

Connecting the Dots Without Top-Down Knowledge: Evidence for Rapidly-Learned Low-Level Associations That Are Independent of Object Identity

Patrick Sadil, Kevin W. Potter, David E. Huber, and Rosemary A. Cowell
University of Massachusetts

Knowing the identity of an object can powerfully alter perception. Visual demonstrations of this—such as Gregory’s (1970) hidden Dalmatian—affirm the existence of both top-down and bottom-up processing. We consider a third processing pathway: lateral connections between the parts of an object. Lateral associations are assumed by theories of object processing and hierarchical theories of memory, but little evidence attests to them. If they exist, their effects should be observable even in the absence of object identity knowledge. We employed Continuous Flash Suppression (CFS) while participants studied object images, such that visual details were learned without explicit object identification. At test, lateral associations were probed using a part-to-part matching task. We also tested whether part-whole links were facilitated by prior study using a part-naming task, and included another study condition (Word), in which participants saw only an object’s written name. The key question was whether CFS study (which provided visual information without identity) would better support part-to-part matching (via lateral associations) whereas Word study (which provided identity without the correct visual form) would better support part-naming (via top-down processing). The predicted dissociation was found and confirmed by state-trace analyses. Thus, lateral part-to-part associations were learned and retrieved independently of object identity representations. This establishes novel links between perception and memory, demonstrating that (a) lateral associations at lower levels of the object identification hierarchy exist and contribute to object processing and (b) these associations are learned via rapid, episodic-like mechanisms previously observed for the high-level, arbitrary relations comprising episodic memories.

Keywords: continuous flash suppression, object identification, state-trace analysis, top-down processing, visual learning

Supplemental materials: <http://dx.doi.org/10.1037/xge0000607.suppl>

The hidden figure of a Dalmatian dog (Gregory, 1970; Figure 1a) is often presented to students of cognitive psychology as a powerful demonstration of top-down processing. Upon first viewing this picture, it is perceived as a jumble of spots (Chang, Baria, Flounders, & He, 2016; van Tonder & Ejima, 2000), but after the hidden Dalmatian is outlined, the dog is readily seen and one’s

perception of the image is permanently altered. Undoubtedly, top-down knowledge plays a role in our ability to perceive the Dalmatian more quickly and easily after prior exposure (Gregory, 1970; Marr, 1982; Newen & Vetter, 2017). To shed light on the mechanisms of object perception and visual memory, we consider whether the widely accepted top-down processing explanation

This article was published Online First May 9, 2019.

Patrick Sadil, Kevin W. Potter, David E. Huber, and Rosemary A. Cowell, Department of Psychological and Brain Sciences, University of Massachusetts.

Kevin W. Potter is now at Center for Addiction Medicine, Massachusetts General Hospital.

These results have been archived as a preprint on PsyArXiv (psyarxiv.com/bqp32/), and the first experiment was presented at the 2018 Annual Vision Sciences Society meeting (Cowell, Sadil, Potter, & Huber, 2018) and the 2017 Context and Episodic Memory Symposium (memory.psych.upenn.edu/CEMS_2017). Additionally, the first experiment was discussed in Sadil, Cowell, and Huber (2019). We thank John Dunn for sharing state-trace analysis code from Kalish et al. (2016), Piotr Winkielman for sharing samples of code to produce CFS, and Andrew Heathcote for his review of this study. The work was supported by National Science Foundation (NSF) award BCS-1431147 to David E.

Huber, NSF CAREER Award 1554871 to Rosemary A. Cowell, and National Institutes of Health (NIH) award 1RF1MH114277-01 to Rosemary A. Cowell and David E. Huber. Author Contributions: Conceptualization, Patrick Sadil, David E. Huber, and Rosemary A. Cowell; Methodology, Patrick Sadil, David E. Huber, and Rosemary A. Cowell; Investigation, Patrick Sadil, David E. Huber, and Rosemary A. Cowell; Data Curation, Patrick Sadil; Formal Analysis, Patrick Sadil, Kevin W. Potter, David E. Huber, and Rosemary A. Cowell; Writing—Original Draft, Patrick Sadil, David E. Huber, and Rosemary A. Cowell; Writing—Review & Editing, Patrick Sadil, Kevin W. Potter, David E. Huber, and Rosemary A. Cowell; Supervision and Funding Acquisition David E. Huber and Rosemary A. Cowell.

Correspondence concerning this article should be addressed to Patrick Sadil, Department of Psychological and Brain Sciences, University of Massachusetts, 135 Hicks Way, Amherst, MA 01003. E-mail: psadil@gmail.com




Figure 1. (a) Hidden Dalmatian in park. (b) Schematic illustration of the kinds of associations that might contribute to recognizing the Dalmatian. The traditional role of top-down knowledge is depicted by the solid arrows from object identity to parts. Bottom-up information flow is shown by dashed arrows from parts to object identity. A central question of the present experiment is whether direct, lateral associations between the parts of the object—unmediated by higher-level associations—can also be learned and retrieved, thereby contributing to object identification (dotted arrows).

provides a full account of this phenomenon. We ask whether perception additionally benefits from rapidly learned lateral associations between the separate parts of previously viewed objects. In the case of the Dalmatian, these lateral associations would literally connect the dots to form the outline of the dog. More generally, these lateral associations would form relations between the different parts or features of an object, such as in Geon theory (Biederman, 1987). Furthermore, we ask whether they can be acquired for specific objects with episodic-like rapidity and independently of feedback from higher-level representations.

Figure 1b schematically depicts a simple theoretical framework consisting of two object parts and an object whole, with three types of association: top-down, bottom-up, and lateral. Bottom-up processing is self-evidently necessary for perception, being required to convey sensory information to the higher-order brain regions in which object recognition is carried out. In addition, the role of top-down processing is amply demonstrated by visual illusions in which conceptual, linguistic, or Gestalt information alters low-level perception (Gregory, 1972; Koffka, 1935; Palmer, 1975; Reicher, 1969). Finally, there is a third possible pathway for information flow in object processing: lateral associations between the separate parts of an object (Figure 1b). Such lateral associations are assumed by successful theories of object identification (e.g., Biederman, 1987), and yet there is little direct, unequivocal evidence that they exist. Moreover, if such associations exist, it is unclear whether they are learned via low-level, perceptual mechanisms or instead via episodic-like encoding mechanisms. These questions are key to the dispute between traditional versus hierarchical models of memory (Cowell, Bussey, & Saksida, 2010a; Henke, 2010; Shimamura, 2010; Squire & Zola-Morgan, 1991; Tulving, 1985).

Object identification is thought to result from hierarchical processes and representations: simple perceptual processing of elemental features occurs first, with information about object wholes and the semantic properties of objects being extracted and encoded in subsequent stages (Cowell, Leger, & Serences, 2017; Felleman & Van Essen, 1991; Hubel & Wiesel, 1965; Kobatake & Tanaka,

1994). Most theories of object processing assume that earlier stages are more visual and implicit, whereas later stages are more conceptual and explicit (but see Hochstein & Ahissar, 2002, for a more nuanced hypothesis). In addition, there is clear evidence that associative learning occurs at the highest levels of this hierarchy—levels that are commonly considered postperceptual—such as learning to associate an object with a scene (Ranganath, 2010) or learning that an object was studied in one temporal context but not another (Johnson, Hashtroudi, & Lindsay, 1993). Indeed, this kind of associative learning is critical to the formation of episodic memories, given that episodic events are typically defined by unique conjunctions of place, time and content (Tulving, 1983). However, associative learning at lower, undeniably perceptual levels of the hierarchy (e.g., between the parts of an object) has not been clearly demonstrated to occur without mediation by higher-level representations. Yet in the absence of such a demonstration, we have not shown that low-level lateral associations exist: when higher-level representations are available, any information putatively transmitted via lateral associations might instead travel via a part-to-whole back down to part route, in which top-down feedback plays a key role (see Figure 1b).

