procedure and Wilcoxon's signed-rank test. We drew from 14 pilot subjects with replacement to generate a range of sample sizes. We analyzed the increase in number of crossings through platform locations from the first to second half of the experiment. We chose this metric to estimate power because it is analogous to the canonical "escape latency" measure of spatial learning (Cornwell et al., 2008; Mueller, Jackson, & Skelton, 2008). We ran 1000 iterations per sample size and calculated the proportion of tests in which the effect was below a significance level of 0.05. Using this proportion as our estimate of power, a sample size of 30 participants was sufficient to detect this effect with greater than 99% power. This sample size also falls in the range of other studies that showed a learning effect on escape latency (range 15-52) (Cornwell et al., 2008; Mueller et al., 2008; Richards et al., 2014).

Stimuli

The virtual environment was designed as a circular arena, adapted from Woolley et al. (2015) and constructed in Blender (www.blender.org). The environment was rendered in MATLAB (Mathworks, Natick, MA, USA), and Psychtoolbox (Brainard & Vision, 1997) was used to display task instructions. The arena was designed graphically with an island theme. The circular floor (radius = 7.85 units) was covered by an image of sand (see Fig. 1a). Mountains appeared on the north end and palm trees on the south end (see Fig. 1b); similar to the classic water maze task, these landmarks served as directional headings. The “platforms” were hidden nests of turtle eggs and, following Richards et al. (2014), their locations were drawn from a normal distribution of distances d from center ($\mu = 3.4$ units, $\sigma = 0.9067$) and a circular normal distribution of angles θ between the platform and the eastern cardinal direction ($\mu = 0.2618$ radians, $\kappa = 8$).

Procedure

Participants were instructed that they were on an island, searching for a total of 20 turtle eggs buried below the sand. They could only find one egg at a time and had to walk around the beach until they found it. Participants used the I key to walk forward, the J and L keys to turn left and right, respectively, and the M key to walk backward.

At the beginning of the search for each new egg, the participant’s location was initialized to the center of the arena. For the first 5 s, the ground was green and the I/M keys were locked only allowing participants to rotate in place left or right. The ground then turned to sand, cuing participants that they could now start moving forward and backward. There was no time limit and thus the searches varied in length (40 mins total on average). Participants successfully located an egg when they walked within a one-unit radius of the platform location of that egg. When this occurred, further navigation was locked and the walls turned green to reveal the outline of an egg. The screen then went black and participants were instructed to return to the same location (with the cover story of needing to return to the site with the appropriate

tools). Their location was then re-initialized at the center of the arena and they navigated to the latest egg location. The purpose of having them repeat the search was to aid episodic encoding of the egg location in the arena with respect to the landmarks. After finding the egg again, participants were told how many new eggs remained to be found before beginning to search for the next egg (see Fig. 1c).

Fig. 1. Task design and performance. (a) The virtual water maze was a circular, beach-themed arena. (b) The mountains and palm trees served as north and south directional headings, respectively. (c) Participants were prompted to begin searching at the start of each trial. They were initialized in the center of the arena and could only rotate left and right initially, while the ground was green. After 5 s, the ground turned to sand and participants could then freely navigate. When they found an egg, they were instructed to return to the same location on the next trial. On the trial after that, they began searching for a new egg location. (d) We quantified location crossings and time in location during participant navigation. An example participant shows increased navigation to both locations over the course of the experiment.

Data Analysis

The analyses below focus on the search paths from trials when participants first found an egg. There was less need to search when they returned to the same location on the next trial. We did not include search paths from the first two eggs because it was not yet possible to extract a pattern and thus episode- and pattern-based navigation could not be distinguished.

Raw data consisted of X and Y coordinates output every 40 ms of navigation. For each search trial, we quantified pattern- and episode-based navigation as the number of crossings through each representative location. We used the mean of the currently known points (i.e., on trial four, the mean of egg locations from trials one through three) for pattern-based navigation, and the most recently learned egg location (i.e., on trial four, crossings through the egg location from trial three) for episode-based navigation. We focused on the most recent location because we expected this to be the strongest memory trace and thus to provide a conservative comparison for navigation to the mean (see Fig. S1).

