Main Manuscript for

Memory-based predictions prime perceptual judgments across head turns in immersive, real-world scenes

Anna Mynick[1], Adam Steel[1], Adithi Jayaraman[1], Thomas L. Botch[1], Allie Burrows[1], Caroline Robertson[1]

[1]Dartmouth College, Hanover NH

Anna Mynick

Email: anna.r.mynick.gr@dartmouth.edu

This PDF file includes:

> Main Text
> Figures 1 to 4
> Tables 1 to 1

Abstract

Our visual world extends beyond our field of view, requiring eye, head, and body movements to see completely. What cognitive mechanisms support our seamless interaction with real-world environments despite our limited field of view? We hypothesized that memory-based predictions enable humans to anticipate upcoming views across head turns, facilitating efficient perceptual behavior. We tested this hypothesis using immersive, head-mounted virtual reality (VR). After learning a set of immersive real-world environments in VR, we tested whether participants (N=104) used memory-based predictions to support scene judgments across head-turns. Specifically, participants in immersive VR were briefly primed with a single view from a studied environment, and then turned left or right to judge whether an adjacent view was navigable. We found that participants responded faster when they were primed with images from the same (compared to neutral or different) environments. Priming was memory based: it only occurred in learned (vs. novel) environments. Further, consistent with a role in supporting active vision, priming only occurred in the direction of intended head turns. Finally, priming only occurred when environment views were presented in their learned locations, suggesting that predicted content is structured by the spatial coordinates of the original environment. Thus, memory-based predictions facilitate rapid perception across large-scale visual actions, such as head and body movements, and may be critical for efficient behavior in complex immersive environments.

Significance Statement

The world we experience is seamless and immersive. Yet, each view of this environment is discrete, fleeting, and limited. A central puzzle for human neuroscience is: what cognitive mechanisms enable us to anticipate and efficiently engage with our immersive environment despite our limited field of view? Here, we tested how efficient perception is enacted across large scale visual actions: head turns. Participants (N=104) learned real-world, 360º environments in immersive virtual reality, then head-turned across scene views to make rapid perceptual judgments. Overall, we found that memory for familiar environments supplies predictions of upcoming scene views across head turns to facilitate seamless perception. Together, these studies provide key insight into how contextual information from memory informs active perceptual behavior.

Main Text

Introduction

Humans constantly make predictions to support ongoing behavior: we predict what will be inside our mailbox, which people will be in the room we are entering, and whether there is

time to get through a yellow light. These predictions are often shaped by our memories of past experiences (1). In traditional studies of memory-based predictions, participants learn associations between stimuli presented sequentially on a screen and are trained to predict upcoming target images (e.g., expect C, given A and B) (2–4). Once learned, these associations facilitate sensory processing of upcoming stimuli, speeding detection (5, 6) and improving recognition of anticipated percepts (7, 8). However, despite their presumed importance in everyday cognition, the content and form of memory-based predictions during naturalistic, active vision is unclear. As we look around the world, exchanging visual content in the current view for content in the next, what information is predicted across head turns?

Several lines of evidence suggest that the visual system predicts the perceptual consequences of visual actions. During saccades, associations between foveal and peripheral percepts support perceptual continuity. For instance, before a saccade to a target stimulus, the pre-saccadic center of gaze already shows enhanced sensitivity to the post-saccadic target's stimulus features, linking processing of the current view to the upcoming one (9). Memory for a previously foveal percept also influences peripheral vision. For example, past work has shown that the percept of a stimulus in the periphery can become biased to resemble the same stimulus' foveal percept (10, 11). These behavioral findings align closely with neural recordings. For example, cells in retinotopic cortex anticipate the visual consequences of saccades by preemptively responding to objects that will fall into their receptive fields following a saccade (12–14). Thus, it is clear that the visual system predicts the consequences of small-scale visual actions — i.e., saccades. But, prediction during saccades is relatively straightforward, because the visual information being predicted is, of course, already available to draw from in the visual periphery. An open question is how predictive vision generalizes to large-scale visual actions like head and body movements, where upcoming visual information is fully out of view.

Two solutions could plausibly address how the brain generates predictions in immersive, real-world contexts. First, schema-level knowledge about the type of environment could be used to extrapolate visual content beyond the current field of view. Such scene extrapolation has been proposed as a mechanism supporting the phenomenon of boundary extension (15–18 but see 19–21). Second, memory of the broader visuospatial environment could supply specific predictions about the content and layout of upcoming views that will result from visual actions. Indeed, recent work has shown that views from the same environment become associated in the brain, which could support this type of behavioral prediction (22, 23).

Here, we sought to understand how memory supports naturalistic scene perception across large-scale visual actions in immersive environments. To do this, we developed an experimental paradigm that used head-mounted virtual reality (VR) to examine memory-based predictions in immersive, real-world scenes. After learning a set of 360° real-world

scenes, participants were primed with a single scene view before head-turning to another view from the studied scenes to make an ecologically-relevant perceptual judgment.

