

Distinct error-correcting and incidental learning of location relative to landmarks and boundaries

Christian F. Doeller* and Neil Burgess*

Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, United Kingdom

Communicated by Jan Bureš, Czech Academy of Sciences, Prague, Czech Republic, December 22, 2007 (received for review November 13, 2007)

Associative reinforcement provides a powerful explanation of learned behavior. However, an unproven but long-held conjecture holds that spatial learning can occur incidentally rather than by reinforcement. Using a carefully controlled virtual-reality object-location memory task, we formally demonstrate that locations are concurrently learned relative to both local landmarks and local boundaries but that landmark-learning obeys associative reinforcement (showing “overshadowing” and “blocking” or “learned irrelevance”), whereas boundary-learning is incidental, showing neither overshadowing nor blocking nor learned irrelevance. Crucially, both types of learning occur at similar rates and do not reflect differences in levels of performance, cue salience, or instructions. These distinct types of learning likely reflect the distinct neural systems implicated in processing of landmarks and boundaries: the striatum and hippocampus, respectively [Doeller CF, King JA, Burgess N (2008) *Proc Natl Acad Sci USA* 105:5915–5920]. In turn, our results suggest the use of fundamentally different learning rules by these two systems, potentially explaining their differential roles in procedural and declarative memory more generally. Our results suggest a privileged role for surface geometry in determining spatial context and support the idea of a “geometric module,” albeit for location rather than orientation. Finally, the demonstration that reinforcement learning applies selectively to formally equivalent aspects of task-performance supports broader consideration of two-system models in analyses of learning and decision making.

associative learning | blocking | hippocampus | overshadowing | striatum

The dominant model of learning from repeated feedback (or “reinforcement”) associates environmental cues with expected reinforcement and with actions, using a single prediction-error signal (the difference between actual and expected reinforcement) to modify these associations (1–5). However, spatial learning, a crucial aspect of daily life, has long been proposed to exemplify a qualitatively different type of learning (6, 7), whereby “incidental” and “latent” learning occur independent of reinforcement. Perhaps surprisingly, formal demonstration that spatial learning deviates from the predictions of associative reinforcement, given the additional assumption that exploration can be rewarding in itself, has not been forthcoming (8–11).

Here, we carefully dissociate two contributions to spatial learning in virtual environments: learning locations relative to a local boundary and learning locations relative to a local landmark. Learning relative to a local landmark has been shown to follow the predictions of associative reinforcement learning (10). By contrast, we predicted that boundary-related learning would be incidental, for two reasons. First, the hippocampus has been specifically implicated in incidental learning in spatial (12) and other (13–19) domains. Second, the hippocampus seems to specifically process location relative to local environmental boundaries. The firing of hippocampal place cells reflects distances and directions to local boundaries (20, 21) but not to intramaze landmarks (22), and human hippocampal activation corresponds to learning of locations relative to local boundaries but not to local landmarks (23).

Associative reinforcement operating on multiple cues, using a single error term, predicts “blocking” (24) and “overshadowing” (24, 25) between cues. We give a standard analysis, based on ref. 1,

but note that these predictions also apply to other formulations of associative learning based on an error-correcting principle (2, 4, 26). Expected reinforcement v is related to the vector of stimuli \underline{u} via the vector of associative strengths or “weights” \underline{w} according to:

$$v = \underline{w} \cdot \underline{u} \quad [1]$$

and the learning rule adjusts these weights according to:

$$\underline{w} \rightarrow \underline{w} + \varepsilon \delta \underline{u} \quad [2]$$

where ε is a constant and δ , the prediction-error term, is the difference between expected and actual reinforcement r :

$$\delta = (r - v). \quad [3]$$

It can be seen that learning to cue 2 (i.e., adjusting weight w_2 from element u_2 of the stimulus vector) is reduced by the extent to which reinforcement is already predicted by another cue 1 (e.g., if $w_2 = 0$ and $v = w_1 u_1 = r$, then $\delta = 0$ and w_2 will not be modified). Put simply, if one cue already accurately predicts feedback, the error term δ is reduced, blocking the learning of associations from other cues. Similarly, when learning to two cues occurs concurrently, learning to one of the cues can overshadow learning to the other by reducing δ . In models of learning to act through reinforcement, the above mechanism is referred to as the “critic,” whereas an “actor” learns to associate stimuli to actions, using the same prediction-error and learning rule to modify these associations [more generally, Eq. 3 is adapted to include the effects of actions on the prediction of future reward, e.g., via the temporal difference rule: $\delta(t) = r(t) + v(t+1) - v(t)$; see refs. 5 and 27]. Thus, learned behavior (“instrumental” conditioning) is predicted to show blocking and overshadowing in the same way as learned expectation of reward (Pavlovian or “classical” conditioning). We predicted that learning relative to the landmark would show blocking and overshadowing, whereas learning relative to the boundary would continue irrespective of prior or concurrent learning to other cues.

To test these predictions, we designed a virtual environment in which locations can be learned relative to a local circular boundary or to a local intramaze landmark or landmarks. Distal cues were always present to provide orientation but could not be associated with locations as rendered at infinity. The virtual environment (a modified video game; see *Methods*) was presented in first-person perspective on a screen; the participant could navigate through it by pressing buttons to turn left or right or to move forwards. Participants initially collected several objects in turn from the virtual arena by navigating to them. In each subsequent trial, they were cued to replace an object (the object appeared on a blank background, they navigated to where they thought they had collected it

Author contributions: C.F.D. and N.B. designed research; C.F.D. performed research; C.F.D. analyzed data; and C.F.D. and N.B. wrote the paper.

The authors declare no conflict of interest.

*To whom correspondence may be addressed. E-mail: n.burgess@ucl.ac.uk or c.doeller@ucl.ac.uk.

This article contains supporting information online at www.pnas.org/cgi/content/full/0711433105/DC1.

© 2008 by The National Academy of Sciences of the USA

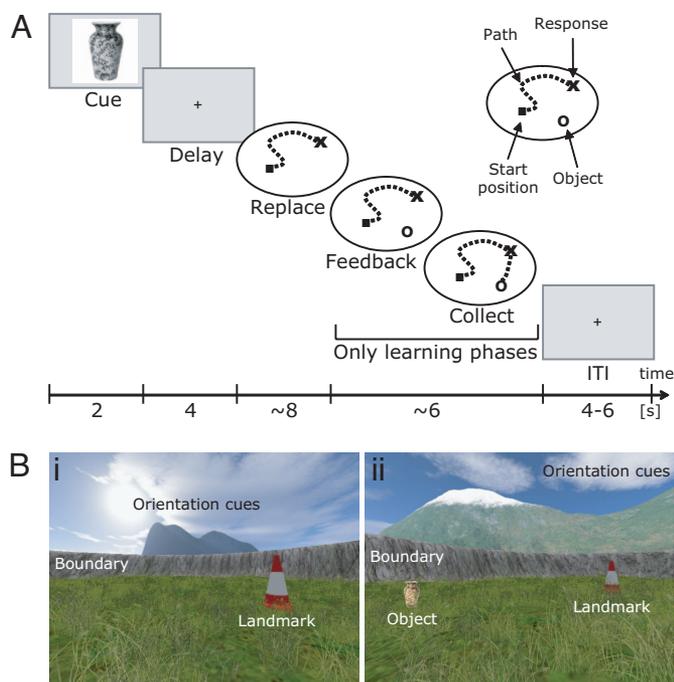


Fig. 1. Virtual reality task. (A) Trial structure (after initial collection of objects). Participants replace the cued object after a short delay phase and received feedback during learning trials (object appears in correct location immediately after the response and is collected) but not testing trials. (B) Virtual arena from the participant's perspective (*i*, replace phase; *ii*, feedback phase; different viewpoints) showing the intramaze landmark (traffic cone), the boundary (circular wall), the extramaze orientation cues (mountains, which were projected at infinity), and one object (vase). ITI, intertrial interval.

and pressed a button). Performance was measured as the distance between the participant's response and the correct location. Learning trials included feedback (the object appearing in the correct location to be collected again) from which participants learned to improve performance. "Test" trials omitted this feedback. See Fig. 1. Start positions varied between trials so that path integration could not be used.