In order to generate a mean location radius through which to quantify a crossing (e.g. an instance in which a pair of coordinates are outside the radius, and the immediately following pair are inside the radius), we wanted a value that would encapsulate the change in representation across trials as more data points (egg locations) were acquired. We thus, for each trial, quantified the Full-Width, Half-Maximum of the currently known distribution, and halved this value to generate a radius around the coordinates of the current mean. In order to be maximally conservative and avoid any bias due to the increasing size of this mean location radius across the task, we used the same radius value for the recently learned location (see Fig. 1d).

In addition to crossings, we also analyzed the total time spent in the most recent and the current mean locations, for episode- and pattern-based navigation, respectively. We calculated this as the number of raw data points (output every 40ms) generated within these location radii. To account for differences in trial length, we normalized

each measure by the total distance traveled during that trial. We used total distance rather than total time to normalize because trial length *per se* does not increase crossings or time in (non-center) locations. Finally, we binned the data into first and second experiment halves in order to examine learning.

Results

Human Behavior

We first investigated changes in crossings through, and time spent in, the most recent egg location (episode-based navigation) and the cumulative mean egg location (pattern-based navigation), from the first to second halves of the experiment. Our analyses were based on our *a priori* prediction that pattern-based navigation would exceed episode-based navigation in the second half. Thus, we ran mixed-effects linear regressions with condition (most recent and cumulative mean), experiment half (first and second), and their interaction as regressors predicting number of location crossings and the time in each location. To control for effects of distance traveled per trial, we included total distance as a regressor, as well as its interactions with each of our variables of interest. Random intercepts for each participant were also included in the models.

The analysis of location crossings revealed main effects of condition ($\beta = -0.01$, $t = -2.21$, $p = 0.030$), experiment half ($\beta = 0.02$, $t = 4.61$, $p < 0.001$), and distance ($\beta = -0.02$, $t = -2.77$, $p = 0.007$). There were interactions between condition and experiment half ($\beta = -0.01$, $t = -2.06$, $p = 0.043$) and between experiment half and distance ($\beta = -0.01$, $t = -2.51$, $p = 0.013$); importantly, we did not find an interaction between condition and distance ($\beta = 0.01$, $t = 1.25$, $p = 0.216$), indicating that differences in crossings through the cumulative mean and most recent location were not merely a function of overall trial duration and number of steps taken.

Comparison against a simple main effects model revealed that our model better predicted participant behavior ($\chi^2(3) = 13.39$, $p = 0.004$). We then applied Wilcoxon's signed-rank test to our predicted condition x experiment half interaction and found

more cumulative mean vs. most recent location crossings during the second half ($Z = -2.30$, $p = 0.021$, $r = 0.48$, 95% CI = [0.16, 0.72]) but not the first half ($Z = -0.45$, $p = 0.655$, $r = -0.085$, 95% CI = [-0.45, 0.27]) (Figure 2a).

The analysis of time in location similarly revealed a main effect of experiment half ($\beta = 19.13$, $t = 5.15$, $p < 0.001$), but only a marginal effect of distance ($\beta = -7.84$, $t = -1.84$, $p = 0.065$), and no main effect of condition ($\beta = -5.59$, $t = -1.13$, $p = 0.264$). There were again interactions between condition and experiment half ($\beta = -10.91$, $t = -2.10$, $p = 0.039$) and between experiment half and distance ($\beta = -10.37$, $t = -3.10$, $p = 0.002$), but no interaction between condition and distance ($\beta = 3.18$, $t = 0.61$, $p = 0.541$). Compared to a simple main effects model, our model again provided a better fit to participant data ($\chi^2(3) = 14.58$, $p = 0.002$). Subsequent tests of the condition by experiment half interaction revealed that participants spent significantly more time in the cumulative mean vs. most recent location during the second half ($Z = -2.84$, $p = 0.005$, $r = 0.56$, 95% CI = [0.25, 0.82]) but not the first half ($Z = -0.69$, $p = 0.49$, $r = -0.13$, 95% CI = [-0.46, 0.22]).

In both of our metrics, participants showed increased pattern-based navigation relative to episode-based navigation over the course of a single session (see Fig. 2a).