To preview our results, we found that memory-based predictions are generated for upcoming scene views across head turns (Experiment 1), and that these predictions depend on memory for out-of-sight views of the current scene (Experiment 2). Crucially, these predictions appear to occur in service of action plans (Experiment 3). In addition, predicted content appears to rely on highly specific spatial knowledge of a scene, adhering to the spatial coordinates of the original environment at a scene-wide (Experiment 4) and local-view (Experiment 5) level. Together, our results suggest a role for memory-based predictions in facilitating rapid perception as we sample the visual world around us.

Results

Upcoming scene views are primed across head turns

We first asked whether memory-based predictions are generated for upcoming scene views across head turns in immersive environments (Experiment 1). In the Training Phase, participants (N=18) learned a set of 360° real-world scenes in virtual reality (VR) (Fig. 1a). To facilitate faster learning, scenes depicted familiar locations around the Dartmouth College campus (Experiment 1 familiarity per participant, M=85.5%, SD=11.29%) (Supplementary Fig. 1 & 2). Scenes fully surrounded participants, so that they needed to use head turns to see a scene completely (Fig.1b). On each trial (20s), participants were simply told to naturally explore each scene (i.e., "look around like you normally would"). Each scene featured an open view (e.g. a view of an open parking lot) and a closed view (e.g. a view of a building's façade) on either side of an initial facing direction (Fig. 1c). Importantly, participants could accurately recall the position of individual scene views (Explicit Memory Test, M = 93.70±5.80 s.d.; t(17) = 44.17, p < .001) (Supplemental Fig. 3) following the Training Phase.

After the Training Phase, we assessed participants' memory-based predictions in the Priming Test. On each trial, participants saw a full-field scene image ("prime") from a studied scene. Following the prime image, participants turned approximately 90° left or right to make a perceptual judgment (open or closed) on a second scene view ("target") (Fig. 1d). Primes either depicted: (1) a view from the same scene (Same-scene prime), (2) a blank gray rectangle (Neutral prime) or (3) a view from a different scene (Different-scene prime) (Fig. 1e). We considered memory retrieved in response to the prime image the "predicted content", and we assessed the behavioral impact of this prediction by comparing response time across the priming conditions. We hypothesized that if participants generate memory-based predictions for upcoming scene views, response times would be faster after Same-scene primes compared with Neutral or Different-scene

primes. We compared the differences across conditions using repeated measures ANOVAs with priming condition (Same-scene / Neutral / Different-scene) as a factor.

Our results revealed striking evidence for memory-based predictions across head turns (Main effect of priming condition: $F(2,34) = 15.80$, $p = <.001$, $\eta_p^2 = 0.48$) (Fig. 2b). Critically, perceptual judgments were faster after Same-scene primes compared to Neutral primes ($t(17) = -2.77$, $p = .01 * 3$ tests $= p_{corr} = .04$, $d = -0.43$), and were slower after Different-scene primes compared to Neutral primes ($t(17) = -3.37$, $p_{corr} = .01$, $d = 0.57$). Explicit Memory accuracy (see Methods) covaried marginally with the extent of the priming effect (i.e. difference in response time between Same- and Different-scene primes ($r(16) = .42$, $p = .079$)) (Fig. 2c). Taken together, these results demonstrate that upcoming scene views are primed across head turns in familiar environments and suggest that the extent of priming may be related to memory for the spatial structure of tested scenes.

We next tested whether priming across head turns requires memory for the specific scene, or whether predictive information arises from scene extrapolation (akin to boundary extension (Intraub & Richardson, 1989). Participants in Experiment 2 (N=20) performed the same Priming Test as in Experiment 1 on a set of novel, unfamiliar scenes with no Training Phase (Supplementary Fig. 4 for scene images; Methods Table 1). Compared to familiar scenes (Experiment 1), response times were slower overall for unfamiliar scenes (Experiment 2) (main effect of experiment, $F(1,108) = 51.54$, $p = <.001$, $\eta^2 = 0.32$). Critically, in unfamiliar scenes, there was no difference in response times across priming conditions (No main effect of priming condition: $F(2, 38) = 0.50$, $p = .61$, $\eta_p^2 = 0.03$) (Fig. 2d), suggesting that no predictive information was available to speed up perceptual judgments. Thus, our results indicate that behavioral priming across head turns requires memory-based prediction and does not arise from scene extrapolation.

Action plans bias direction of priming

So far, our results suggest that memory-based predictions support efficient scene perception across head turns. In Experiment 3, we asked whether predictions specifically serve intended action plans, or whether predictions are generated for all scene views surrounding the current one, regardless of action plans. Participants (N=26) completed a version of our Priming Test in which a brief left or right arrow (300 ms) preceded the prime images (Same-scene or Neutral) (Fig. 3a), indicating a direction to plan a head turn in. Arrows correctly indicated the upcoming target location ("valid arrow") on two-thirds of trials, and were incorrect ("invalid arrow") in the remaining one-third of trials. If predictions serve intended action plans, a planned head turn should strengthen the extent of priming in the intended – but not unintended – direction.