This gap in the empirical record is not only relevant to models of perception: lateral, intra-object associations are also predicted by hierarchical theories of memory (Cowell, Bussey, & Saksida, 2006; Cowell et al., 2010a; Shimamura, 2010). Hierarchical memory theories assume, like hierarchical theories of object vision, that processing in the ventral visual pathway unfolds along a continuum from simple visual features to object identity. But hierarchical memory theories make the extra assumption that the medial temporal lobe constitutes an additional, higher-level station in the hierarchy, containing associative, episodic representations (Bussey & Saksida, 2002, 2005; Graham, Barense, & Lee, 2010). Furthermore, such theories predict that intra-object, lateral associations, which reside at intermediate levels, can be acquired and retrieved using the same mechanisms proposed for the higher-level associations comprising episodic memory. That is, whereas traditional theories of memory assume encoding and retrieval processes that

are specialized for episodic events—mechanisms such as binding and recollection that allow rapid formation of associations between disparate elements of an event and later retrieval via part-cued pattern-completion (e.g., Diana, Yonelinas, & Ranganath, 2007; Eichenbaum, Yonelinas, & Ranganath, 2007; Tulving, 1985)—hierarchical memory models claim that cognitive operations (e.g., for encoding and retrieval) are the same at all levels of the hierarchy. This claim entails not only that lateral, intra-object associations exist but that—as for episodic event associations—they can be rapidly learned (e.g., Weinberger, 2007), and their learning and retrieval can occur without top-down support from higher level representations.

To examine whether the rapid learning and subsequent retrieval of intra-item associations can be separated from any top-down influence of higher-level information, we asked whether intra-object learning can be dissociated from object identification. For this, we required one task in which visual object knowledge can be retrieved in the absence of object identification, and another in which object identity must be retrieved. In a part-matching task, participants indicated whether a pair of parts came from the same or different objects. In a naming task, participants attempted to provide the object's name from a visual part cue. In the part-matching task, each part was a circular patch drawn from an object, chosen to be unrecognizable when viewed in isolation, as if viewing the object through an opaque sheet with two punched-out holes. Specifically, participants were given a two-alternative forced choice (2AFC) between a pair of patches (e.g., A_1 and A_2) that came from the same object (the two patches were located in their correct positions with respect to each other), A , versus a pair of patches (e.g., A_1 and B_2) that came from two different objects, A and B . With a common patch in both choices (A_1), the task amounted to identifying which other patch (A_2 vs. B_2) was associated with the common patch. In the Naming task, which was

administered immediately after each 2AFC Part-matching trial, participants were presented with a single circular patch—always the common patch (e.g., A_1)—and asked to name the object from which it was drawn. Analogous to how participants can know which dots in a Mooney image belong to the same object without knowing the identity of that object (van Tonder & Ejima, 2000), the 2AFC part-matching task can be performed intuitively in the absence of object identification (e.g., “those patches just seem to go together”). In contrast, the naming task requires access to object identity knowledge, and the identity can be produced in the absence of lateral associations between separate parts of an object. Thus, to assess whether intra-item associations can be retrieved separately from top-down associations, we tested for a dissociation between 2AFC part-matching and naming.

In addition to needing two retrieval tasks—one assessing top-level information (Naming) and the other assessing intra-item associations (2AFC Part-matching)—the demonstration of a dissociation between top-down and lateral associations requires prior learning conditions that selectively boost the strength of each type of association (Figure 2a). Specifically, the flow of top-down information might be selectively enhanced by studying an object's name in the absence of visual details, whereas lateral associations might be selectively enhanced by studying an object's visual details without awareness of its identity. We therefore combined the two retrieval tasks with two key study conditions: Continuous Flash Suppression (CFS, an experimental technique described below) and Word. The CFS condition allowed us to expose participants to the visual details of the objects while reducing awareness of object identity. The Word condition, in which subjects studied object names with full awareness (i.e., without CFS), provided object identity but no visual details. We also included a No Study baseline condition, and a Binocular Image control condition, in which objects were studied as visual images without




Figure 2. Experimental Design and Predictions. (a) Four study conditions provided different types of information (visual details, object identity, neither or both). CFS presents a stimulus to the nondominant eye only, whereas a flashing mask of colored squares presented to both eyes eliminates or greatly reduces awareness of the stimulus, masking object identity. (b) State-trace plot of predicted benefit of the 4 study conditions to performance on 2 tasks—one testing visual knowledge and the other object identity knowledge. See the online article for the color version of this figure.

CFS, providing visual details as well as full awareness of object identity (these were easily recognized, everyday objects).

Brief, subliminal presentations have been used to produce learning without awareness, but brief durations may be insufficient to support the learning of durable lateral associations. To provide long study durations in the absence of object identification, we developed a novel use of CFS (Tsuchiya & Koch, 2005) as a technique for implicit visual learning. CFS takes advantage of binocular rivalry to block awareness of a visually presented stimulus for up to several seconds. During CFS, different images are presented to each eye. The dominant eye is shown only a highly salient, dynamic mask (e.g., overlapping, flashing squares) whereas the nondominant eye is shown an image of the to-be-learned object superimposed on the mask. For an extended duration, participants remain unaware of the object shown to the nondominant eye and are aware of only the masking images. Thus, participants may be able to study the visual details of the masked object without attaching a verbal identity to those details.

In our experiment, CFS was predicted to produce learning of lateral, intra-item associations, but provide less learning (if any) between object parts and object identity (Figure 1b). Therefore, study of objects masked by CFS should boost 2AFC Part-matching performance with little or no benefit to the Naming task. In the complementary Word study condition, participants studied with full awareness only the word-name of each object. This condition should support learning of object identity, which may influence subsequent test performance via top-down associations (i.e., from Object Identity to Part 1 or Part 2, Figure 1b), to the extent that the test materials (i.e., Part 1 and Part 2) contain visual information that is typically associated with this identity. However, the Word condition should not support learning of lateral visual associations, because there is no viewing of the to-be-tested object parts. Therefore, word study should boost naming performance with little or no boost to performance in the 2AFC part-matching task.

As outlined above, this pattern of predicted results constitutes a double dissociation (i.e., one manipulation selectively affects one performance measure, whereas a different manipulation selectively affects a different performance measure). Such a finding would support the conclusion that lateral, intra-object associations can be learned separately from top-down influence of object identity information. However, this support would not constitute unequivocal evidence that two distinct types of association (lateral, top-down) exist—double dissociations can arise as an artifact of comparing two performance measures that differ in sensitivity to a single underlying representation (Bamber, 1979; Dunn & Kirsner, 1988; Loftus, 1978; Newell & Dunn, 2008; Wagenmakers, Krypotos, Criss, & Iverson, 2012). To address this concern, we performed a state-trace analysis (e.g., Bamber, 1979; Kalish, Dunn, Burdakov, & Sysoev, 2016). A state-trace analysis determines whether more than one latent variable contributes to a set of behavioral findings in a manner that is not rendered ambiguous by possible differences in task sensitivity. A state-trace analysis therefore provides a principled way to determine whether the different kinds of prior study—for example, CFS, Word—boost at least two latent variables (lateral associations, top-down associations). The No Study and Binocular conditions were included to provide constraint for the state trace analysis.

To summarize, we hypothesized that lateral, intra-item associations can be rapidly learned and subsequently retrieved, indepen-

dently from the higher-level representations required for object identification. To test this hypothesis, we used two tasks (2AFC Part-matching and Naming) thought to rely on two different kinds of associations (see Figure 2). Because 2AFC part-matching does not require explicit identification of the object, intra-item visual associations can support this measure, and learning under CFS should preferentially boost 2AFC part-matching relative to naming. In contrast, studying an object's name does not provide visual details of the object to be encountered at test, but does supply information about object identity. Thus, studying the name alone under full awareness (Word condition) should boost naming more than 2AFC part-matching. Finally, because binocular study provides both the name and the visual details, it is likely to boost both naming and 2AFC part-matching. Performance in these conditions is compared with performance on objects that are novel at the time of test (No Study).

To foreshadow the results, participants in this study were able to rapidly learn and subsequently retrieve lateral, intra-item associations that were distinct from object-identity information. To assess the reliability of these results, we replicated the experiment, which bolstered the evidence in favor of this inference. These findings expand the inventory of paths contributing to perception beyond bottom-up and top-down processing to include lateral, part-to-part associations and provide support for hierarchical theories of memory.

Method

Participants

We aimed for a minimum of 1,000 data points per condition to achieve sufficient precision in the parameter estimates of a statistical Bayesian model. Participants were recruited through the University of Massachusetts, Amherst's SONA account and campus fliers. In total, 52 participants completed the original experiment, and 62 participants completed the replication. Participants provided written, informed consent in accordance with the University of Massachusetts Amherst Institutional Review Board and were compensated either with course credit or at a rate of \$10 per hour.