Fig. 2. Human behavior and model fits. (a) The number of recently learned and mean location crossings in the first and second half of the experiment are represented for each participant by a blue and red line, respectively. The bold lines represent the mean number of crossings for each location, and the error bars are ± 1 within-subject SE. Amount of time in each location is displayed similarly. (b) The average beta weights for each model from 1000 regressions, using simulated agent performance on each experiment half to predict human performance, is indicated by the bar heights. Blue bars represent weights for the Episode Model, and red bars represent weights for the Pattern Model. A subset of 250 beta weights for each model indicate the distribution of fits. Error bars are 95% confidence intervals.

Control Distribution

As a sanity check that our key results were location-specific and not an artifact of averaging or other analysis choices, we repeated the above analyses after rotating the spatial distribution of locations by 90 deg.

If the apparent increase in pattern-based navigation reflected learning of the true

distribution, the critical interaction between condition and experiment half should be eliminated. Indeed, using the same regressors as our main analyses, neither the rotated model of location crossings ($\beta = -0.00090$, $t = -3.81$, $p = 0.705$) nor the rotated model of time in location ($\beta = -3.31$, $t = -1.14$, $p = 0.258$) revealed a condition by experiment half interaction (See Fig. S2).

We additionally expected fewer crossings through, and less time spent in, the rotated locations overall. Comparison of the true and rotated distributions revealed robustly more crossings through real locations ($Z = -9.12$, $p < 0.001$, $r = 0.74$, 95% CI = $[0.66, 0.91]$), as well as more time spent in real locations ($Z = -6.11$, $p < 0.001$, $r = 0.53$, 95% CI = $[0.39, 0.65]$).

Simulated Agents

We lastly wanted to verify that our main findings reflected pattern-based navigation and could not be explained as incidental to episode-based navigation. That is, because the shortest path to certain egg locations was inadvertently through the mean, and the number of such locations increased over time, an agent with only knowledge of previous locations may still show an increase in cumulative mean location navigation from the first to second halves of the experiment.

To address this possibility, we simulated task performance with two biased Gaussian random walk models — an episode model and a pattern model — and fit participant behavior with each model’s simulated navigational routes. Random walks have previously been shown to be capable of imitating certain optimal behaviors (Abbott, Austerweil, & Griffiths, 2015). We predicted that if participant behavior was driven by an explicit representation of the distribution mean, it would be better fit by a randomly navigating agent with access only to the cumulative mean, compared to an agent with access only to the individual past locations (which controls for the amount of incidental mean navigation en route to specific locations).

To approximate participant behavior, each model’s search duration was constrained to recapitulate the group average number of location crossings per trial; for

the episode model, this location constraint was the average number of platform crossings, whereas for the pattern model, it was the average number of mean location crossings. We determined which egg locations were crossed by the episode model on each trial, and in what order, by randomly drawing from previously learned locations weighted by recency. Thus, the most recently learned egg location was the most likely to be crossed first. For the pattern model, the only location represented was the mean of the currently known locations.

Like human participants, each model was initialized at the origin of a grid. The episode model navigates with Gaussian noise through previously learned locations in the order they were drawn, and the pattern model navigates with Gaussian noise through the mean. Steps in two-dimensional space were generated by:

$$X_{s+1} = X_s + K \quad (1)$$

$$Y_{s+1} = Y_s + K \quad (2)$$

Where K was drawn from a Gaussian distribution ($\mu = 0, \sigma = 1$). Distance from the next step to the goal location (a randomly drawn previous egg location for the episode model and the mean for the pattern model) was constrained such that:

$$D_{s+1} \leq D_s + D_s \phi \quad (3)$$

Where ϕ is a temperature parameter initialized at 0.5 that monotonically decreased on each trial. This parameter was used to capture participants' decreased randomness and increased navigation toward previous locations across trials, as evidenced by the reduction in distance traveled per trial over the course of the experiment.

We simulated 1000 agents per model and repeated the mixed-effects linear regression on each model's output. For the episode model (see Fig S3a), the critical interaction between condition and experiment half was found in analyses of both location crossings ($\beta = -0.02, t = -19.71, p < 0.001$) and time in location ($\beta = -1.89, t = -25.81, p < 0.001$). This was reflected in more crossings through the cumulative mean location in the second half ($Z = -16.64, p < 0.001, r = 0.53, 95\% \text{ CI} = [0.48, 0.57]$)

relative to first half ($Z = -15.07$, $p < 0.001$, $r = -0.46$, 95% CI = $[-0.51, -0.41]$), and more time spent in the cumulative mean location in the second half ($Z = -20.48$, $p < 0.001$, $r = 0.61$, 95% CI = $[0.57, 0.65]$) relative to the first half ($Z = -21.92$, $p < 0.001$, $r = -0.64$, 95% CI = $[-0.68, -0.60]$).