Our results were consistent with this hypothesis: Priming was stronger in the direction of intended head turns, evidenced by a significant interaction between (Same/Neutral) and arrow (valid/invalid) conditions (Two-way repeated measures ANOVA $F(1, 25) = 4.75$, $p =$

.04, η_p^2 = 0.16). Follow up tests indicated that Same-scene primes resulted in faster response times than Neutral primes in trials with valid arrows (t(25) = -4.26, p_{corr} = .002, d = -0.29) but not invalid arrows (t(25) = -0.65, p_{corr} = 3.11, d = 0.07) (Fig. 3b). This finding suggests that planned actions evoke predictive visual representations that facilitate efficient perceptual judgments in immersive environments.

Predicted content obeys a scene's global spatial structure

Having established that visual information is primed across head turns for Same-scene views, we next turned to the question of the predicted content's format. We addressed this question using two simple manipulations that systemically disrupted the match between the visual features of the upcoming scene view and the location of the scene view in space.

We first asked whether scene's immersive spatial structure determines the associations between discrete views in memory, or whether the associations are based on purely associative information. For instance, using purely associative information, a prime image might prime both neighboring (left/right) scene views equally because both are likely equally viewed in association with the prime image as a result of their spatial proximity. Participants in Experiment 4 (N=18) completed the Training Phase as in Experiment 1, and then performed the original Priming Test (i.e., no arrows) with three conditions: Same-scene, Neutral, and Globally-reversed trials (Fig. 4a; Methods Table 1). On Globally-reversed trials, the prime and target were drawn from the same scene, but the target image depicted the view 180º opposite the expected view, disrupting the global structure of the scene. If predicted content is purely associative, then priming should occur in both Same scene and Globally-reversed trials. Alternatively, if predictions are structured by the global spatial coordinates of the environment, then Globally-reversed trials should interfere with priming.

Indeed, our results suggest that predicted content is structured by a scene's global spatial coordinates (Main effect of priming condition: F(2, 34) = 23.01, p < .001, η_p^2 = 0.58) (Fig. 4b). As expected, we replicated the advantage for Same-scene primes relative to Neutral identified for familiar scenes in Experiment 1 (t(17) = -2.78, p_{corr} = .04, d = -0.39). Critically, predicted content does not apply to any view from the same scene: response times suffered when targets were Globally-reversed compared to the Neutral condition (t(17) = -5.44, p_{corr} < .001, d = 0.74). The extent of priming (quantified by the difference in response time between Same-scene and Globally-reversed trials) covaried with participants 'ability to recall the position of individual scene views (i.e., Explicit Memory Test performance; r(16) = .62, p = .006) (Fig. 4c), suggesting that the extent of priming may be related to memory for the spatial structure of tested scenes.

Predicted content obeys the local spatial structure of individual scene views

6

Based on the evidence that predicted content is structured by a scene's global spatial coordinates, we lastly sought to test the granularity of the predicted content within a scene view. In Experiment 5 (N=22), we tested whether memory-based predictions carry information about the local structure of content within a scene view (e.g. the configuration of objects, buildings, and navigational affordances), by introducing Locally-reversed trials, in which the target images were flipped across the vertical axis (Fig. 4d). This manipulation disrupted the granular spatial configuration of content within the scene view, while keeping its open/closed structure intact. Participants completed the Training Phase as in Experiment 1, and then performed the Priming Test with Same-scene, Neutral and Locally-reversed trials (Methods Table 1). We expected that if predicted content carries granular information about the configuration of an upcoming scene view, Locally-reversed trials should reduce priming. Alternatively, if the prediction reflects the content of the scene view without spatial organization, there should be no impact of local reversal on priming.

Our results support the conclusion that priming arises from prediction of spatially-organized information about the upcoming scene view. Priming was sensitive to the local structure of a scene view: response times differed significantly across priming conditions (repeated measures ANOVA ($F(2, 42) = 6.99$, $p = .002$, $\eta_p^2 = 0.25$) (Fig. 4e). As expected, we again replicated the advantage for Same-scene relative to Neutral primes ($t(21) = -3.28$, $p_{corr} = .011$, $d = -0.40$). Critically, we found no priming effect for Locally-reversed primes relative to Neutral primes ($t(21) = -0.71$, $p_{corr} = 1.46$, $d = 0.12$), indicating that priming is sensitive to the granular configuration of a scene view rather than to content of the scene view regardless of spatial organization. More broadly, this result suggests that memory-based predictions of upcoming scene views carry detailed information about the local spatial structure of individual scene views. Interestingly, unlike previous experiments, the extent of priming (quantified by the difference in response time between Same-scene and Locally-reversed trials) did not significantly covary with Explicit Memory accuracy on this task ($r(20) = .34$, $p = .12$), which could reflect a difference in mnemonic demands between the local and global reversal manipulations (Fig. 4f).

Discussion

Here, we show that memory-based predictions of upcoming, out-of-sight scene views prime perception of those views across head turns to enable rapid, ecologically-relevant perceptual judgments. This priming only occurs in familiar scenes, signaling its dependence on memory, and is skewed in the direction of planned head turns, underlining its utility for active vision. Interestingly, predicted content from memory seems to retain its spatial structure at both a global (scene-wide) and local (within-view) level. Based on these results, we propose that memory-based predictions could allow us to overcome our limited field of view by supporting perception across large-scale visual actions.