Participants were excluded in both data sets if they did not exhibit above chance performance for objects that were studied Binocularly (i.e., images of objects that were studied with full awareness) on either of the 2AFC or the naming tasks, as determined by a binomial test with an alpha level of 0.05. Chance was defined for the 2AFC task as 50%, and for the naming task as $1/N$, where N was the number of objects per condition (32). This translates to participants being excluded if they did not score more than 20 of 32 on 2AFC trials or more than two of 32 on Naming trials. These criteria excluded four participants from the original dataset (leaving 48) and five from the replication (leaving 57).

Materials

Images of 144 unique, everyday objects (e.g., chair, couch, monitor) were gathered from the Internet. Images of objects were presented in gray-scale, cropped, and resized (maintaining the aspect ratio of the original object) such that the longest cardinal axis (horizontal or vertical) of the object extended to 295 pixels. Objects were superimposed upon a gray (RGB values: [161, 161,

161]) background of 300×300 pixels and framed by a 150 pixel white border. They were displayed on a 24" LED monitor at a resolution of 1920×1080 (120 Hz refresh rate).

Object patches were created using three aperture views (circles of radius 80 pixels) of each object. The view through an aperture could potentially include some background as well as part of the object. The choice of aperture location was constrained such that the ratio of object to background visible through each aperture was no less than the ratio of the object to the background in the entire 300×300 pixel image. Fixed pairings of the objects were determined in advance (creating 77 pairs from the 144 objects), with each pairing chosen to yield a pair of mismatched apertures (i.e., apertures from two different objects to be used as the foil stimulus in the 2AFC task) that were not obviously mismatched. Each pair of objects was used to create two different intact-rearranged 2AFC trials, with both trials created from a given object pair being assigned to the same study condition. For instance, object *A* had three apertures (A_1 , A_2 , and A_3) and object *B* had three apertures (B_1 , B_2 , and B_3). From these apertures, two different 2AFC Part-matching trials were created: $A_1 - A_2$ (intact) versus $A_1 - B_3$ (rearranged), and $B_1 - B_2$ versus $B_1 - A_3$. Both trials were then assigned to the same condition (e.g., Binocular).

Mondrian masks for CFS were constructed from overlapping, colored rectangles. The rectangles were presented in a cropped, 300×300 pixel square, centered on the presentation screen. The hue, saturation, and value (HSV space) of the colors of the squares were determined by uniformly random samples between [1/6, 1]. The width and height of the rectangles were randomly chosen to be between 15 and 40 pixels, and 1,000 such rectangles were drawn for each mask.

Procedure

Experimental code was written with MATLAB (MathWorks, 2015) scripts using the Psychophysics Toolbox (Brainard, 1997; Kleiner et al., 2007). Prior to the main experiment, ocular dominance of each participant was measured via the Porta test (as performed by Roth, Lora, & Heilman, 2002). Figure 3 presents an overview of the main experimental procedure.

The experiment contained nine blocks, each with a study phase and test phase. Participants encountered 16 objects per block (eight pairs of objects). Each block contained 12 study trials and 16 test trials (there were only 12 study trials because 4 objects in each block were assigned to "No Study"). Each study trial presented just one object. The two objects in each pair were assigned to the same study condition (or both were not studied). The assignment of object pairs to blocks was fixed across participants, but the study condition for each object pair was randomly determined for each participant, with the constraint that each block contain 4 test trials for each of the 4 conditions. The study conditions were (a) No prior study, (b) Word (the name of the object was presented binocularly in text form, with no image), (c) CFS (the image of the object was presented, masked by CFS), and (d) Binocular (the image of the object was presented without masking).

Study Phase

During the study phase of each block, participants saw 12 objects, one at a time (four from each of the three conditions that

contained prior study, with the remaining four objects in the list of 16 being assigned to "No Study"). The order of the 12 presentations was pseudo-random, disallowing more than three presentations of the same study condition in a row. Participants studied the 12 items in the same order twice, with the second presentation of the list of items immediately following the first, in the same order as the first. Each study trial began with a 500-ms fixation square, followed by a CFS mask (presented binocularly), with a 150-pixel white border. On Word trials, the name of the object gradually appeared in front of the mask, presented to both eyes. On Binocular trials, the image of the object gradually appeared in front of the mask, presented to both eyes. On CFS trials, the image of the object gradually appeared in front of the mask for the participant's nondominant eye, whereas the dominant eye continued to view only masks.

Gradual appearance of the stimulus in Word and CFS trials was accomplished by changing the transparency of the stimulus. The object increased from 0% to 100% opacity over the first 1000 ms of a study trial, in 10 evenly spaced steps of 100 ms that were synchronized with the 100 ms cycles of the Mondrian mask. Opacity increased in linear steps, that is, by an equal amount at each step (i.e., 10% approximately every 100 ms).


Presentation of different images to the different eyes was accomplished by using NVidia 3D shutter glasses, which synchronize with the display monitor to allow only one eye to see through the glasses on a given refresh cycle; by alternating between two displays on each refresh cycle, each eye views a different display. On all trials, once the object or word reached full opacity it was presented for 2,500 ms, and then gradually decreased to full transparency over 500 ms.

Throughout study trials, the Mondrian masks changed every 100 ms (i.e., at 10 Hz), by altering the position, size, and color of the rectangles in the mask. A set of masks was randomly created for each participant, and this same set was used for all trials for that participant, with each trial presenting the masks in the same order. Each 100 ms duration mask was presented only once per trial.

At the start of each CFS study trial, participants were instructed to press a button as soon as they detected that an object was present. They were instructed to give this response if they were confident that there was an object, but before they could identify the object. As soon as they gave this response, the trial ceased so as to minimize learning with awareness. To encourage fast responses, they were shown their reaction time (RT). To encourage accuracy, one-third of study trials were catch trials, in which no object appeared. If they responded on catch trials, participants saw the message: "CAREFUL! No object appeared." This self-termination of study on CFS trials was included for two reasons: first, to provide the maximum opportunity for learning with limited awareness and, second, control the degree of awareness. Regarding the second rationale, a fixed duration of study would likely result in a mixture of trials, with some experiencing full breakthrough from CFS and thus full awareness of object identity, whereas other trials would be completely without any awareness that anything was shown. By allowing participants to self-terminate study at the first stage of breakthrough (i.e., when they first started to become aware that an object was present), the aim was to place all study trials at the same level of limited awareness.

At the start of Word trials, participants were instructed to "please imagine the following object in detail, as if it were pre-

A. Overall Task Structure



B. Example Study Trial




Figure 3. Schematic of the experiment. (A) Overall task structure. The experiment was organized into nine blocks, of which the first served as a practice. Each block contained a study phase (in which a list of 12 items was presented, twice) immediately followed by a test phase (containing 16 test trials). In the Test Phase, each trial comprised a 2AFC part-matching task followed by a part-to-name task. Participants completed 16 test trials (one for each item from the Study Phase, plus four items from the No Study condition), before moving to the next study-test block. (B) Example study trial. On each study trial, participants saw an item for up to 4 seconds, then provided a Perceptual Awareness Scale (PAS) rating, and finally answered a question designed encourage attention to the details of the object. See the online article for the color version of this figure.

sented over the flashing squares.” At the start of binocular trials, participants were instructed “please study the details of the following object.”

At the end of every study trial for all study conditions, participants were asked two questions about the object. First, they gave a four-valued Perceptual Awareness Scale (PAS; Ramsøy & Overgaard, 2004) rating, with values 0–3. Participants were instructed to use the scale as follows: “If you CLEARLY SAW something besides the squares, AND COULD NAME IT, answer 3,” “If you DEFINITELY saw something, but are unsure what (though you might be able to guess), answer 2,” “If you only POSSIBLY saw

something, but COULDN’T accurately say what it was, answer 1,” and “If you didn’t see anything besides the squares, answer 0.” These instructions were presented at the beginning of the experiment, and a reminder was presented on all study trials. The second question was included to encourage attention to the visual details of each object (in the Word condition, participants were instructed to respond based on *imagined* visual details of the object). The question was one of four randomly assigned questions: (a) “Was the object symmetric across its horizontal axis?,” (b) “Did the object fill more than one-quarter of the flashing squares?,” (c) “Was the object reflective?,” or (d) “Did the object contain multiple parts?”