Results

Overshadowing Experiments. In everyday life, locations are learned in the presence of multiple stable cues of both types [local landmarks (L) and boundaries (B)]. How do both types of learning interact? Four groups of 12 participants each learned four object-locations over four trials per object with either a landmark or a boundary present alone ("simple learning," groups L and B) or both cues present together ("compound learning," groups LB1 and LB2). Performance on all four objects was then tested with one or other cue alone (group LB1 tested with L; group LB2 tested with B) to compare the strengths of associations formed to that cue during compound learning with the strengths of equivalent associations formed during simple learning (i.e., group L tested with L; group B tested with B) (see Fig. 2). Associative reinforcement predicts a reduced associative strength after compound learning compared with simple learning because of overshadowing.

As predicted, learning to the landmark was reduced (overshadowed) by the presence of the boundary during compound learning, consistent with associative reinforcement, while learning to the boundary occurred incidentally, unaffected by the presence or absence of the landmark. ANOVA of the test performance of the four groups as a function of learning-situation (simple vs. compound) \times test-cue (L vs. B) revealed a significant interaction ($F_{1,44} = 5.57$; $P < 0.05$) and no main

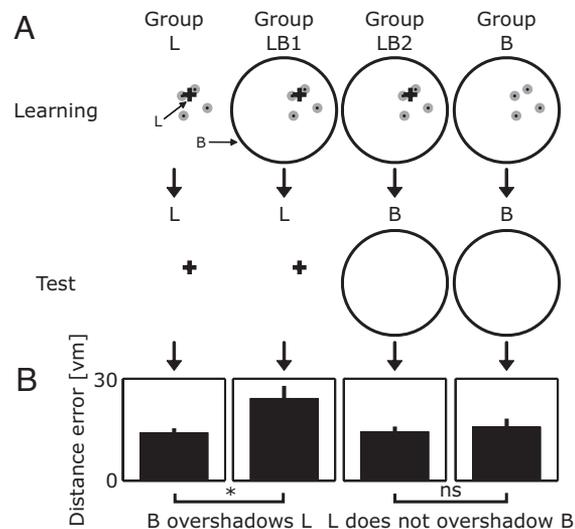


Fig. 2. Overshadowing experiment. The landmark is overshadowed by the boundary but not vice versa. (A) Four different groups (columns, 12 participants per group) learned four object-locations with either one of (simple learning); or both of the landmark (L) and boundary (B) present (compound learning); *Upper* and were tested with either landmark or boundary alone (*Lower*). (B) Boundary overshadows landmark (i.e., the presence of the boundary during learning reduces learning to the landmark in group LB1 compared with group L) but not vice versa (i.e., the presence of the landmark during learning does not reduce learning to the boundary in group LB2 compared with group B). Bars indicate mean distance between response and correct location during test phase in virtual meters, \pm SEM; *, $P < 0.05$; ns, not significant. +, landmark; circle, boundary; dots, object locations.

effects of learning-situation ($F_{1,44} = 3.16$; $P > 0.08$) or test-cue ($F_{1,44} = 2.79$; $P > 0.1$). Posthoc tests verified the greater replacement error of group LB1 than group L (both tested with landmark; $t_{22} = 2.58$, $P < 0.05$), whereas groups LB2 and B did not differ (both tested with boundary; $t_{22} < 1$, $P > 0.6$). These differential effects were not due to differences in performance during the learning phase with either cue [no difference in simple learning to L and B; $t_{22} < 1$, $P > 0.7$; see [supporting information \(SI\) Fig. 5](#) and *SI Text*].

Thus, the presence of the boundary during compound learning overshadowed learning to the landmark, as predicted by associative reinforcement. By contrast, the presence of the landmark during compound learning did not affect learning to the boundary. Simple learning to either cue alone did not differ, ruling out differences in the salience of the two cues. Nonetheless, we also explicitly manipulated the visual salience of the landmark: running new groups L', LB1', and LB2' in a second experiment with a landmark three times as tall as previously (group B never saw the landmark). We found identical effects; overshadowing of the landmark was unaffected by its visual salience (see [SI Fig. 4](#) and *SI Text*).

Blocking Experiments. In three further experiments, we examined blocking between two local landmarks L1 and L2 (Experiment LL), between a local landmark L and a local boundary B (Experiment LB), and between opposite sections B1 and B2 of the local boundary (Experiment BB). Each experiment occurred in three phases: (*i*) During an initial learning phase (known as "prelearning"), the two cues (L1 and L2, L and B, or B1 and B2) moved relative to each other at the start of each block of trials (eight blocks, 14 trials per object in total) with four object-locations paired with either cue; i.e., after the two cues moved, four objects kept a fixed bearing to the first cue (but not to the second), whereas the other four objects kept a fixed bearing to the second cue (but not to the first). (*ii*) Next, during compound learning, both cues and object-locations remained fixed (six trials per object), allowing associations

occurs incidentally, irrespective of the level of error: Object-locations already successfully predicted by another cue were nonetheless learned relative to the boundary and to the same accuracy as object-locations paired with the boundary throughout.

As with the overshadowing results, the differential blocking of landmarks but not boundaries does not reflect differences in learning rates (see above), or greater salience of boundaries (initial performance is, if anything, slightly better for objects paired with landmarks (see *SI Fig. 6* and *SI Text*). In further analyses, of objects equidistant to both cues during compound learning, we additionally ruled out that differential blocking effects were due to differential cue proximity (see *SI Fig. 7* and *SI Text*).

Finally, we note that the absence of any performance differences, when tested with a boundary in experiments LB and BB, between objects initially paired or unpaired with the boundary also rules out any effects of learned irrelevance. Both learned irrelevance and blocking should produce a performance difference in this situation, and none was found.

Discussion

Although associative reinforcement via a single prediction-error signal provides a powerful, almost ubiquitous, model for learning over repeated experience, it captures the acquisition of some types of knowledge, but not others, even when both are learned concurrently in formally identical conditions (see ref. 29; O. Hardt, A. Humpbach, and L. Nadel, unpublished data). More specifically, learning object-locations relative to intramaze landmarks obeys associative reinforcement (showing overshadowing and blocking or learned irrelevance), whereas learning relative to environmental boundaries is incidental, occurring independently of behavioral error or the presence of other predictive cues (showing neither overshadowing nor blocking nor learned irrelevance).

The two types of learning occur in parallel within the same task, without differences in the time-courses of learning and stimulus presentation, performance levels, instructions, location-cue proximity, or cue salience. Thus, learning to the boundary seems fundamentally inconsistent with associative reinforcement based on a single prediction-error term, and would potentially require separate error signals for landmarks and for each segment of boundary in the environment (see also ref. 30).