Our work significantly advances our understanding of how predictive processes influence perception across large-scale visual actions by showing that memory-based predictions are generated for upcoming scene views across head turns in immersive environments (Experiments 1-3). Previous studies of predictive vision have primarily used basic stimuli (e.g. oriented gratings or single objects) to examine small-scale visual actions like saccades within a single field of view. In these contexts, others have found evidence for predictive vision in foveal and peripheral vision (9–14). For example, before making a saccade, the center of gaze becomes more sensitive to the visual features of the stimulus that the saccade will bring into view (9). Likewise, across saccades, the appearances of peripheral objects become biased toward their associated foveal percepts, which is thought to smooth the transition between low-acuity peripheral- and high-acuity foveal vision (10, 11). Interestingly, saccades are necessary to evoke this predictive bias – biased perception does not occur for percepts that become associated without a saccade (10), which mirrors our finding that intended actions modulate the extent of priming. Notably, prediction across saccades differs fundamentally from predictions across head turns: across saccades, information in the periphery can be used to inform predictions (10), while head turns often bring new information into the field of view and should therefore require additional information to inform predictions. Indeed, our results indicate that in immersive contexts, memory supports predictions of upcoming scene views across head turns. Thus, an open question is the degree of overlap between the cognitive mechanisms underpinning predictions across saccades (operating within a field of view) versus across large-scale visual actions like head turns (operating across many fields of view).

As for predicting information across head turns, our finding that priming occurs across head turns for familiar (but not unfamiliar) scenes is surprising. This is because two known phenomena which suggest that scene content can be extrapolated from the visual features of the current scene view – boundary extension and scene layout extrapolation (24–26 but see 27) – occur robustly in response to unfamiliar images. These findings argue against the need for memory in predictive scene perception. Critically, boundary extension and scene layout extrapolation are fundamentally different from predictions across large-scale visual actions like head turns: boundary extension and scene layout extrapolation operate only within a small spatial window around the scene (16, 24). Thus, the neural mechanisms underpinning these forms of scene extrapolation are unlikely to be universal "predictive" phenomena (i.e., pervasive cognitive mechanism supporting predictive vision across head turns). In addition, the nature of boundary extension appears to depend on image features (19–21). For example, the extent of boundary extension is influenced by the expanse (19), blur (21) and depth of field (20). This susceptibility to image features could indicate that boundary extension operates more on image features than image content, which juxtaposes with our finding that memory-based priming across head turns relies on scene content above and beyond scene features (Experiments 4 and 5). Taken together, the findings described in the present study likely complement the cognitive phenomena of boundary extension or scene layout extrapolation, offering a memory-

based cognitive mechanism to support efficient perceptual action across head turns. We hypothesize that this cognitive mechanism could arise from associations between scene views from the same environment in the brain during learning (22, 23).

Our finding that perceptual judgments are primed across head turns suggests that memory-based predictions play a role in facilitating perception during active vision, and raises the question of the format and content of those predictions. Two of our findings speak directly to this issue. First, we showed evidence that predictions are mapped to the spatial coordinate frame in which they were learned (Experiment 4). Second, we found evidence that within scene views, memory-based predictions contain granular information about the configuration of the scene view as opposed to a simple gist-like representation (Experiment 5). These findings agree with past behavioral and neuroimaging work demonstrating the surprisingly fine-grained detail of our memories for scene images (28–30). Past work has also highlighted an interplay between memory and spatial position. For instance, memory for object details in scenes improves when objects are tested in their learned locations within a scene (31). These highly-specific spatial and visual details could help support diverse actions including navigation (32, 33) and memory-based visual search in immersive environments (34). Though the present work provides evidence for detailed memory-based predictions to the degree tested here, it is unlikely that memory-based predictions are "photorealistic" given the tendency for stimuli to be to compressed (35) or distorted (36–39) in memory. Perhaps instead, stable aspects of a scene view (e.g. locations of paths and buildings) are predicted with greater specificity, while more malleable types of visual information (e.g. small moveable objects) are represented less faithfully.

A key challenge for memory-based predictions during active vision is that perception and memory often operate in distinct spatial reference frames (40). For example, an agent who sees a mud puddle in the center of her retinotopic coordinate frame must convert its position to spatiotopic coordinates for memory of the puddle to be useful after she looks away; otherwise, she risks stepping in the puddle once it is out of view. Importantly, because memory-based predictions have primarily been studied in screen-based contexts (2–6, 41–43), it was previously unknown whether they are in fact implemented in a spatiotopic reference frame. Our findings show that memory-based predictions are mapped to the spatial coordinate frame of the global environment. This aligns with recent virtual reality (VR) studies indicating that objects maintained in immersive working memory are primarily remembered based on their location in the environment's global coordinate frame (rather than on their egocentric position relative to the viewer during encoding) (40, 44, 45). Together with our findings, these studies underscore the role of memory in supporting ecologically-relevant perceptual judgments (1, 45) and highlight the promise of head mounted VR as a tool for studying how perceptual information is represented – and predicted – during naturalistic, active perception (23, 34, 46, 47).