Test Phase

Each test phase occurred immediately after the corresponding study phase. The test phase was self-paced, and the 16 items were presented in random order. On each test trial, participants were first asked to choose which of two pairs of parts came from the same object. Immediately after each 2AFC Part-matching test, the Naming test for the corresponding object occurred (i.e., the two test formats were interleaved, trial-by-trial). Specifically, the part that was common to both forced-choice options (A_1) reappeared by itself, and the participant was asked to name the corresponding object. In both tasks, participants were encouraged to “use your memory from the items that you studied, if that helps.”

There was accuracy feedback after each 2AFC trial, designed to keep people on task, and to let them appreciate that the correct responses could be made even in the absence of explicit knowledge for the identity of the object. No feedback was provided after naming responses. After analyzing the results from the initial experiment, a programming error was discovered, which had led to a failure to randomly counterbalance the side on which the correct 2AFC pair of parts was presented in the No Study condition. More specifically, just for this one condition, the correct pair was always on the left, although the trial-by-trial feedback for this condition was based on a randomly determined side. Because of this random (inaccurate) feedback, it is unlikely that any participants realized that the correct answer was always on the left for this one condition (and, furthermore, knowledge of how to use this information to benefit performance would require knowing that the currently tested object was not previously studied, which in itself would require identification of the object, negating the need for a guessing strategy). This programming error should not have affected the other conditions and, moreover, the baseline No Study condition is the least important in terms of establishing a dissociation. However, because this programming error may have affected performance in the No Study condition (and, more generally, to better assess the reliability of our initial results), a direct replication of the experiment was performed using experimental software that did not contain this error.

For all participants, the first block was used as a practice. The practice block was excluded from all analyses. The entire experiment lasted approximately 2 hr.

Results

Data were analyzed in three ways. First, we tested for statistically significant differences in performance between the critical conditions (i.e., CFS vs. Word, treating Task—i.e., 2AFC vs. Naming—as a two-level fixed-effects factor) using linear mixed modeling (Baayen, Davidson, & Bates, 2008; Bates, Mächler, Bolker, & Walker, 2015). Second, another frequentist technique was used to conduct a State-Trace Analysis of our results (Kalish et al., 2016). Finally, a recently developed hierarchical Bayesian State-Trace Analysis was used to calculate a pseudo-Bayes Factor (Sadil, Cowell, & Huber, 2019).

To assess the effect of study condition on 2AFC part-matching and naming performance, a linear mixed model was fit to the data from the initial experiment ($n = 48$). We used the lme4 (Bates et al., 2015) R (R Development Core Team, 2016) package, treating condition and task as fixed-effects and giving each participant a random-effect intercept for each task (random-effects were al-

lowed to be correlated between tasks). Parametric bootstraps were used to test for a main effect of condition in each task using the R package pbkrtest (Davison & Hinkley, 1997, Chapter 4; Halekoh & Højsgaard, 2014). The bootstraps were conducted by generating 1,000 samples from a reduced model (single intercept for each task), comparing the log-likelihood of the full model to the reduced model on each simulated dataset, and comparing this bootstrapped distribution of log-likelihoods to the observed difference. This analysis revealed that there was a main effect of condition, across both tasks ($\chi^2 = 347.2$, difference in $df = 6$, $p < .001$). Post hoc comparisons assuming asymptotic degrees of freedom revealed that average performance in the Word and CFS conditions were different from each other in both the 2AFC Part-matching and Naming tasks, but that the order was reversed for the two tasks. The average 2AFC Part-matching accuracy for an object studied as a Word was 0.74 ($SE = 0.014$), which was lower than the average 2AFC accuracy for an object studied under CFS ($M = 0.79$, $SE = 0.014$; $z = 3.1$, $p < .01$). In contrast, objects studied as a Word were more often named ($M = 0.41$, $SE = 0.016$) as compared with objects studied under CFS ($M = 0.38$, $SE = 0.016$), though that difference was only marginally significant ($z = 1.8$, $p < .07$). Finally, and most importantly, the difference in performance between objects studied as Words and objects studied under CFS was significantly different for the two tasks ($M = 0.078$, $SE = 0.022$; $z = 3.5$, $p < .001$). This dissociation between the effects of studying the identity of an object and studying its visual details suggests that the study conditions allowed learning of two distinct kinds of information. Furthermore, these two distinct kinds of information cannot be top-down (e.g., learned during Word study) versus bottom-up (learned from visual, CFS study) in the absence of lateral associations. Without lateral associations, if CFS study benefitted part-matching performance by enhancing bottom-up processing, that benefit could only be realized via the top (i.e., Part 1 $\rightarrow$ whole object $\rightarrow$ Part 2, Figure 1b), in which case naming performance should have been equally enhanced.

However, as mentioned above, dissociations in performance are ambiguous as to how many cognitive systems are required to account for the data (Bamber, 1979; Dunn & Kirsner, 1988; Loftus, 1978; Newell & Dunn, 2008; Wagenmakers et al., 2012). Therefore, we conducted a state-trace analysis to ask whether lateral, part-to-part associations were dissociable from top-down associations between object identity and visual parts.

We used the state trace analysis technique developed by Kalish et al. (2016) to test for a dissociation. This procedure first finds the best-fitting, single-latent-variable model. The model assumes that performance on the two tasks (i.e., 2AFC part-matching and naming) depends on a single kind of representation (e.g., memory strength), in which case the rank ordering of conditions is required to be the same for both dependent measures. Next, the data are fit with a multi-latent-variable model, which allows for different rank orders for each dependent measure. This latter model corresponds to a model in which lateral associations make a distinct contribution to task performance as compared with top-down associations. Finally, the fits of these two models are compared using bootstrap resampling (Wagenmakers, Ratcliff, Gomez, & Iverson, 2004), which penalizes a model for being too flexible and thus fitting noise. In this case, the multi-latent-variable model is more flexible. If the single-variable model results in a significant loss of fit, then a single-representation account of the data can be rejected (Kalish

et al., 2016). Unlike a traditional double dissociation, such a demonstration of dissociability would be free from concern over whether the two dependent variables are equally sensitive to the latent variable(s) of interest. For this initial experiment ($n = 48$), the single-latent-variable model fit significantly worse ($p < .01$; Figure 4a), indicating that performance on the 2AFC Part-matching and Naming tasks depended on more than one latent variable. This supports our hypothesis that lateral associations can be learned and retrieved separately from object identity.

As with most other techniques designed to assess a state trace analysis statistically, the technique developed by Kalish et al. (2016) assumes independence between the two measures. However, this assumption may be invalid. For example, participants may devote more effort to one task or the other on each trial, with the task that is assigned greater effort differing from trial to trial. In this case, there would be a negative dependence between the two memory measures, rather than the assumed independence. In contrast, item-effects and subject-effects are likely to impose a positive dependence (e.g., some items are easier for both tasks and some people perform better at both tasks). Accounting for such dependencies between measures is important because strong dependencies between the two measures can bias the state-trace results (Sadil, Cowell, & Huber, 2019). To address this, we applied a hierarchical-Bayesian model that estimated trial-level dependence between Naming and 2AFC Part-matching in each condition, as well as correlations between the two performance measures at the level of subject-effects and item-effects.

This hierarchical-Bayesian state trace analysis estimates the probability of each rank order of conditions under consideration, given the data that were observed. If there is only one latent variable underlying performance, the two performance measures must follow the same rank order across conditions. This produces a monotonic function when one performance measure is plotted against the other (i.e., in the state trace). If

there are two latent variables, the rank orders can be different and the function may be non-monotonic. The model considered plausible monotonic and nonmonotonic orders of the four conditions across the dependent measures (see the [online supplemental materials](#) and Sadil, Cowell, & Huber, 2019). Orders were yoked across the two dependent measures such that tonicity of a given order was determined by the relationship between the two (e.g., the order No Study < Word < CFS < Binocular for 2AFC accuracy paired with the order No Study < CFS < Word < Binocular for Naming accuracy produces a nonmonotonic order). Given the study conditions, only the four orders in which performance in the No Study condition was the worst and performance in the Binocular condition was the best were deemed plausible (see the [online supplemental materials](#)). The posterior probabilities of each order can be combined into probability odds to give a pseudo-Bayes factor to assess evidence for or against each order (Geisser & Eddy, 1979; Gelfand, Dey, & Chang, 1992; Kass & Raftery, 1995; Morey, Romeijn, & Rouder, 2016; Yao, Vehtari, Simpson, & Gelman, 2018; [online supplemental materials](#)). In these data, the pseudo-Bayes Factor indicated very strong evidence in support of just a single order: the one that matched the condition means, that is, No Study < Word < CFS < Binocular for 2AFC accuracy along with No Study < CFS < Word < Binocular for Naming accuracy. The odds of this winning order to the other three orders were 99:1. Because the winning order is nonmonotonic, the logic of state-trace analysis dictates that it must have been produced by a model with more than one latent dimension. Hence, this analysis, which relaxes assumptions inherent to other state-trace analyses, also provides evidence that participants used two distinct kinds of associations in this experiment. We claim that these two kinds are top-down associations versus novel, rapidly learned lateral associations.