For the pattern model (see Fig. S3b), there was again an interaction between condition and experiment half for both location crossings ($\beta = -0.05$, $t = -30.13$, $p < 0.001$) and time in location ($\beta = -2.31$, $t = -33.07$, $p < 0.001$). This was again reflected in more crossings through the cumulative mean location in the second half ($Z = -26.46$, $p < 0.001$, $r = 0.84$, 95% CI = $[0.83, 0.85]$) relative to first half ($Z = -5.05$, $p < 0.001$, $r = 0.16$, 95% CI = $[0.10, 0.22]$), and more time spent in the cumulative mean location in the second half ($Z = -31.10$, $p < 0.001$, $r = 0.82$, 95% CI = $[0.81, 0.84]$) relative to the first half ($Z = -8.52$, $p < 0.001$, $r = -0.27$, 95% CI = $[-0.32, -0.21]$).

Given that both models demonstrated more pattern-based than episode-based navigation in the second half of the experiment, it is critical for our analyses that this difference be significantly larger in the pattern model than the episode model. If both models display analogous increases in mean-based navigation, then we have failed to generate models that "learn" differently over the course of the task as a function of their different representations. Indeed, in the second half, the ratio of mean navigation to recently learned location navigation was significantly larger in the pattern model than the episode model, both in crossings ($Z = -25.05$, $p < 0.001$, $r = -0.71$, 95% CI = $[-0.74, -0.68]$) and time in location ($Z = -22.70$, $p < 0.001$, $r = -0.66$, 95% CI = $[-0.69, -0.62]$). This shows that knowledge of episodes alone yielded qualitatively less pattern-based navigation and was not sufficient to explain the full increase in navigation to the mean.

To further support the conclusion that participants extracted the mean, we attempted to fit participant behavior with both models (see Fig 2b). Specifically, we divided each model's crossings and time in location per trial into the first and second halves of the experiment (nine trials each after excluding the first two trials). We then ran linear regressions for the 1000 simulations, separately using each model's "behavior" to predict participant behavior. Specifically, we z-scored each model's number of

crossings and amount of time in location and input the z-scores into the regressions. We chose to use z-scores instead of raw model performance to ensure a common scale between the two models. The result of this procedure was a distribution of 1000 beta weights averaged across participants for each model, experiment half, and metric (crossings and time). We analyzed the data for each metric using an ANOVA with two factors (model and experiment half).

In predicting crossings, there were main effects of model ($F(1,999) = 26.66$, $p < 0.001$) and experiment half ($F(1,999) = 2031.36$, $p < 0.001$), as well as an interaction between model and experiment half ($F(1,999) = 128.30$, $p < 0.001$). This interaction reflected a switch from a greater fit for the episode model in the first half of the experiment ($t(999) = -4.61$, $p < 0.001$, $d = 0.15$, 95% CI $[-0.0033, 0.0013]$) to a greater fit for the pattern model in the second half ($t(999) = 11.45$, $p < 0.001$, $d = 0.36$, 95% CI $[0.0050, 0.0070]$).

Likewise when predicting time in location, there were main effects of model ($F(1,999) = 5.51$, $p < 0.001$) and experiment half ($F(1,999) = 249.18$, $p < 0.001$), as well as an interaction ($F(1,999) = 108.83$, $p < 0.001$). We again observed a switch from dominance of the episode model in the first half ($t(999) = -10.21$, $p < 0.001$, $d = 0.32$, 95% CI $[-2.96, -2.01]$) to the pattern model in the second half ($t(999) = 5.09$, $p < 0.001$, $d = 0.16$, 95% CI $[0.93, 2.17]$).

This change over time in the best fitting model, for both crossings and time in location, provides support for our hypothesis of increased pattern-based navigation as a result of statistical learning within the session (see Fig. 2b).