The differential learning relative to the boundary compared with that relative to the landmark over the three blocking experiments is also inconsistent with the many elaborations of basic associative reinforcement. Thus, increased learning rates to more salient stimuli (2) cannot explain our results, because performance indicates that the landmark was always at least as salient as the boundary. There was also no effect of tripling the landmark's height. Goal-directed reinforcement learning, in which information about the nature of the reinforcer is learnt (refs. 31 and 32; see also ref. 33) cannot explain our differential results, because learning to both cues results from the same feedback. Other potential explanations are that poor predictors during initial learning will have an increased learning rate during compound learning (4), that the presence of a good predictor will aid the learning to a poor predictor by increasing reward frequency ("feature enhancement") (34), that a cue's associability varies with its associative history (35), or that performance changes after removal of a cue reflect a "generalization decrement" (36, 37) cannot easily explain the differential blocking seen in experiment LL compared with experiment BB in which the contingencies of the two landmarks and the two boundaries are identical.

What causes the distinct characteristics of learning to landmarks and boundaries in our experiments? One potential explanation would be that the two cues interact differently with visual behavior (see also ref. 34) such that the previously paired landmark is not seen during some compound learning trials, allowing learning to the boundary, whereas the (larger) boundary is nearly always seen and therefore provides more potent blocking. However, participants

started each trial from a position toward the edge of the arena looking inwards, with all cues visible, and often made an entire rotation during the feedback phase. In addition, this explanation would predict at most a graded difference (i.e., with some learning for the unpaired landmark in experiment LL and some blocking of the unpaired boundary in experiment BB), rather than the stark contrast observed. Thus, this explanation applies more readily to failures of transiently presented stimuli in blocking continuous "contextual" cues (38) (i.e., temporal or environmental cues, although incidental learning to environmental boundaries might also contribute to the latter).

Our interpretation, and the hypothesis behind our design, is that the contributions of landmarks and boundaries to spatial memory are supported by distinct striatal and hippocampal systems operating different learning rules. To test for this, participants in an fMRI experiment (23), performed four blocks of trials, each block identical to the learning phase of compound groups LB1 and LB2 in our overshadowing experiment. However, at the start of each block, the landmark and boundary were moved relative to each other, with two of the four objects paired with either cue (similarly to the initial learning phase of our blocking experiments). The association of each object location to landmark or boundary was learned at similar rates for either cue. Significantly, learning relative to the boundary corresponded to activation in the right posterior hippocampus, whereas right dorsal striatum activation reflected learning of landmark-related locations. In this view, the incidental learning to boundaries comprises unsupervised Hebbian association between hippocampal place cells [whose firing represents conjunctions of bearings to boundaries (20, 21)] and object locations, whenever both representations are coincidentally active.

Previous studies of associative reinforcement in spatial learning have produced mixed results. Our present data in combination with our identification of distinct systems for processing local boundaries versus local landmarks (23) helps to clarify these previous findings. Thus, rats in the water maze task use the boundary to locate the platform (39), a task that is hippocampal-dependent (40), and boundary-learning did not show blocking in our study. Conversely, learning the platform location relative to an intramaze landmark in the water maze is not hippocampal-dependent (41), and blocking has been observed in the processing of intramaze landmarks both in rats (10) and our study. Note that we avoided using intramaze landmarks placed exactly at the target location, because these can impede learning of more distant cues simply by focusing attention on the target location alone (11, 42). We also ensured that orientation depended solely on distal cues. Otherwise, orientation might depend on many types of cue, combining outside of the hippocampus in the head-direction circuit (43) and possibly by associative reinforcement. Thus, asymmetrical boundary geometry and distal cues combine to orient participants (44) [and place cell firing (45)] and can overshadow and block each other in doing so (8, 9, 46, 47). Whether or not distal cues do block each other in determining orientation depends on many factors, including task instructions (O. Hardt, A. Humpbach, and L. Nadel, unpublished data). Even though boundary geometry can dominate intramaze cues to orientation in some circumstances (7, 48) and this can result in an absence of blocking or overshadowing of the boundary in determining orientation (49), it is hard to rule out associative explanations invoking the differential salience of the cues (2) or feature enhancement because of interactions between cues mediated by behavioral choice (34).

What distinguishes a local boundary from a local landmark? The finding that firing of hippocampal place cells is well described as a match to the distances to obstacles in all directions around the animal (20, 21) indicates that an important attribute for hippocampal processing is simply the horizontal angle subtended by the obstacle at the animal. Thus, "boundaryness" may simply reflect extent as seen from the participant's viewpoint. Note that the distinction is not due to differential proximity to the object-location,

because we found the same differential pattern of blocking between boundaries and landmarks equidistant to the object (see *SI Fig. 7* and ref. 23 for further discussion).

How do our results relate to findings that disoriented rats (48) and toddlers (50) reorient themselves using only boundary geometry—ignoring salient local features? These are interpreted as evidence of a “geometric module” for reorientation (refs. 7, 48, and 50 but see also ref. 51). Our results are not explainable by this type of nonspecific dominance of the boundary over the landmark: Performance and learning rates to both types of cue are matched. In addition, orientation is controlled by the distal cues rather than the circular boundary. By contrast, we show that there may be specialized learning of locations relative to surface geometry, even when it does not dominate behavior, in that it proceeds via a different rule than learning relative to intramaze landmarks.

How do our results relate to the involvement of the hippocampal and striatal systems in respectively supporting “place” versus “response” learning (12, 52), or flexible–relational/declarative/episodic versus procedural memory (13–15, 17, 18, 53)? Consistent with many of these suggestions, we formally demonstrate that the two systems’ distinct roles may result from differences in the learning rule implemented in either system—and not necessarily from other differences, such as in learning rate, that were not controlled in previous less formal approaches. In this view, striatal synaptic plasticity is controlled by a single error signal reflecting deviations from expected reinforcement (1, 3, 5) possibly mediated by dopamine (54–56). This produces “trial-and-error learning” or “learning by doing” appropriate for procedural memory (15). Unsupervised hippocampal Hebbian synaptic plasticity occurs independent of error or reinforcement, although it may be boosted by novelty [i.e., inconsistency between perception and relationships stored in hippocampus (12, 17)]. This produces incidental learning or “learning by observation” appropriate with maintenance of a flexible mental model (17), “mediating representation” (19), or “cognitive map” (6, 12), and with efficient encoding of experience into episodic memory (15, 16). In addition, our results suggest that the local boundaries may dominate over local landmarks in defining the spatial context of hippocampally mediated episodic memories. More generally, our demonstration of concurrent incidental and reinforcement learning requires a broadening of the models of learning and decision-making applied to social and economic neuroscience (see also ref. 33).

Methods

Participants. One hundred thirty-two male participants (aged 18–37, mean age 21.9 years) took part in this study, which was approved by the local Research Ethics Committee. Data from three additional participants had to be excluded because of technical problems or misunderstanding of the task. Participants gave full informed written consent and were paid for participating.

Virtual Reality Environment. We used UnrealEngine2 Runtime software (Epic Games) to display a first-person perspective view of a grassy plane surrounded by a landscape of two mountains, clouds, and the sun (created by using Terragen; Planetside Software) and projected at infinity, to provide distal cues to orientation but not to location. The arena was surrounded by a circular boundary (a cliff) in some experiments (or contained a semicircular portion of it in the BB blocking test). In the absence of the cliff, the grassy plane extended to infinity. A traffic cone was used as an intramaze landmark (and a bowling pin as the second landmark in the LL blocking experiment). Both the boundary and landmarks were rotationally symmetric, leaving the distal cues as the main source of orientation. Participants moved the viewpoint by pressing keys to move forward and turn left or right. The viewpoint is ≈ 2 virtual meters above ground, the standard boundary was ≈ 180 virtual meters in diameter (this varied in the BB blocking experiment), and the virtual heading and location were recorded every 100 ms. Participants received training in an unrelated virtual environment before performing the experiment (see *SI Text*).