To better estimate the reliability of our inferences, and in light of the failure to randomly counterbalance left/right side during 2AFC trials for the No Study condition (see Methods), the experiment was replicated (Figure 4B). The replication data appear less nonmonotonic than the original experiment (see Figure 4). Indeed, when the Kalish et al. method was applied the null hypothesis of monotonicity was not rejected ($p = .14$). However, a statistically significant classic dissociation was observed: the difference of differences in performance on objects studied as Words versus objects studied under CFS between the 2AFC and Naming tasks was 0.045 ($SE = 0.02$, $z = 2.17$, $p < .03$). It is therefore unclear from these frequentist tests alone whether the replication provides evidence for or against the nonmonotonicity of the data (especially given that a failure to reject the null hypothesis does not constitute evidence for the null hypothesis—a limitation of statistics based on p -values; e.g., Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010). The combined data (i.e., treating this exact replication as an addendum to the original, producing a single dataset with over twice as many participants as in the original experiment) again allowed rejection of monotonicity using the Kalish et al. technique ($p = .01$). To circumvent the limitations of null hypothesis significance testing and instead assess the evidence for or against monotonicity within the combined da-




Figure 4. Average performance on the two tasks in the four conditions, for both original (left) and replication (right) experiments. Error bars indicate within-subject confidence intervals, including correction by (Morey, 2008). See the online article for the color version of this figure.

taset, we returned to the hierarchical Bayesian State-Trace analysis.¹

Bayes Factors provide a principled way of combining evidence across data sets (e.g., Kass & Raftery, 1995; Morey et al., 2016), and pseudo-Bayes Factors can be used in a similar way. Interpreted as the posterior odds of one hypothesis making better predictions than another, the odds from two experiments can be multiplied to update the evidence in favor of a hypothesis with each new source of data. In the replication dataset, the odds in favor of the same nonmonotonic order that was preferred in the original dataset were 4.75:1. Multiplying the two pseudo-Bayes Factors (or odds) together gives 469.97:1, which is very strong evidence in favor of the nonmonotonically arranged data (Kass & Raftery, 1995), and therefore very strong evidence that performance in the 2AFC and Naming tasks was supported by at least two kinds of information.

In the CFS condition, participants were instructed to self-terminate study trials when they first detected that an object was appearing, but before they could identify the object. This likely resulted in some encoding of object identity on a subset of CFS trials, specifically those trials in which termination occurred too late. Thus, although it is assuredly true that study under CFS resulted in less awareness of the object's identity than binocular study, it is unlikely that CFS study eliminated object identity awareness entirely. Critically, however, our inferences do not require that CFS prevent all such awareness. Instead, our conclusion is reached by observing a dissociation between the two dependent measures, when the relative amounts of object identity and visual object information are manipulated. We do not claim that performance in any individual condition was supported by a single kind of information (e.g., object identity information in the Word condition or lateral associations in the CFS condition), but we can infer, using state trace analyses, that more than one kind of information contributed to performance across the four conditions. These considerations notwithstanding, we repeated the analysis excluding all 2AFC and Naming trials in which participants provided a PAS rating of 3 during either of the two study presentations. Reassuringly, the results of the hierarchical Bayesian analysis were largely unchanged. That is, the data without these PAS 3 trials yielded a preference for the same nonmonotonic order, with a pseudo Bayes Factor of 54.5 for the original dataset, 5.1 for the replication, and a combined pseudo Bayes Factor of 278.1. This again constitutes very strong evidence in favor of nonmonotonically arranged data (Kass & Raftery, 1995).

Discussion

Our key finding is that representations of intra-item, part-to-part associations can be learned and retrieved separately from the representation of a visual object's identity. In other words, we can rapidly learn, from brief exposure, to connect the dots of a visual object without using knowledge of what the object is. Returning to the hidden Dalmatian, our results suggest that prior experience with this picture makes it easy to see the dog not only because of top-down expectations to find a Dalmatian in the otherwise unrecognizable pattern, but also because we have learned lateral associations between the visual parts of the image (we have connected the dots).

A large literature attests to the importance of top-down processes in vision, across a range of tasks and stimulus types (e.g.,

Gregory, 1972; Palmer, 1975; Pinto, van der Leij, Sligte, Lamme, & Scholte, 2013; Reicher, 1969; Rubin, 1915). In line with this, many models of visual processing contain both bottom-up and top-down mechanisms (e.g., Graboi & Lisman, 2003; Marr, 1982; Rutishauser, Walther, Koch, & Perona, 2013). Furthermore, the distinction between bottom-up and top-down aspects of cognition finds support beyond the vision literature, in memory research: low-level, implicit mechanisms of learning and retrieval have been dissociated from high-level, explicit mechanisms by comparing, for example, repetition priming with conscious recall (Cave & Squire, 1992). Broadly speaking, such dissociations have been interpreted as evidence for separable contributions of bottom-up (automatic, unconscious, implicit) versus top-down (intentional, effortful, explicit) routes to memory (Squire & Zola-Morgan, 1991; Tulving & Schacter, 1990). But, to our knowledge, this is the first study to demonstrate that intra-item, lateral connections can be learned and retrieved independently of top-down feedback and, thus, that they exist separately from representations at higher levels. This provides evidence for a third, lateral pathway in object processing, complementing the classic bottom-up and top-down routes.

Hierarchical models of memory posit that processing in the ventral visual pathway and adjacent medial temporal lobe unfolds along a continuum, with early processing of simple visual features giving way to later processing of object identity, semantics and associative, episodic representations (Cowell et al., 2010a; Kent, Hvoslef-Eide, Saksida, & Bussey, 2016; Tulving & Schacter, 1990). Evidence for the learning and retrieval of lateral associations at the highest levels of the hierarchy is abundant: arbitrary, interitem associations are the stuff of episodic memories, and empirical demonstrations of their existence are innumerable. These high-level associations are predicted by hierarchical models and indeed most current classes of memory theory. However, hierarchical memory theories make a further claim: that the same processes of encoding and retrieval that support episodic memories should, in principle, occur at all levels of the hierarchy from visual cortex through to the medial temporal lobe (e.g., Cowell, Bussey, & Saksida, 2010b). Therefore, it is the unique claim of hierarchical theories that rapidly acquired lateral associations exist not only at a high (episodic) level but at all levels (including perceptual). To date, the question of whether analogous lateral associations exist between relatively low-level, perceptual elements within an object has remained unanswered.

Prior studies have found evidence of intra-item learning (Amano, Shibata, Kawato, Sasaki, & Watanabe, 2016; Biederman & Cooper, 1991; Diana, Yonelinas, & Ranganath, 2008; Shibata, Watanabe, Sasaki, & Kawato, 2011; Yonelinas, 2013), but these studies did not determine whether this reflected lateral associations versus top-down feedback. To assess the role of top-down feedback, we compared fully aware study with study with limited awareness under CFS, thereby manipulating the degree to which participants were aware of high-level object identity. The logic of

¹ We note that the state-trace analysis techniques developed by Davis-Stober, Morey, Gretton, and Heathcote (2016) and Prince, Brown, and Heathcote (2012) produce Bayes Factors. One of these approaches could also be used to assess the reliability of the nonmonotonicity observed here in a Bayesian manner. However, owing to those techniques' inability to explicitly model correlations between the two tasks, we did not use them here.

the study was that if participants were not aware of the object identity at study, there could be no opportunity to strengthen bottom-up and top-down associations between the object parts and the object identity; thus, any increased ability to match the object parts despite a failure to name the object must reflect part-to-part lateral associations. Of course, in our study, as in all studies using CFS or subliminal presentation, we may not have eliminated awareness of object identity entirely, so we have not unequivocally demonstrated learning of lateral associations in the *complete absence* of learning of object identity. However, our claim is not that the learning in our study is implicit, but rather that there exist part-to-part representations that are separate from part-to-whole representations. Because we make no claims about the implicitness of the learned representations, our conclusions do not require a demonstration that learning was fully implicit. Instead, our conclusions depend only on the assumption that the relative amounts of verbal object-identity information and visual information are manipulated by our different conditions (i.e., CFS and Word). Therefore, even without demonstrating a complete lack of awareness of object identity, the logic of the study still applies: by manipulating the availability of object identity versus visual information in the two key conditions (CFS, Word), we can ask whether the resultant data indicate separability between bottom-up/top-down versus lateral associations.