To conclude, here we developed a novel VR priming paradigm to investigate the role of memory-based predictions during naturalistic behavior. Our results indicate that memory-based predictions of specific, upcoming scene views enable rapid perceptual judgments across head turns. We propose that memory-based predictions may be used to overcome our limited field of view by facilitating visual perception across large-scale visual actions (head turns and body movements). Future work should aim to understand how memory-based predictions retain their spatial structure, perhaps through interplay between the perceptual and head-direction systems.

Materials and Methods

Participants
150 participants were recruited across five experiments in this study. Adult participants were recruited from Dartmouth College or the local community and received payment ($15/hour) or course credit for participation. Participants in Experiment 2, 3, and 5 were recruited based on the following criteria: 1) normal or corrected-to-normal vision 2) no neurological or psychiatric conditions 3) no history of epilepsy. Participants in Experiment 1 and 4 participated for course participation. Participants in Experiment 1 and 4 completed the experiment remotely from their homes using virtual reality headsets that were delivered to them in the mail. Participants in Experiments 2 and 3 were mixed between remote testing and in-person testing at Dartmouth College in the following amounts: Experiment 2: 7 remote, 17 in-person; Experiment 3: 4 remote, 42 in-person. Participants in Experiment 5 completed the experiment in person. Written consent was obtained from all participants in accordance with a protocol approved by the Dartmouth College Institutional Review Board. 46 participants were excluded due to data quality concerns associated with remote VR data collection (see Inclusion criteria section below), leaving data from 104 participants in the final analyses: 18 in Experiment 1, 20 in Experiment 2, 26 in Experiment 3, 18 in Experiment 4 and 22 in Experiment 5.

Materials

Hardware
This study took place entirely in head-mounted virtual reality (VR). Stimuli were displayed through head-mounted virtual reality (Oculus Quest 2, resolution 1832x1920 per eye, approximately 97º horizontal by 93º vertical field of view, 120 Hz refresh rate) using a project designed in Unity (www.unity3d.com) with custom scripts written in C#, and deployed to headsets using mobile device management software ManageXR (www.managexr.com). Experimental data was transferred from the HMD to lab servers via custom data transfer pipeline written in C# and PHP (34).

Stimuli: Experiments 1,3, 4, and 5
Stimuli were comprised of 18 real-world 360º panoramic scenes drawn from unique locations around Dartmouth's campus downloaded from Google Maps (9 indoor/outdoor

locations). We selected scenes with a specific layout: each contained an open side (e.g. a field or parking lot) and a closed side (e.g. a wall or facade) on opposite sides of an initial facing direction (i.e. separated by 180º) (Supplementary Fig. 1).

Stimuli: Experiment 2
36 scenes in Experiment 2 were chosen on matched criteria to those in Experiments 1,3,4, and 5 except that the scenes were selected to be unfamiliar to participants (Supplementary Fig. 4). These scenes were drawn from diverse locations around the world via Google Maps.

Procedure

Our study consisted of 5 experiments that each tested a specific question related to memory-based predictions in immersive environments.

In the sections that follow, we describe the procedure for each experiment. We first outline the experimental timeline for all experiments. Next, we outline the Training Phase and the Priming Tests. Finally, we outline the experimental manipulations used in each Priming Test. See Methods Table 1 for an overview of the procedure across all experiments.

All experiments used a common *Priming Test* to evaluate the presence or absence of prediction-based facilitation of an ecologically-relevant perceptual judgment. Experiments 1, 3, 4, and 5 investigated the nature of memory-based predictions using a priming test. To this end, these experiments were conducted over two consecutive days. On Day 1, participants studied the scenes in detail (Training Phase). Then, they received instructions on distinguishing between open and closed scene views (Open/Closed Instructions). On Day 2, participants performed the Priming Test. Experiment 2 specifically investigated whether priming occurred in unfamiliar environments. For this reason, Experiment 2 took place on a single day in which participants performed the Priming Test on unfamiliar scenes.

Training Phase (Experiments 1, 3, 4, 5)
During the Training Phrase, participants actively viewed the 18 immersive photorealistic scenes taken from the area around Dartmouth's campus for 20 seconds each (see Supplemental Figure 1). The open/closed sides of a scene were always positioned to participants' left/right relative to their starting facing direction (presentation side counterbalanced), and the back-most 90º of each scene was occluded to provide a clear sense of left/center/right in each scene. Participants were asked to "complete this study standing at a desk or countertop where you can comfortably reach your keypad." For each scene, participants were instructed to "Look around like you normally would". The instructions told participants that there would be a "gray wall" (occluder) behind them in each scene, and that they only needed to explore what wasn't blocked by the wall. After viewing each scene, they indicated whether the scene was familiar in real life via wireless

keypad (1=Yes, 2=No, 3=Unsure). In Experiment 3, 3 participants were not included in the calculation of percent familiarity (Supplementary Fig. 2) because of missing data due to technical challenges with remote data collection.

On Day 2, participants completed a refresher training in which they viewed the 18 scenes again for 20 seconds each.