Stimuli, Task, and Trial Structure. Pictures of everyday objects were used as experimental stimuli. Before each experiment, participants familiarized them-

selves with the arena by exploring for 2–3 min. Next, objects were presented sequentially (once each) within the arena and participants collected the objects by running over them and were instructed to remember their locations. At the beginning of each subsequent trial, a picture of an object was presented for 2 s on a blank background (the cue phase), followed by a variable delay period (fixation cross for 2–6 s; mean 4 s). Participants then started at a variable position within the arena (see *SI Text*) and had to move to where they thought the cued object had been (the replace phase; mean duration, 8.20 s). After participants had indicated their response by a button press, feedback was provided, i.e., the object appeared in its correct position and participants collected it by running over it (the feedback phase; mean duration 6.23 s). Feedback was not provided in the test phases of the experiments (see below); in the blocking experiments, participants waited in the arena for 6 s (mean) after their response. A fixation cross was then presented for a variable intertrial interval (ITI) (overshadowing experiments: 2–10 s, mean 6 s; blocking experiments: 2–6 s, mean = 4 s) before the start of the next trial.

Overshadowing Experiments. Procedure and design. The overshadowing experiments comprised two phases: learning and test. Each phase was one block of 16 trials (4 trials per object; four objects). Objects were presented in pseudorandom order within blocks. Feedback was provided in the learning phase but not in the test phase. To reduce the error variance in the planned between-group comparisons, each participant of a particular group had the same configuration of object order and delay and ITI times as a matching participant in the other groups. Participants were randomly assigned to different groups.

Basic experiment. Forty-eight participants were tested in four groups of 12. In the learning phase of groups LB1 and LB2, both the landmark and the boundary were present. For group L only the landmark was present and in group B only the boundary was present. In the test phase, either solely the landmark (groups LB1 and L) or solely the boundary (groups LB2 and B) was present [see Fig. 2; see *SI Text* for details of the control experiment with a taller landmark (three groups of 12 participants)].

Blocking Experiments. Procedure and design. The three blocking experiments (16 participants in each experiment) comprised three phases: initial learning (or prelearning), separated into eight blocks of one to three trials per object (14 trials per object in total; eight experimental objects); compound learning (one block of six trials per object); test phase, comprising two blocks in which memory for four experimental object-locations was tested in the presence of one cue alone, followed by testing of memory for the other four object locations in the presence of the other cue alone (six trials per object). Feedback about the correct object position was provided in the initial learning and compound learning phases but not in the test phase. Objects were presented in pseudorandom order within blocks (trials with two control objects were interspersed among regular trials; these trials do not require spatial memory: the objects are collected from an infinite grassy plane with blue background and are always visible, as used in our fMRI experiment, see ref. 23 for details). **Initial learning phase.** The experimental objects were separated into two sets: four objects associated with cue 1 (L1 in Experiment LL; L in Experiment LB; or B1 in Experiment BB) and four objects associated with cue 2 (L2 in Experiment LL; B in Experiment LB; or B2 in Experiment BB).

In Experiment LL, the two landmarks were moved relative to each other between blocks (using three different spatial configurations of the two landmarks). After each change, four objects maintained a fixed bearing to landmark L1 (but not to landmark L2); the other four objects maintained a fixed bearing to landmark L2 (but not to landmark L1).

In Experiment LB, the landmark and boundary moved relative to each other between blocks (using four different landmark positions approximately in the middle of the NE, SE, SW, and NW sectors of the arena, as defined by the distal orientation cues). After each change, four objects maintained a fixed bearing to the boundary B (but not to the landmark L), the other four objects maintained a fixed bearing to the landmark L (but not to the boundary B).

In Experiment BB, the radius of the circular boundary changed between blocks (three sizes: the one used in all other experiments including a boundary, a 20% smaller radius, and a 20% larger radius). After each change, four objects maintained a fixed bearing to the Western part of the cliff, boundary B1 (but not to boundary B2); the other four objects maintained a fixed bearing to the Eastern part of the cliff, boundary B2 (but not to boundary B1).

Thus, over the course of the initial learning phase, the positions of four objects were predicted by cue 1 but not by cue 2, whereas the positions of the other four objects were predicted by cue 2 but not by cue 1 (see Fig. 3).

Compound Learning Phase. During compound learning, both cues predicted the position of all eight experimental objects, remaining in the same locations throughout.

Test Phase. Memory for four of the object locations defined in the compound learning phase was tested with only one cue present (L1 in Experiment LL, L in Experiment LB, B1 in Experiment BB), and the other four were tested with only the other cue present (L2 in Experiment LL, B in Experiment LB, B2 in Experiment

BB). Two cue 1-related objects and two cue 2-related objects entered each test. The order of test runs was counterbalanced across subjects.

ACKNOWLEDGMENTS. This work was funded by the Biotechnology and Biological Sciences Research Council and the U.K. Medical Research Council. We thank J. King for technical help; M. Lehmann for help with testing; and P. Dayan, U. Frith, M. Good, R. Honey, M. Mishkin, J. O'Keefe, and R. Rescorla for useful discussions.

- Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. *Classical conditioning II. Current research and theory*, eds Black AH, Prokasy WF. (Appleton-Century-Crofts, New York), pp 64–99.
- Mackintosh NJ (1975) A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychol Rev* 82:276–298.
- Dickinson A (1980) Contemporary animal learning theory. (Cambridge Univ Press, Cambridge, UK).
- Pearce JM, Hall G (1980) A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev* 87:532–552.
- Sutton RS, Barto AG (1988) Reinforcement learning: An introduction. (MIT Press, Cambridge, MA).
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208.
- Gallistel R (1990) The organization of learning. (MIT Press, Cambridge, MA).
- Rodrigo T, Chamizo VD, McLaren IP, Mackintosh NJ (1997) Blocking in the spatial domain. *J Exp Psychol Anim Behav Process* 23:110–118.
- Hamilton DA, Sutherland RJ (1999) Blocking in human place learning: Evidence from virtual navigation. *Psychobiol* 27:453–461.
- Biegler R, Morris RGM (1999) Blocking in the spatial domain with arrays of discrete landmarks. *J Exp Psychol Anim Behav Process* 25:334–351.
- Chamizo VD (2003) Acquisition of knowledge about spatial location: Assessing the generality of the mechanism of learning. *Q J Exp Psychol B* 56:102–113.
- O'Keefe J, Nadel L (1978) The hippocampus as a cognitive map. (Oxford Univ Press, Oxford, UK).
- Hirsh R (1974) The hippocampus and contextual retrieval of information from memory: A theory. *Behav Biol* 12:421–444.
- Mishkin M, Malamut B, Bachevalier J (1984) Memories and habits: Two neural systems. *Neurobiology of Learning and Memory*, eds Lynch G, McGaugh JL, Weinberger M (Guilford, New York), pp 65–77.
- Squire LR, Zola SM (1996) Structure and function of declarative and nondeclarative memory systems. *Proc Natl Acad Sci USA* 93:13515–13522.
- Morris RGM, Frey U (1997) Hippocampal synaptic plasticity: Role in spatial learning or the automatic recording of attended experience? *Philos Trans R Soc Lond B* 352:1489–1503.
- Eichenbaum H, Cohen NJ (2001) *From Conditioning to Conscious Recollection: Memory Systems of the Brain* (Oxford Univ Press, Oxford).
- Poldrack RA, Clark J, Pare-Blagoev EJ, Shohamy D, Creso MJ, Myers C, Gluck MA (2001) Interactive memory systems in the human brain. *Nature* 414:546–550.
- Gluck MA, Meeter M, Myers CE (2003) Computational models of the hippocampal region: Linking incremental learning and episodic memory. *Trends Cogn Sci* 7:269–276.
- O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381:425–428.
- Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J (2000) Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus* 10:369–379.
- Cressant A, Muller RU, Poucet B (1997) Failure of centrally placed objects to control the firing fields of hippocampal place cells. *J Neurosci* 17:2531–2542.
- Doeller CF, King JA, Burgess N (2008) Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proc Natl Acad Sci USA* 105:5915–5920.
- Kamin LJ (1969) Predictability, surprise, attention and conditioning. *Punishment and aversive behavior*, eds Campbell BA, Church RM (Appleton-Century-Crofts, New York), pp 279–296.
- Pavlov IP (1927) *Conditioned Reflexes* (Oxford Univ Press, Oxford).
- Wagner AR (1981) SOP: A model of automatic memory processing in animal behaviour. *Information Processing in Animals: Memory Mechanism*, eds Spear NE, Miller RR (Lawrence Erlbaum, Hillsdale, NY), pp 5–47.
- Dayan P, Balleine BW (2002) Reward, motivation, and reinforcement learning. *Neuron* 36:285–298.
- Baker AG, Mackintosh NJ (1977) Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Anim Learn Behav* 5:315–319.
- Raby CR, Alexis DM, Dickinson A, Clayton NS (2007) Planning for the future by western scrub-jays. *Nature* 445:919–921.
- Rescorla RA (2002) Effect of following an excitatory-inhibitory compound with an intermediate reinforcer. *J Exp Psychol Anim Behav Process* 28:163–174.
- Dickinson A (1985) Actions and habits: The development of behavioural autonomy. *Philos Trans R Soc Lond B* 308:67–78.
- Balleine BW, Dickinson A (1998) Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37:407–419.
- Daw ND, Niv Y, Dayan P (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8:1704–1711.
- Miller NY, Shettleworth SJ (2007) Learning about environmental geometry: An associative model. *J Exp Psychol Anim Behav Process* 33:191–212.
- Le Pelley ME (2004) The role of associative history in models of associative learning: A selective review and a hybrid model. *Q J Exp Psychol B* 57:193–243.
- Pearce JM (1987) A model for stimulus generalization in Pavlovian conditioning. *Psychol Rev* 94:61–75.
- Brandon SE, Vogel EH, Wagner AR (2000) A componential view of configural cues in generalization and discrimination in Pavlovian conditioning. *Behav Brain Res* 110:67–72.
- Williams DA, LoLordo VM (1995) Time cues block the CS, but the CS does not block time cues. *Q J Exp Psychol B* 48:97–116.
- Maurer R, Derivaz V (2000) Rats in a transparent morris water maze use elemental and configural geometry of landmarks as well as distance to the pool wall. *Spatial Cognit Comput* 2:135–156.
- Morris RGM, Garrud P, Rawlins JN, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297:681–683.
- Pearce JM, Roberts AD, Good M (1998) Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature* 396:75–77.
- Redhead ES, Roberts A, Good M, Pearce JM (1997) Interaction between piloting and beacon homing by rats in a swimming pool. *J Exp Psychol Anim Behav Process* 23:340–350.
- Taube JS (1998) Head direction cells and the neuropsychological basis for a sense of direction. *Prog Neurobiol* 55:225–256.
- Hartley T, Trinkler I, Burgess N (2004) Geometric determinants of human spatial memory. *Cognition* 94:39–75.
- Jeffery KJ, Donnett JG, Burgess N, O'Keefe JM (1997) Directional control of hippocampal place fields. *Exp Brain Res* 117:131–142.
- Fenton AA, Arofo MP, Nerad L, Bureš J (1994) Place navigation in the Morris water maze under minimum and redundant extramaze cue conditions. *Behav Neural Biol* 62:178–189.
- Pearce JM, Graham M, Good MA, Jones PM, McGregor A (2006) Potentiation, overshadowing, and blocking of spatial learning based on the shape of the environment. *J Exp Psychol Anim Behav Process* 32:201–214.
- Cheng K (1986) A purely geometric module in the rat's spatial representation. *Cognition* 23:149–178.
- Hayward A, McGregor A, Good MA, Pearce JM (2003) Absence of overshadowing and blocking between landmarks and the geometric cues provided by the shape of a test arena. *Q J Exp Psychol B* 56:114–126.
- Hermer L, Spelke E (1996) Modularity and development: The case of spatial reorientation. *Cognition* 61:195–232.
- Pearce JM, Good MA, Jones PM, McGregor A (2004) Transfer of spatial behavior between different environments: Implications for theories of spatial learning and for the role of the hippocampus in spatial learning. *J Exp Psychol Anim Behav Process* 30:135–147.
- Packard MG, McGaugh JL (1996) Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol Learn Mem* 65:65–72.
- Knowlton BJ, Mangels JA, Squire LR (1996) A neostriatal habit learning system in humans. *Science* 273:1399–1402.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599.
- Waelti P, Dickinson A, Schultz W (2001) Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412:43–48.
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD (2006) Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442:1042–1045.



Sign up for PNAS Online eTocs

Get notified by email when new content goes on-line

[Info for Authors](#) | [Editorial Board](#) | [About](#) | [Subscribe](#) | [Advertise](#) | [Contact](#) | [Site Map](#)

PNAS

Proceedings of the National Academy of Sciences of the United States of America

[Current Issue](#)

[Archives](#)

[Online Submission](#)

GO

[advanced search >>](#)

Institution: [UCL Library Services](#) | [Sign In via User Name/Password](#)

Doeller and Burgess. 10.1073/pnas.0711433105.

Supporting Information

Files in this Data Supplement:

[SI Figure 4](#)

[SI Text](#)

[SI Figure 5](#)

[SI Figure 6](#)

[SI Figure 7](#)

[SI Figure 4](#)

Fig. 4.

Overshadowing and landmark salience. Three new groups L1', LB1', and LB2' (columns, 12 participants per group; see magenta bars in yellow area) learned four object-locations with a more salient landmark (3× taller) and were tested with either landmark or boundary alone (bottom row). The asymmetric overshadowing effect of the main experiment was also observed with this more salient landmark: the boundary overshadows the landmark (cf. replacement error for groups LB1' vs. L') but not vice versa (cf. groups LB2' vs. B). For comparison we also depict the findings of the original 3 groups with the normal landmark (blue; note that group B learned and was tested with the boundary only: black), see Fig. 2. These data show that the differential overshadowing effects do not depend on the visual salience of the landmark.

[SI Figure 5](#)

Fig. 5.

Overshadowing experiments: performance during learning. The figure shows mean replace errors for the seven groups of the two overshadowing experiments (original and salient landmark) during the learning phase. There was no difference between learning with the landmark and with the boundary alone, nor was there a difference between learning with the original and with the more salient landmark alone. Performance was better when both cues, landmark and boundary, were present than when only one cue was present (see Fig. 2 and the main text for details).