In the literature, the paradigm most similar to the novel experimental design we present is the word-stem completion task as tested with medial temporal lobe amnesic patients. When patients studied lists of words and were subsequently asked to complete three-letter word stems, they completed these stems with a previously studied word much more often than would be expected without study, and at a comparable level with controls, despite being unable to recall those words explicitly (Graf & Schacter, 1985; Graf, Squire, & Mandler, 1984). However, this paradigm differs from the paradigm we have developed in a key respect: it did not test for lateral, part-to-part associations. Although the stem-completion and explicit recall tasks presented to patients may appear analogous to our part-to-part and part-to-name tasks, respectively, this is not quite the case. The stem-completion task can be achieved via strengthened bottom-up and top-down associations from the letters to the word, and thus amounts to a part-to-whole test, analogous to our part-to-name task. In turn, the explicit recall tasks depended on associations between the to-be-recalled word and a cue extrinsic to that word (recall was cued by, e.g., the context of the study episode, or by another word previously paired with the target word). Thus, the explicit recall tasks tested whole-to-whole or whole-to-context retrieval. In sum, because no task in these patient studies probed learning of associations between the letters of a word, they did not establish whether lateral, part-to-part associations exist separately from bottom-up or top-down, part-whole associations.

Related to the present findings, some studies have shown that novel, rapidly acquired associations can be learned and retrieved implicitly, for example, in healthy participants exposed to subliminal study presentations (Duss, Oggier, Reber, & Henke, 2011; Reber & Henke, 2011). However, these studies examined the learning of unrelated word-word or arbitrary word-face pairs, which correspond to associations in the highest levels of the hierarchy rather than at lower-level, traditionally perceptual stages. In line with the notion that such arbitrary associations reside at

higher levels of the hierarchy (but challenging the notion that the neural substrates of declarative and nondeclarative memory are distinct) this learning was linked to high-level brain regions in the medial temporal lobe (Degonda et al., 2005; Henke, 2010; Henke et al., 2003). We note that others have attributed similar implicit associative memory effects to the learning of links between perceptual features of the paired stimuli, residing at earlier, more sensory levels of processing (e.g., Cohen, Poldrack, & Eichenbaum, 1997; Goshen-Gottstein & Moscovitch, 1995). In our view, this account of the data is plausible but not confirmed, because these studies did not dissociate learning at the level of word-word associations from learning at the perceptual-feature level. That is, these prior studies did not adjudicate between high-level versus low-level associative learning accounts of their findings.

In addition to advancing theory by demonstrating the existence of lateral perceptual associations that are separate from top-down processing, our study makes a methodological advance: We have demonstrated that CFS can be used to provide long duration study exposures (on the order of seconds) that support measurable learning despite limited awareness during study. This novel experimental paradigm provides a means of testing hypotheses about implicit learning for complex visual stimuli. As compared with methods for inducing learning without awareness through brief subliminal presentation (e.g., the use of multiple, binocularly masked brief presentations by Duss et al., 2011; Reber & Henke, 2011), it may be that the longer durations afforded by CFS study were critical to the formation of these lateral associations. Relatedly, it is known that different types of masks can be used to manipulate the kind of information available for encoding (Cohen, Nakayama, Konkle, Stantić, & Alvarez, 2015; Gelbard-Sagiv, Faivre, Mudrik, & Koch, 2016; Yang & Blake, 2012). Our results suggest that CFS masking allows sufficient encoding to enable lateral, intra-item associations, while limiting the learning of bottom-up or top-down associations that involve object identity.

To summarize, our study used CFS to limit access to an object's identity at the time of study, reducing the learning of bottom-up and top-down associations between visual features and object identity knowledge. At test, we measured performance on two measures of memory retrieval, one that assessed lateral associations (2AFC Part-matching) and another that assessed the extent to which top-down associations played a role in retrieval (Naming). Using conventional statistics, we found a dissociation in which CFS study boosted 2AFC part-matching more than naming whereas word study boosted naming more than 2AFC. However, such dissociations can be an artifact of using two dependent measures that differ in sensitivity. We ruled out this possibility by using a state-trace analysis. Furthermore, we ruled out potential concerns over unacknowledged dependencies in the state-trace results by applying a hierarchical Bayesian model to assess the trial-level dependence relationship between the two measures of memory. Together, these results indicate that more than one source of information underlies the two memory measures, supporting the claim that lateral, intra-item associations can be learned and retrieved separately from top-down associations. Thus, our study establishes new links between the perception and memory literatures, demonstrating (a) that lateral associations within lower, perceptual levels of the object identification hierarchy exist and contribute to object processing, and (b) that these associations are

acquired via rapid, episodic-like learning mechanisms previously observed for the high-level, arbitrary relations comprising episodic memories.

Context of the Research

Our broad hypothesis is that visual cognition and declarative memory (in the ventral stream and medial temporal lobe) are underpinned by a hierarchical continuum of representations, rather than by a set of anatomical modules specialized for distinct cognitive functions. Under this hypothesis, every stage in the hierarchy from early visual cortex to hippocampus may contribute to any cognitive task, provided its representations are useful for the task. Moreover, every stage can perform the same cognitive operations (e.g., perceptual discrimination, recollection) upon the representations it contains. To test this, we take the kinds of tasks typically used to tap cognition at the top of the hierarchy (e.g., cued recall, associative recognition) and redesign them using stimuli that tap lower levels, asking whether cognitive processes previously ascribed to only the highest levels emerge at lower levels when the stimuli are appropriately adjusted. We recently demonstrated that object recall recruits perirhinal cortex but not hippocampus, challenging the notion that recollection is an exclusively hippocampal mechanism (Ross, Sadil, Wilson, & Cowell, 2018). Here, we show that associative learning occurs not just at the highest, traditionally mnemonic levels, but at lower, perceptual levels, too. Next, we will test whether these visual associations enable visual recollection—a hypothesized pattern-completion-like retrieval process, analogous to traditionally defined recollection of episodic memories, but which occurs at cognitively and neuroanatomically earlier stages than hippocampal recollection.