Discussion

In the current study, we investigated two kinds of spatial representation that could be used during navigation. We compared the role of memories for individual past locations against generalized knowledge of their spatial pattern. We found evidence of navigation to the mean of the distribution of past locations, and used computational modeling to show that this pattern extraction was not incidental to having more

locations in memory. Consistent with our hypothesis, navigation shifted from being episode-based to pattern-based.

Conventional theories view pattern extraction as occurring over days or weeks (McClelland et al., 1995; Squire & Alvarez, 1995; Winocur & Moscovitch, 2011). These theories have found support in findings from rodents, whose spatial and contextual memories become generalized and lose episodic detail after a week to a month (Richards et al., 2014; Winocur, Moscovitch, & Sekeres, 2007). The timescale of transformation in the current study, within an hour, cannot readily be explained in these frameworks. To the extent that patterns are extracted online, this might suggest an update to conventional theories: offline consolidation may operate not only over individual episodes, but also over patterns extracted within a day that have already begun the process of memory transformation.

A recent theory of memory fits more naturally with our findings (Schapiro et al., 2017). According to this theory, the key brain region that stores episodic memories – the hippocampus – also extracts patterns over the same inputs in parallel. The existence of both kinds of representation in the hippocampus during or after navigation is consistent with some rodent studies (Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004). The hippocampus is thought to accomplish episodic encoding and statistical learning in different pathways, centering on CA3 and CA1 subfields, respectively. The re-use of CA1 place cells across environments encountered in the same session (Cai et al., 2016) provides a potential mechanism for rapid pattern extraction during spatial navigation.

Author Contributions

K. N. Graves and N. B. Turk-Browne developed the study design with input from J. W. Antony. Testing, data collection, and data analysis were performed by K. N. Graves with input from N. B. Turk-Browne. K. N. Graves drafted the manuscript and incorporated feedback from N. B. Turk-Browne and J. W. Antony. All authors approved the final manuscript for submission.

Open Practices Statement

The experiment reported in this article was not formally preregistered. The data will be made available on a permanent third-party archive upon article acceptance, and before then will be made available to reviewers and/or editors upon request. Questions about the data or requests for analysis code should be sent to kathryn.graves@yale.edu.

References

- Abbott, J., Austerweil, J., & Griffiths, T. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, *122*(3), 558–569.
- Barry, D. N., Coogan, A. N., & Commins, S. (2016). The time course of systems consolidation of spatial memory from recent to remote retention: A comparison of the immediate early genes *zif268*, *c-fos* and *arc*. *Neurobiology of Learning and Memory*, *128*, 46–55.
- Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., ... others (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*, *534*(7605), 115.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*(1), 28–71.
- Cornwell, B. R., Johnson, L. L., Holroyd, T., Carver, F. W., & Grillon, C. (2008). Human hippocampal and parahippocampal theta during goal-directed spatial navigation predicts performance on a virtual morris water maze. *Journal of Neuroscience*, *28*(23), 5983–5990.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499–504.
- Judd, S., & Collett, T. (1998). Multiple stored views and landmark guidance in ants. *Nature*, *392*(6677), 710.
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas ca3 and ca1. *Science*, *305*(5688), 1295–1298.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419.

- Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods*, 11(1), 47–60.
- Mueller, S. C., Jackson, C. P., & Skelton, R. W. (2008). Sex differences in a virtual water maze: An eye tracking and pupillometry study. *Behavioural Brain Research*, 193(2), 209–215.
- Richards, B. A., Xia, F., Santoro, A., Husse, J., Woodin, M. A., Josselyn, S. A., & Frankland, P. W. (2014). Patterns across multiple memories are identified over time. *Nature Neuroscience*, 17(7), 981.
- Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language*, 35(4), 606–621.
- Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160049.
- Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*, 5(2), 169–177.
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, 134(4), 552.
- Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5), 766–780.
- Winocur, G., Moscovitch, M., & Sekeres, M. (2007). Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nature Neuroscience*, 10(5), 555.
- Woolley, D. G., Mantini, D., Coxon, J. P., D’Hooge, R., Swinnen, S. P., & Wenderoth, N. (2015). Virtual water maze learning in human increases functional connectivity between posterior hippocampus and dorsal caudate. *Human Brain Mapping*, 36(4), 1265–1277.