This Article

▶ [Abstract](#)

Services

▶ [Email this article to a colleague](#)

▶ [Alert me to new issues of the journal](#)

▶ [Request Copyright Permission](#)

Citing Articles

▶ [Citing Articles via CrossRef](#)

SI Figure 6

Fig. 6.

Blocking experiments: performance in the initial learning and compound learning phases. Plots show mean distance errors in the initial learning phase (separate plots for the mean of all trials and for the last trial) (*A*) and the compound learning phase of all three blocking experiments (Experiments LL, LB, and BB) (*B*). Initial learning to landmark and boundary occurs at similar rates, and reached near-asymptotic levels from which compound learning produced negligible improvement (see Fig. 3 and the main text for details).

SI Figure 7

Fig. 7.

Differential blocking in objects equidistant to both cues during the compound learning and testing phases. The plots show the mean distance errors in the test phase of the LL and BB blocking experiments for two objects that had the same distance to both cues during the compound learning and testing phases. Analyses revealed the same statistical pattern as in the analyses including all objects: landmark L1 blocked landmark L2 and vice versa (Experiment LL); and boundaries B1 and B2 did not block each other (Experiment BB). These data show that the differential blocking effects are not due to differential proximity to the two cues.

SI Text

Methods

Details of Procedure and Design.

Within trial delay and intertrial interval (ITI) durations were selected randomly for each trial with the constraint that the overall distribution (and by this the overall duration) of delay and ITI times were identical for each object. Within experimental blocks, trials were separated into miniblocks (unknown to the participants), each object occurring once per miniblock. Object order within the miniblocks was randomized with the constraint that no object appeared in two successive trials across miniblock boundaries. Trials across miniblock boundaries were separated by the same average ITI as the other trials. In the overshadowing experiments, participants started each trial at a random location and orientation. In the blocking experiments, they started each trial from one of four positions near to the edge of the arena, facing inwards so that all cues were visible. These four starting positions were randomized within miniblocks with the constraint that no starting position was used twice in succession.

Training.

To prepare the participants for the experiment and to familiarize them with the task, participants underwent training immediately before the experiment. To prevent any unspecific learning effects the training environment differed from the experimental environment with respect to the shape of the environment (square vs. circle), the ground (desert-like vs. grass), the intramaze landmark (bush vs. traffic cone), the global configuration of the background mountains and the objects. Participants learned the initial position of four training objects and then performed two trials per object.

Analyses and Control Experiments

Overshadowing Experiments. *Differential overshadowing is not due to differential landmark-salience.* Might the landmark be somehow less salient than the boundary? To test this, we replicated the overshadowing groups who experienced the landmark (groups L, LB1, and LB2) using a more salient landmark. 36 participants separated into three groups (new groups L', LB1', and LB2') performed the identical task and sequence of phases as the corresponding groups L, LB1, and LB2 in the basic experiment: learning with landmark and boundary (new groups LB1' and LB2') or solely with the landmark (new group L') and then tested with either solely the landmark (new groups L' and LB1') or solely the boundary (new group LB2'). The only difference was that the new groups were presented with a different landmark, a traffic cone three times taller than the original traffic cone (but not wider, so as not to affect participant's estimate of distance from the landmark).

The same statistical pattern as in the basic experiment was observed in an ANOVA on groups LB1', L', LB2', and B: a significant interaction between both factors ($F_{1,44} = 5.03$, $P < 0.05$) with the main effects "test cue" ($F_{1,44} = 1.64$, $P > 0.2$) and "learning situation" ($F_{1,44} = 3.00$, $P > 0.08$) not reaching significance. As in the basic experiment, posthoc t tests revealed that the boundary overshadows the landmark ($t_{22} = 2.30$, $P < 0.05$; cf. replacement error for groups LB1' and L') but not vice versa ($t_{22} < 1$, $P > 0.6$; cf.

replacement error for groups LB2' and B). See SI Fig. 4. Note that the fact that performance with landmark alone is not different to that with boundary alone (cf. groups L and B, see below and SI Fig. 5) is further evidence that the landmark is as salient as the boundary (see also the blocking experiments in which performance is actually slightly better to the landmark than to the boundary during initial learning).

Differential overshadowing is not due to performance differences in the learning phase. Mean performance during the learning phase was better with the compound cue LB than with a single cue (L or B), and was not different for groups L and B ($t_{22} < 1$, $P > 0.7$), in the basic experiment (same ANOVA as test phase: effect of learning situation $F_{1,44} = 8.45$, $P < 0.01$; no effect of test cue or interaction $F_{1,44} < 1$), and in the landmark-salience experiment (ANOVA on groups LB1', L', LB2', and B: effect of learning situation $F_{1,44} = 14.01$, $P < 0.001$, no effects of test cue or interaction $F_{1,44} < 1$). Thus, the differential overshadowing effects can not be explained by performance differences in the learning phase. In addition, the better performance with both cues present than only with one indicates that one cue (e.g. L) is not simply ignored when in the presence of the other (e.g. B).

There were also no significant performance differences in the learning phase between the groups with the original landmark and those with the more-salient landmark (group LB1 vs. group LB1': $t_{22} < 1$, $P > 0.5$; group LB2 vs. group LB2': $t_{22} = 1.22$, $P > 0.2$; group L vs. group L': $t_{22} < 1$, $P > 0.8$). Finally, there were also no performance differences between groups L' and B ($t_{22} < 1$, $P > 0.6$). See SI Fig. 5.

Blocking Experiments. *Basic analyses of the test phase.* Blocking effects were analyzed with an object-type (cue 1-related vs. cue 2-related) \times test cue (test cue 1 vs. test cue 2) ANOVA. In Experiment LL a significant object-type \times test cue interaction was observed ($F_{1,15} = 29.93$, $P < 0.001$). Neither a main effect of object type ($F_{1,15} = 1.66$, $P > 0.20$) nor a main effect of test cue ($F_{1,15} = 2.66$, $P > 0.10$) reached significance. In Experiment LB the ANOVA revealed a significant interaction effect ($F_{1,15} = 72.96$, $P < 0.001$), a significant main effect of object type ($F_{1,15} = 42.60$, $P < 0.001$), but no main effect of test cue ($F_{1,15} = 1.08$, $P > 0.30$). In Experiment BB neither the interaction term ($F_{1,15} < 1$, $P > 0.35$) nor the main effects (object-type: $F_{1,15} < 1$, $P > 0.55$; test cue: $F_{1,15} = 2.36$, $P > 0.1$) reached significance. We report the posthoc t tests between both object types in the main text, separately for each test cue, showing the effect of blocking of learning to landmarks in Experiments LL and LB (i.e., significantly reduced performance for objects paired with the landmark during compound learning but not initial learning). See Fig. 3.

Performance in initial learning and compound learning phases. We conducted ANOVAs for both learning phases with the factors object-type and learning-phase (initial learning vs. compound learning), see SI Fig. 6. In all three

experiments mean performance was better in the compound learning phase than mean performance in the initial learning phase for both object types, but not significantly better than in the last trial of initial learning (with the exception of marginally significant additional learning to one of the landmarks in Experiment LL, see below). Mean performance did not differ between objects paired with cues of the same type in either phase, but there was a slight overall performance advantage for landmark-related objects compared to boundary-related objects. See details below, and SI Fig. 6.