References

- Akaike, H. (1978). On the likelihood of a time series model. *The Statistician*, 27(3/4), 217. <http://dx.doi.org/10.2307/2988185>
- Amano, K., Shibata, K., Kawato, M., Sasaki, Y., & Watanabe, T. (2016). Learning to associate orientation with color in early visual areas by associative decoded fMRI neurofeedback. *Current Biology*, 26, 1861–1866. <http://dx.doi.org/10.1016/j.cub.2016.05.014>
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390–412. <http://dx.doi.org/10.1016/j.jml.2007.12.005>
- Bamber, D. (1979). State-trace analysis: A method of testing simple theories of causation. *Journal of Mathematical Psychology*, 19, 137–181. [http://dx.doi.org/10.1016/0022-2496\(79\)90016-6](http://dx.doi.org/10.1016/0022-2496(79)90016-6)
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*. Advance online publication. <http://dx.doi.org/10.18637/jss.v067.i01>
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115–147. <http://dx.doi.org/10.1037/0033-295X.94.2.115>
- Biederman, I., & Cooper, E. E. (1991). Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393–419. [http://dx.doi.org/10.1016/0010-0285\(91\)90014-F](http://dx.doi.org/10.1016/0010-0285(91)90014-F)
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436. <http://dx.doi.org/10.1163/156856897X00357>
- Burdakov, O., Sysoev, O., Grimvall, A., & Hussian, M. (2006). An $O(n^2)$ algorithm for isotonic regression. In G. Di Pillo & R. Massimo (Eds.), *Large-scale nonlinear optimization* (pp. 25–33). New York, NY: Springer.
- Bussey, T. J., & Saksida, L. M. (2002). The organization of visual object representations: A connectionist model of effects of lesions in perirhinal cortex. *European Journal of Neuroscience*, 15, 355–364. <http://dx.doi.org/10.1046/j.0953-816x.2001.01850.x>
- Bussey, T. J., & Saksida, L. M. (2005). Object memory and perception in the medial temporal lobe: An alternative approach. *Current Opinion in Neurobiology*, 15, 730–737. <http://dx.doi.org/10.1016/j.conb.2005.10.014>
- Cave, C. B., & Squire, L. R. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 509–520. <http://dx.doi.org/10.1037/0278-7393.18.3.509>
- Chang, R., Baria, A. T., Flounders, M. W., & He, B. J. (2016). Unconsciously elicited perceptual prior. *Neuroscience of Consciousness*. Advance online publication. <http://dx.doi.org/10.1093/nc/niw008>
- Cohen, M. A., Nakayama, K., Konkle, T., Stantić, M., & Alvarez, G. A. (2015). Visual awareness is limited by the representational architecture of the visual system. *Journal of Cognitive Neuroscience*, 27, 2240–2252. http://dx.doi.org/10.1162/jocn_a_00855
- Cohen, N. J., Poldrack, R. A., & Eichenbaum, H. (1997). Memory for items and memory for relations in the procedural/declarative memory framework. *Memory*, 5(1–2), 131–178. <http://dx.doi.org/10.1080/741941149>
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2006). Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *The Journal of Neuroscience*, 26, 12186–12197. <http://dx.doi.org/10.1523/JNEUROSCI.2818-06.2006>
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010a). Components of recognition memory: Dissociable cognitive processes or just differences in representational complexity? *Hippocampus*, 20, 1245–1262. <http://dx.doi.org/10.1002/hipo.20865>
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010b). Functional dissociations within the ventral object processing pathway: Cognitive modules or a hierarchical continuum? *Journal of Cognitive Neuroscience*, 22, 2460–2479. <http://dx.doi.org/10.1162/jocn.2009.21373>
- Cowell, R. A., Leger, K. R., & Serences, J. T. (2017). Feature-coding transitions to conjunction-coding with progression through human visual cortex. *Journal of Neurophysiology*, 118, 3194–3214. <http://dx.doi.org/10.1152/jn.00503.2017>
- Cowell, R. A., Sadil, P., Potter, K., & Huber, D. (2018). Implicit visual recollection: Connecting the dots without top-down knowledge. *Journal of Vision*, 18(10), 410. <http://dx.doi.org/10.1167/18.10.410>
- Cox, G. E., & Kalish, M. L. (2019). Dial M for monotonic: A kernel-based Bayesian approach to state-trace analysis. *Journal of Mathematical Psychology*. Advance online publication. <http://dx.doi.org/10.1016/j.jmp.2019.02.002>
- Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap methods and their application* (Vol. 1). Cambridge, UK: Cambridge university press. <http://dx.doi.org/10.1017/CBO9780511802843>
- Davis-Stober, C. P., Morey, R. D., Gretton, M., & Heathcote, A. (2016). Bayes factors for state-trace analysis. *Journal of Mathematical Psychology*, 72, 116–129. <http://dx.doi.org/10.1016/j.jmp.2015.08.004>
- Degonda, N., Mondadori, C. R. A., Bosshardt, S., Schmidt, C. F., Boesiger, P., Nitsch, R. M., . . . Henke, K. (2005). Implicit associative learning engages the hippocampus and interacts with explicit associative learning. *Neuron*, 46, 505–520. <http://dx.doi.org/10.1016/j.neuron.2005.02.030>
- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: A three-component model. *Trends in Cognitive Sciences*, 11, 379–386. <http://dx.doi.org/10.1016/j.tics.2007.08.001>

- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2008). The effects of unitization on familiarity-based source memory: Testing a behavioral prediction derived from neuroimaging data. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 730–740. <http://dx.doi.org/10.1037/0278-7393.34.4.730>
- Dunn, J. C. (2008). The dimensionality of the remember-know task: A state-trace analysis. *Psychological Review*, 115, 426–446. <http://dx.doi.org/10.1037/0033-295X.115.2.426>
- Dunn, J. C., & James, R. N. (2003). Signed difference analysis: Theory and application. *Journal of Mathematical Psychology*, 47, 389–416. [http://dx.doi.org/10.1016/S0022-2496\(03\)00049-X](http://dx.doi.org/10.1016/S0022-2496(03)00049-X)
- Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, 95, 91–101. <http://dx.doi.org/10.1037/0033-295X.95.1.91>
- Duss, S. B., Oggier, S., Reber, T. P., & Henke, K. (2011). Formation of semantic associations between subliminally presented face-word pairs. *Consciousness and Cognition*, 20, 928–935. <http://dx.doi.org/10.1016/j.concog.2011.03.018>
- Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, 30, 123–152. <http://dx.doi.org/10.1146/annurev.neuro.30.051606.094328>
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47. <http://dx.doi.org/10.1093/cercor/1.1.1>
- Geisser, F., & Eddy, W. (1979). A predictive approach to model selection. *Journal of the American Statistical Association*, 74, 153–160. <http://dx.doi.org/10.1080/01621459.1979.10481632>
- Gelbard-Sagiv, H., Faivre, N., Mudrik, L., & Koch, C. (2016). Low-level awareness accompanies “unconscious” high-level processing during continuous flash suppression. *Journal of Vision*, 16(1), 3. <http://dx.doi.org/10.1167/16.1.3>
- Gelfand, A., Dey, D. K., & Chang, H. (1992). *Model determination using predictive distributions with implementation via sampling-based methods*. Arlington, VA: Office of Naval Research.
- Goshen-Gottstein, Y., & Moscovitch, M. (1995). Repetition priming for newly formed and preexisting associations: Perceptual and conceptual influences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1229–1248. <http://dx.doi.org/10.1037/0278-7393.21.5.1229>
- Graboi, D., & Lisman, J. (2003). Recognition by top-down and bottom-up processing in cortex: The control of selective attention. *Journal of Neurophysiology*, 90, 798–810. <http://dx.doi.org/10.1152/jn.00777.2002>
- Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 501–518. <http://dx.doi.org/10.1037/0278-7393.11.3.501>
- Graf, P., Squire, L. R., & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 164–178. <http://dx.doi.org/10.1037/0278-7393.10.1.164>
- Graham, K. S., Barense, M. D., & Lee, A. C. H. (2010). Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, 48, 831–853. <http://dx.doi.org/10.1016/j.neuropsychologia.2010.01.001>
- Gregory, R. L. (1970). *The intelligent eye*. New York, NY: McGraw-Hill Book Company.
- Gregory, R. L. (1972). Cognitive contours. *Nature*, 238, 51–52. <http://dx.doi.org/10.1038/238051a0>
- Halekoh, U., & Højsgaard, S. (2014). A Kenward-Roger approximation and parametric bootstrap methods for tests in linear mixed models - The R package pbrtest. *Journal of Statistical Software*, 59, 9. <http://dx.doi.org/10.18637/jss.v059.i09>
- Henke, K. (2010). A model for memory systems based on processing modes rather than consciousness. *Nature Reviews Neuroscience*, 11, 523–532. <http://dx.doi.org/10.1038/nrn2850>
- Henke, K., Mondadori, C. R. A. A., Treyer, V., Nitsch, R. M., Buck, A., & Hock, C. (2003). Nonconscious formation and reactivation of semantic associations by way of the medial temporal lobe. *Neuropsychologia*, 41, 863–876. [http://dx.doi.org/10.1016/S0028-3932\(03\)00035-6](http://dx.doi.org/10.1016/S0028-3932(03)00035-6)
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36, 791–804. [http://dx.doi.org/10.1016/S0896-6273\(02\)01091-7](http://dx.doi.org/10.1016/S0896-6273(02)01091-7)
- Hubel, D. H., & Wiesel, T. N. (1965). Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28, 229–289. <http://dx.doi.org/10.1152/jn.1965.28.2.229>
- Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114, 3–28. <http://dx.doi.org/10.1037/0033-2909.114.1.3>
- Kalish, M. L., Dunn, J. C., Burdakov, O. P., & Sysoev, O. (2016). A statistical test of the equality of latent orders. *Journal of Mathematical Psychology*, 70, 1–11. <http://dx.doi.org/10.1016/j.jmp.2015.10.004>
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795. <http://dx.doi.org/10.1080/01621459.1995.10476572>
- Kent, B. A., Hvoslef-Eide, M., Saksida, L. M., & Bussey, T. J. (2016). The representational-hierarchical view of pattern separation: Not just hippocampus, not just space, not just memory? *Neurobiology of Learning and Memory*, 129, 99–106. <http://dx.doi.org/10.1016/j.nlm.2016.01.006>
- Kleiner, M., Brainard, D. H., Pelli, D. G., Broussard, C., Wolf, T., & Niehorster, D. (2007). What’s new in Psychtoolbox-3? *Perception*, 36, S14.
- Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *Journal of Neurophysiology*, 71, 856–867. <http://dx.doi.org/10.1152/jn.1994.71.3.856>
- Koffka, K. (1935). *Principles of Gestalt psychology*. Oxford, England: Harcourt, Brace & Company.
- Kruschke, J. K., & Vanpaemel, W. (2015). Bayesian estimation in hierarchical models. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *The Oxford handbook of computational and mathematical psychology* (pp. 279–299). Oxford, UK: Oxford University Press. <http://dx.doi.org/10.1093/oxfordhdb/9780199957996.013.13>
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6, 312–319. <http://dx.doi.org/10.3758/BF03197461>
- Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, 1, 476–490. <http://dx.doi.org/10.3758/BF03210951>
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY: Henry Holt and Co.
- MathWorks. (2015). MATLAB (R2015a). Natick, MA: The MathWorks Inc.
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4, 61–64. <http://dx.doi.org/10.20982/tqmp.04.2.p061>
- Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72, 6–18. <http://dx.doi.org/10.1016/j.jmp.2015.11.001>
- Newell, B. R., & Dunn, J. C. (2008). Dimensions in data: Testing psychological models using state-trace analysis. *Trends in Cognitive Sciences*, 12, 285–290. <http://dx.doi.org/10.1016/j.tics.2008.04.009>