Open/Closed Instructions

Our primary interest was participants' ability to make an ecologically-relevant behavioral judgment: was a scene view navigable or not. To this end, after viewing each scene, we instructed participants how to distinguish between navigable vs. non-navigable (i.e., open vs. closed) scene views. Before receiving these instructions, participants were informed that open/closed judgments would be performed in a subsequent part of the experiment. Our specific instructions were: "If a snapshot shows an open space or pathway where you could walk into the scene and move forward freely, that snapshot is Open. If a snapshot shows something blocking your way – something that would prevent you from walking into the scene and moving forward freely – that snapshot is Closed." We then showed participants several examples of open versus closed scene views in novel, unstudied scenes. After viewing these examples, participants completed 6 open/closed judgment practice trials with feedback.

Standard Priming Test Procedure (Experiments 1, 2, 4, 5)

On Day 2, participants completed the Priming Test (Experiments 1, 4, 5: after refresher training; Experiment 2: no pre-training (novel scenes)). Each trial of the Priming Test consisted of four events (Fig. 1): 1) Prime image (110° wide, 180° tall) appears directly in front of the participant (display time: 300 ms); 2) Target image (110° wide, 180° tall) appears to the left or right of the participant; 3) Participant turns 90° to face the target image; 4) Participant indicates whether the target is open or closed via button press on a wireless keypad.

Before testing, participants completed several practice Priming Test trials (5 in Experiments 1,4 and 5; 4 in Experiment 2) with feedback on novel, unstudied scenes. We instructed participants to respond as quickly as they knew the answer, and after each trial, the participants' response time was displayed to emphasize the importance of responding quickly.

To encourage participants to use the prime image to predict the upcoming target, more trials contained Same-scene primes (50% of 144 total trials) than Neutral (25%) or Manipulated (25%) primes. See *Priming Test Stimulus* section for additional details regarding stimulus manipulations.

Direction-cued Priming Test Procedure (Experiment 3)

Experiment 3 tested whether priming was skewed in the direction of intended head turns. For this reason, we modified the standard Priming Test (used in Experiments 1, 2, 4, and 5) to manipulate the anticipated direction needed to turn and face the target image.

The trial sequence in Experiment 3 consisted of 5 stages (Fig. 3): 1) Participants saw an arrow indicating the direction to plan a head turn in to see the target image (300 ms display time). On 2/3 of trials, the arrow correctly indicated the direction of the upcoming target (i.e., was a valid cue), and 1/3 of trials the arrow indicated the incorrect direction (i.e., was invalid. 2). Participants then saw the prime image (300 ms display time; 50% Same-scene, 50% Neutral). 3) The target image (110° wide, 180° tall) appears to the participants' left or right. 4) Participants turn 90° to face the target. 5) Participants respond using a button press.

For Experiment 3, the Priming Test had a 2 x 2 design with 4 priming conditions: (1) Same-scene prime, valid arrow (50%) (2) Neutral prime, valid arrow (16.7%), (3) Same-scene prime, invalid arrow (16.7%), (4) Neutral prime, invalid arrow (16.7%). No Different-scene primes were used in Experiment 3.

Priming Test Stimuli, Experiments 1-5
The stimuli used as prime/target images for each experiment are detailed below.

Experiment 1 Priming Test stimuli

In Experiment 1, prime and target images were from known locations on Dartmouth's campus. The key manipulation was whether prime images came from the same scene as the target image. The prime images were: Same-scene (50% of 144 total trials), Neutral (25%) or Different-scene (25%). The Neutral prime for this study was simply a grey screen.

Experiment 2 Priming Test stimuli

Experiment 2 prime and target images were taken from novel scenes in locations unfamiliar to participants. As in Experiment 1, prime images were: Same-scene primes (50% of 144 total trials), Neutral (25%) or Different-scene (25%) primes. The Neutral prime for this study was simply a grey screen.

Experiment 3 Priming Test stimuli

As described above, prime and target scenes were from known locations on Dartmouth's campus. Prime images in Experiment 3 were either from the Same-scene (50%) or Neutral images (50%). Neutral prime images were scrambled scenes (32 x 16 grid).

Experiment 4 Priming Test stimuli

In Experiment 4, prime and target images were from known locations on Dartmouth's campus. The key manipulation was whether the target image appeared on its learned side (Same-scene prime) or whether it appeared in the location opposite its learned location in

360° space (Globally-reversed). For instance, a target that had appeared on the left during the Training Phase would be presented on the right in a Globally-reversed trial. Priming conditions were distributed following the convention established in Experiment 1: Same-scene primes (50% of 144 total trials), Neutral (25%) or Globally-reversed (25%) primes. The Neutral prime for this study was simply a grey screen.

In Experiment 5, prime and target images were from known locations on Dartmouth's campus. The key manipulation was whether the target image appeared as it was studied (Same-scene prime) or whether it was mirror-reversed along the vertical axis (Locally-reversed). For instance, a target that had appeared on the left during the Training Phase would be presented on the left, but the image itself would be flipped such that content learned on the right side of a scene view was presented on the left. Priming conditions were distributed following the convention established in Experiment 1: Same-scene primes (50% of 144 total trials), Neutral (25%) or Locally-reversed (25%) primes. Neutral prime images were scrambled scenes (32 x 16 grid).