In Experiment LL, we found a significant effect of learning-phase ($F_{1,15} = 80.31, P < 0.001$). Neither the object-type ($F_{1,15} = 3.15, P > 0.09$) nor the interaction ($F_{1,15} < 1, P > 0.8$) reached significance. Posthoc t tests revealed that performance for both object types (L1-related and L2-related objects) did not differ during initial learning ($t_{15} < 1, P > 0.4$) or compound learning ($t_{15} = 2.08, P > 0.05$). In Experiment LB, the ANOVA revealed significant main effects of learning-phase ($F_{1,15} = 88.21, P < 0.001$) and object-type ($F_{1,15} = 36.78, P < 0.001$) and a nonsignificant interaction term ($F_{1,15} < 1$). Posthoc t tests revealed higher performance for the landmark-related objects relative to the boundary-related objects during initial learning ($t_{15} = 4.40, P < 0.001$) and compound learning ($t_{15} = 4.35, P < 0.001$), see further analyses below. In Experiment BB, we observed a main effect of learning phase ($F_{1,15} = 63.67, P < 0.001$). The main effects of object-type ($F_{1,15} < 1, P > 0.7$) and the interaction ($F_{1,15} = 1.05, P > 0.3$) did not reach significance. Posthoc t tests showed that performance for both object types (B1-related and B2-related objects) did not differ during initial learning ($t_{15} < 1, P > 0.7$) or compound learning ($t_{15} < 1, P > 0.3$).

Comparison of the last trial of initial learning with the mean performance during the compound phase revealed the following pattern: Experiment LL: L1-related objects: $t_{15} = 1.76, P > 0.09$; L2-related objects: $t_{15} = 2.33, P < 0.05$. Experiment LB: L-related objects: $t_{15} = 1.11, P > 0.25$; B-related objects: $t_{15} < 1, P > 0.35$. Experiment BB: B1-related objects: $t_{15} = 1.90, P > 0.07$; B2-related objects: $t_{15} = 1.60, P > 0.1$.

Learning to landmark and boundary occur at similar rates (differential blocking is not due to differential learning rates).

To compare the rates of learning relative to landmark and boundary in Experiment LB we fitted an exponential function to the performance data over the 14 trials of initial learning of each participant, separately for landmark-related and boundary-related objects. The performance data were well approximated by the exponential function. We then compared the "exponential slope" of the fit function between landmark-related and boundary-related objects. We did not find any difference between both object types ($t_{15} = 1.09, P > 0.25$), indicating that learning relative to landmark and boundary occurred at similar rates.

Finally, we checked whether participants' better performance on landmark-related than boundary-related objects during initial learning explained any of the blocking effect seen when testing with the landmark alone (in which performance is much better for landmark-related than boundary-related objects), but found no sign of correlation across participants (performance difference in initial learning versus performance difference at test: Pearson's $r = -0.15, P > 0.5$).

Differential blocking effects are not due to differential cue proximity. We conducted an additional analysis separately for pairs of objects with the same distance to both cues during the compound and test phases (one paired with either cue during initial learning) in Experiments LL and BB. These analyses revealed the same statistical pattern as the basic analyses including all objects: Landmark L1 blocked landmark L2 and vice versa (Experiment LL; test L1: $t_{14} = 5.47, P < 0.001$; test L2: $t_{14} = 5.25, P < 0.001$), and boundaries B1 and B2 did not block each other (Experiment BB; test B1: $t_{14} < 1, P > 0.5$; test B2: $t_{14} = 1.23, P > 0.2$). See SI Fig. 7.

This Article

▶ [Abstract](#)

Services

▶ [Email this article to a colleague](#)

▶ [Alert me to new issues of the journal](#)

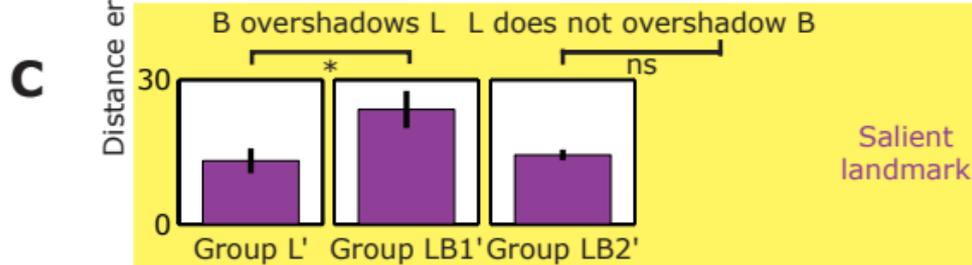
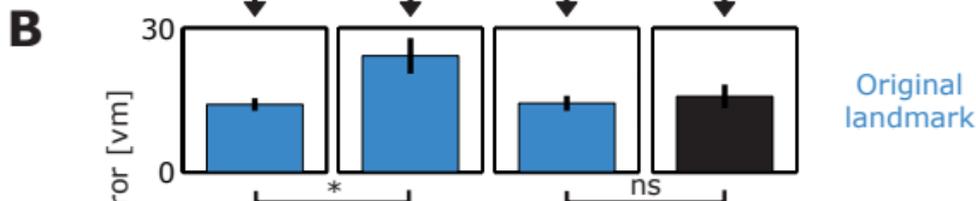
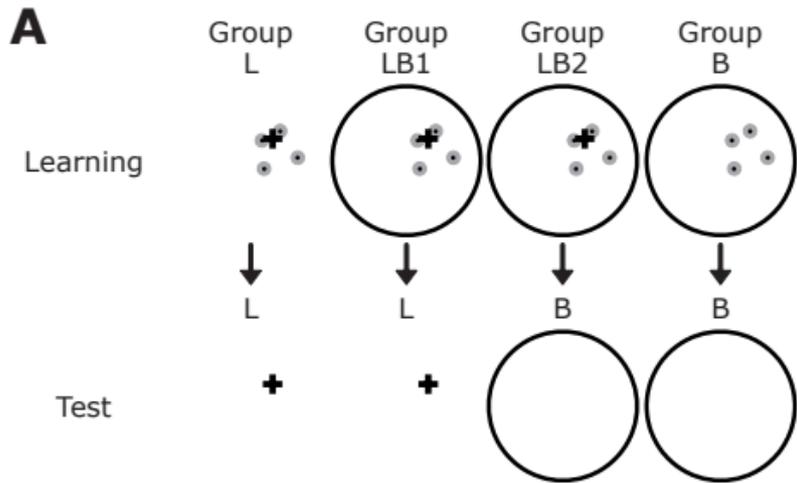
▶ [Request Copyright Permission](#)