- Newen, A., & Vetter, P. (2017). Why cognitive penetration of our perceptual experience is still the most plausible account. *Consciousness and Cognition*, 47, 26–37. <http://dx.doi.org/10.1016/j.concog.2016.09.005>
- Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, 3, 519–526. <http://dx.doi.org/10.3758/BF03197524>
- Pinto, Y., van der Leij, A. R., Sligte, I. G., Lamme, V. A. F., & Scholte, H. S. (2013). Bottom-up and top-down attention are independent. *Journal of Vision*, 13(3), 16. <http://dx.doi.org/10.1167/13.3.16>
- Pratte, M. S., & Rouder, J. N. (2011). Hierarchical single- and dual-process models of recognition memory. *Journal of Mathematical Psychology*, 55, 36–46. <http://dx.doi.org/10.1016/j.jmp.2010.08.007>
- Pratte, M. S., & Rouder, J. N. (2012). Assessing the dissociability of recollection and familiarity in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 1591–1607. <http://dx.doi.org/10.1037/a0028144>
- Prince, M., Brown, S., & Heathcote, A. (2012). The design and analysis of state-trace experiments. *Psychological Methods*, 17, 78–99. <http://dx.doi.org/10.1037/a0025809>
- Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, 3, 1–23. <http://dx.doi.org/10.1023/B:PHEN.0000041900.30172.e8>
- Ranganath, C. (2010). Binding items and contexts. *Current Directions in Psychological Science*, 19, 131–137. <http://dx.doi.org/10.1177/0963721410368805>
- R Development Core Team. (2016). *R: A language and environment for statistical computing* [Computer Software]. Vienna, Austria: R Foundation for Statistical Computing.
- Reber, T. P., & Henke, K. (2011). Rapid formation and flexible expression of memories of subliminal word pairs. *Frontiers in Psychology*, 2, 343. <http://dx.doi.org/10.3389/fpsyg.2011.00343>
- Reicher, G. M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81, 275–280. <http://dx.doi.org/10.1037/h0027768>
- Ross, D. A., Sadil, P., Wilson, D. M., & Cowell, R. A. (2018). Hippocampal engagement during recall depends on memory content. *Cerebral Cortex*, 28, 2685–2698. <http://dx.doi.org/10.1093/cercor/bhx147>
- Roth, H. L., Lora, A. N., & Heilman, K. M. (2002). Effects of monocular viewing and eye dominance on spatial attention. *Brain: A Journal of Neurology*, 125, 2023–2035. <http://dx.doi.org/10.1093/brain/awf210>
- Rouder, J. N., & Lu, J. (2005). An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review*, 12, 573–604. <http://dx.doi.org/10.3758/BF03196750>
- Rubin, E. (1915). *Synsoplevede figurer (visually experienced figures)*. Copenhagen, Denmark: University of Copenhagen.
- Rutishauser, U., Walther, D., Koch, C., & Perona, P. (2013). Is bottom-up attention useful for object recognition? In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004*. (Vol. 2, pp. 37–44). IEEE. <http://dx.doi.org/10.1109/CVPR.2004.1315142>
- Sadil, P., Cowell, R. A., & Huber, D. E. (2019). A hierarchical Bayesian state trace analysis for assessing monotonicity while factoring out subject, item, and trial level dependencies. *Journal of Mathematical Psychology*. Advance online publication. <http://dx.doi.org/10.1016/j.jmp.2019.01.003>
- Shibata, K., Watanabe, T., Sasaki, Y., & Kawato, M. (2011). Perceptual learning incepted by decoded fMRI neurofeedback without stimulus presentation. *Science*, 334, 1413–1415. <http://dx.doi.org/10.1126/science.1212003>
- Shimamura, A. P. (2010). Hierarchical relational binding in the medial temporal lobe: The strong get stronger. *Hippocampus*, 20, 1206–1216. <http://dx.doi.org/10.1002/hipo.20856>
- Squire, L. R., & Dedee, A. J. O. (2015). Conscious and unconscious memory systems. *Cold Spring Harbor Perspectives in Biology*, 7, a021667. <http://dx.doi.org/10.1101/cshperspect.a021667>
- Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, 253, 1380–1386. <http://dx.doi.org/10.1126/science.1896849>
- Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8, 1096–1101. <http://dx.doi.org/10.1038/nn1500>
- Tulving, E. (1983). *Elements of episodic memory*. New York, NY: Clarendon Press.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26, 1–12. <http://dx.doi.org/10.1037/h0080017>
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, 247, 301–306. <http://dx.doi.org/10.1126/science.2296719>
- van Tonder, G. J., & Ejima, Y. (2000). Bottom-up clues in target finding: Why a Dalmatian may be mistaken for an elephant. *Perception*, 29, 149–157. <http://dx.doi.org/10.1068/p2928>
- Wagenmakers, E. J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192–196. <http://dx.doi.org/10.3758/BF03206482>
- Wagenmakers, E.-J., Krypotos, A.-M., Criss, A. H., & Iverson, G. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & Cognition*, 40, 145–160. <http://dx.doi.org/10.3758/s13421-011-0158-0>
- Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. *Cognitive Psychology*, 60, 158–189. <http://dx.doi.org/10.1016/j.cogpsych.2009.12.001>
- Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & Iverson, G. J. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology*, 48, 28–50. <http://dx.doi.org/10.1016/j.jmp.2003.11.004>
- Weinberger, N. M. (2007). Associative representational plasticity in the auditory cortex: Resolving conceptual and empirical problems. *Debates in Neuroscience*, 1, 85–98. <http://dx.doi.org/10.1007/s11559-007-9011-9>
- Yang, E., & Blake, R. (2012). Deconstructing continuous flash suppression. *Journal of Vision*, 12(3), 8. <http://dx.doi.org/10.1167/12.3.8>
- Yao, Y., Vehtari, A., Simpson, D., & Gelman, A. (2018). Using stacking to average Bayesian predictive distributions (with discussion). *Bayesian Analysis*, 13, 917–1007. <http://dx.doi.org/10.1214/17-BA1091>
- Yonelinas, A. P. (2013). The hippocampus supports high-resolution binding in the service of perception, working memory and long-term memory. *Behavioural Brain Research*, 254, 34–44. <http://dx.doi.org/10.1016/j.bbr.2013.05.030>

Received January 10, 2018

Revision received January 4, 2019

Accepted March 5, 2019 ■