Explicit Memory Test (Experiments 1, 4, 5)

In Experiments 1, 4, and 5, after the Priming Test, we tested participants' explicit memory of each scene's spatial structure. On each trial of the Explicit Memory Test, participants saw a 45º view from a studied scene directly in front of them. Participants then used a wireless keypad to indicate whether that view appeared on the Left, Right, or Center of the scene during the Training Phase (Supplemental Figure 3). Participants were tested on each left/right/center view for each studied scene.

Data Analysis

We analyzed data using custom MATLAB scripts (www.mathworks.com).

Inclusion Critera

Due to the novelty of the VR testing environment for remote testing and the complex trial structure, we applied a stringent trial-level inclusion criteria to ensure data quality.

Response time data from a Priming Test trial was excluded from analysis if the open/closed response was:
1. incorrect;
2. faster than 250ms, potentially indicating insufficient time spent to perform the judgment accurately;
3. outside of the participants' mean response time ± 3 standard deviations.

Trials were also excluded if:
1. the headset failed to transfer head-orientation data;
2. the participant turned before at least 50% of the prime image was within view;

3. the participant did not turn far enough toward the target to view at least 50% of the target scene segment;
4. the participant turned the wrong way (i.e. away from the target image) by more than 15 degrees following the prime (Experiments 1,2,4 and 5 only).

In Experiment 3, in trials where participants turned the wrong way (i.e. away from the target image) by more than 15 degrees:
1. If the participant corrected their head turn to the correct direction, the timepoint at which the participant corrected their head turn served as the start of the corrected trial.
2. If the timepoint at which the participant corrected their head turn could not be identified, the trial was excluded.

Finally, for a given participant, if more than 50% of trials in any priming condition were excluded or if their response time exceeded ±2 standard deviations from the group median, we removed that participant from data analysis entirely.

Data Analysis - Statistics
For all experiments, we calculated the median response time for each Priming Test condition. Response times were calculated from the appearance of the target and included the time taken to turn to face the target.

For priming tests in Experiments 1-5, priming effects across priming conditions were measured using either one-way (Experiments 1,2,4,5) or two-way (Experiment 3) repeated-measures ANOVAs. Where appropriate, we used pairwise Bonferroni-corrected t-tests for post-hoc tests.

Acknowledgments

References

1. N. C. Rust, S. E. Palmer, Remembering the Past to See the Future. *Annu. Rev. Vis. Sci.* **7**, 349–365 (2021).

2. N. C. Hindy, F. Y. Ng, N. B. Turk-Browne, Linking pattern completion in the hippocampus to predictive coding in visual cortex. *Nat. Neurosci.* **19**, 665–667 (2016).

3. P. Kok, L. I. Rait, N. B. Turk-Browne, Content-based Dissociation of Hippocampal Involvement in Prediction. *J. Cogn. Neurosci.* **32**, 527–545 (2020).

4. P. Kok, N. B. Turk-Browne, Associative Prediction of Visual Shape in the Hippocampus. *J. Neurosci.* **38**, 6888–6899 (2018).

5. A. M. Cravo, G. Rohenkohl, K. M. Santos, A. C. Nobre, Temporal Anticipation Based on Memory. *J. Cogn. Neurosci.* **29**, 9 (2017).

6. M. G. Stokes, K. Atherton, E. Z. Patai, A. C. Nobre, Long-term memory prepares neural activity for perception. *Proc. Natl. Acad. Sci.* **109** (2012).

7. M. M. Chun, Y. Jiang, Contextual Cueing: Implicit Learning and Memory of Visual Context Guides Spatial Attention. *Cognit. Psychol.* **36**, 28–71 (1998).

8. S. Jaramillo, A. Zador, Auditory cortex mediates the perceptual effects of acoustic temporal expectation. *Nat. Preced.* (2010) https:/doi.org/10.1038/npre.2010.5139.1 (April 28, 2023).

9. L. M. Kroell, M. Rolfs, Foveal vision anticipates defining features of eye movement targets. *eLife* **11**, e78106 (2022).

10. A. Herwig, W. X. Schneider, Predicting object features across saccades: Evidence from object recognition and visual search. *J. Exp. Psychol. Gen.* **143**, 1903–1922 (2014).

11. C. Osterbrink, A. Herwig, Prediction of complex stimuli across saccades. *J. Vis.* **21**, 10 (2021).

12. J.-R. Duhamel, C. L. Colby, M. E. Goldberg, The Updating of the Representation of Visual Space in Parietal Cortex by Intended Eye Movements. *Science* **255**, 90–92 (1992).

13. M. M. Umeno, M. E. Goldberg, Spatial Processing in the Monkey Frontal Eye Field. I. Predictive Visual Responses. *J. Neurophysiol.* **78**, 1373–1383 (1997).

14. M. F. Walker, E. J. Fitzgibbon, M. E. Goldberg, Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J. Neurophysiol.* **73**, 1988–2003 (1995).