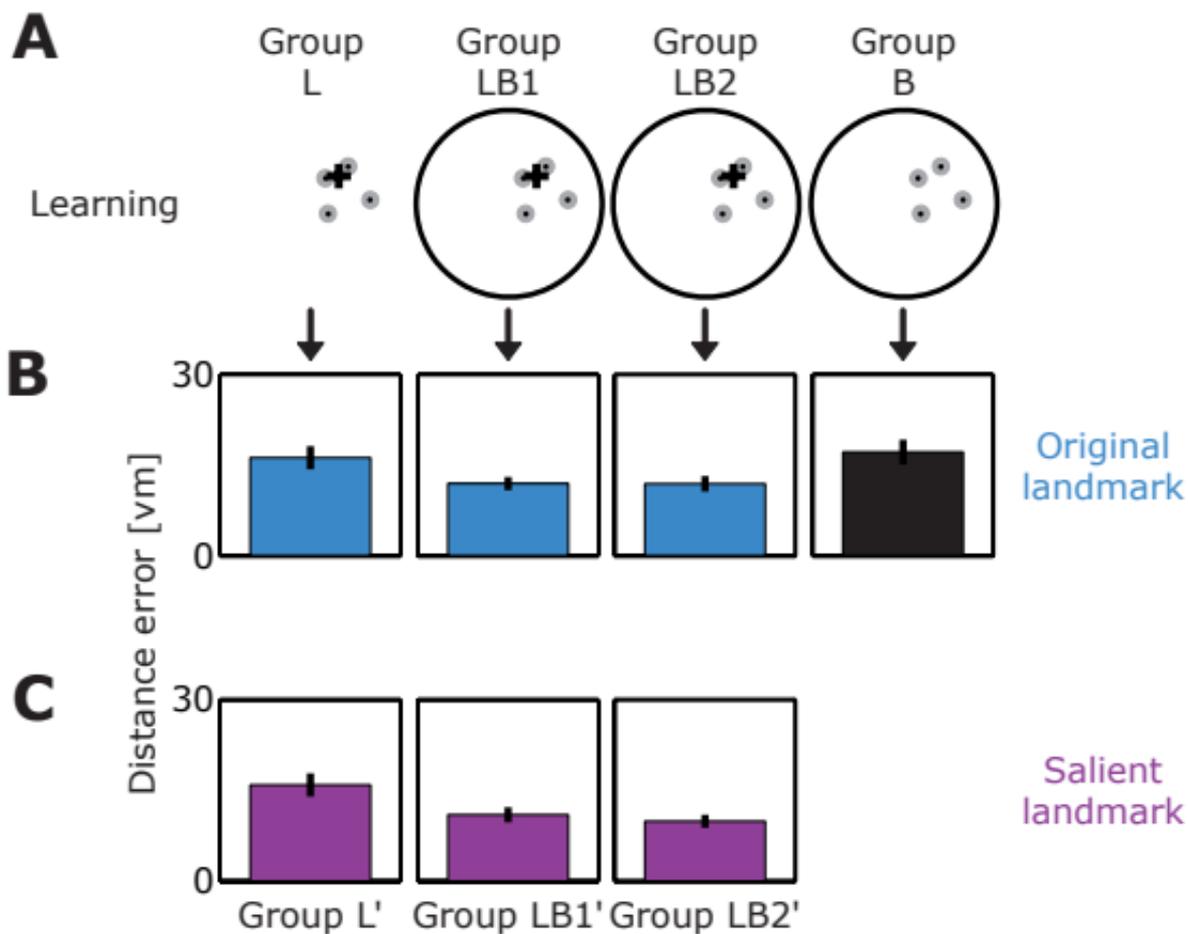
Citing Articles

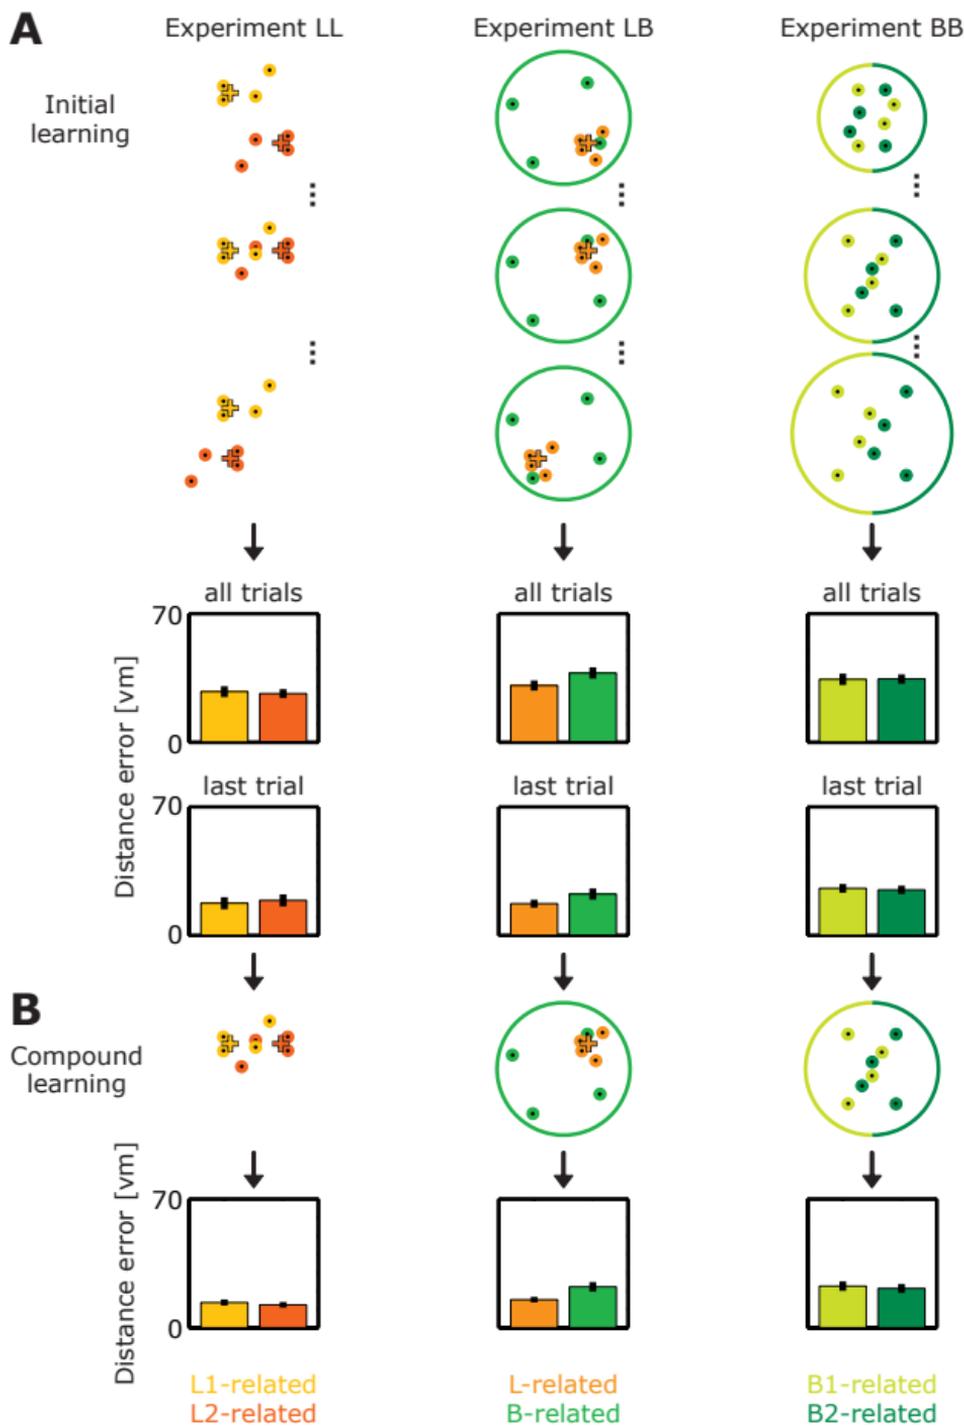
▶ [Citing Articles via CrossRef](#)

[Current Issue](#) | [Archives](#) | [Online Submission](#) | [Info for Authors](#) | [Editorial Board](#) | [About](#)
[Subscribe](#) | [Advertise](#) | [Contact](#) | [Site Map](#)

[Copyright © 2008 by the National Academy of Sciences](#)







Compound learning



Test L1 ↓ Test L2



Test B1 ↓ Test B2

Tests



Distance error [μm]

80

0

L2 blocks L1

L1 blocks L2

B2 does not block B1

B1 does not block B2

L1-related
L2-related

B1-related
B2-related

**

**

ns

ns