15. M. J. Chadwick, S. L. Mullally, E. A. Maguire, The hippocampus extrapolates beyond the view in scenes: An fMRI study of boundary extension. *Cortex* **49**, 2067–2079 (2013).

16. H. Intraub, M. Richardson, Wide-Angle Memories of Close-Up Scenes. 9 (1989).

17. S. L. Mullally, H. Intraub, E. A. Maguire, Attenuated Boundary Extension Produces a Paradoxical Memory Advantage in Amnesic Patients. *Curr. Biol.* **22**, 261–268 (2012).

18. S. Park, H. Intraub, D.-J. Yi, D. Widders, M. M. Chun, Beyond the Edges of a View: Boundary Extension in Human Scene-Selective Visual Cortex. *Neuron* **54**, 335–342 (2007).

19. W. A. Bainbridge, C. I. Baker, Boundaries Extend and Contract in Scene Memory Depending on Image Properties. *Curr. Biol.* **30**, 537-543.e3 (2020).

20. M. Gandolfo, H. Nägele, M. V. Peelen, Predictive Processing of Scene Layout Depends on Naturalistic Depth of Field. *Psychol. Sci.* (2023).

21. A. Hafri, S. Wadhwa, M. F. Bonner, Perceived Distance Alters Memory for Scene Boundaries. *Psychol. Sci.* (2022).

22. S. C. Berens, B. H. Joensen, A. J. Horner, Tracking the Emergence of Location-based Spatial Representations in Human Scene-Selective Cortex. *J. Cogn. Neurosci.* **33**, 445–462 (2021).

23. C. E. Robertson, K. L. Hermann, A. Mynick, D. J. Kravitz, N. Kanwisher, Neural Representations Integrate the Current Field of View with the Remembered 360° Panorama in Scene-Selective Cortex. *Curr. Biol.* **26**, 2463–2468 (2016).

24. M. S. Castelhano, A. Pollatsek, Extrapolating spatial layout in scene representations. *Mem. Cognit.* **38**, 1018–1025 (2010).

25. T. Sanocki, Representation and perception of scenic layout. *Cognit. Psychol.* **47**, 43–86 (2003).

26. T. Sanocki, W. Epstein, Priming spatial layout of scenes. *Psychol. Sci.*, 374–378 (1997).

27. A. Shafer-Skelton, T. F. Brady, Scene layout priming relies primarily on low-level features rather than scene layout. *J. Vis.* **19**, 14 (2019).

28. W. A. Bainbridge, E. H. Hall, C. I. Baker, Drawings of real-world scenes during free recall reveal detailed object and spatial information in memory. *Nat. Commun.* **10**, 5 (2019).

29. W. A. Bainbridge, W. Y. Kwok, C. I. Baker, Disrupted object-scene semantics boost scene recall but diminish object recall in drawings from memory. *Mem. Cognit.* **49**, 1568–1582 (2021).

30. E. H. Hall, W. A. Bainbridge, C. I. Baker, Highly similar and competing visual scenes lead to diminished object but not spatial detail in memory drawings. *Memory* **30**, 279–292 (2022).

31. A. Hollingworth, Scene and position specificity in visual memory for objects. *J. Exp. Psychol. Learn. Mem. Cogn.* **32**, 58–69 (2006).

32. J. S. Matthis, J. L. Yates, M. M. Hayhoe, Gaze and the Control of Foot Placement When Walking in Natural Terrain. *Curr. Biol.* **28**, 1224-1233.e5 (2018).

33. A. Steel, K. Goyal, A. Mynick, C. E. Robertson, Convergence zones for scene perception and visuospatial memory at the anterior edge of visually-responsive cortex. *Rev.* (2022).

34. T. L. Botch, B. D. Garcia, Y. B. Choi, N. Feffer, C. E. Robertson, Active visual search in naturalistic environments reflects individual differences in classic visual search performance. *Sci. Rep.* **13**, 631 (2023).

35. J. Chen, *et al.*, Shared memories reveal shared structure in neural activity across individuals. *Nat. Neurosci.* **20**, 115–125 (2017).

36. J. L. S. Bellmund, *et al.*, Deforming the metric of cognitive maps distorts memory. *Nat. Hum. Behav.* (2019) https:/doi.org/10.1038/s41562-019-0767-3 (November 22, 2019).

37. I. Moar, G. H. Bower, Inconsistency in spatial knowledge. *Mem. Cognit.* **11**, 107–113 (1983).

38. B. Tversky, Distortions in memory for maps. *Cognit. Psychol.* **13**, 407–433 (1981).

39. Y. Zhao, A. J. H. Chanales, B. A. Kuhl, Adaptive Memory Distortions Are Predicted by Feature Representations in Parietal Cortex. *J. Neurosci.* **41**, 3014–3024 (2021).

40. D. Draschkow, A. C. Nobre, F. van Ede, Multiple spatial frames for immersive working memory. *Nat. Hum. Behav.* (2022) https:/doi.org/10.1038/s41562-021-01245-y (March 20, 2022).

41. F. Aitken, P. Kok, Hippocampal representations switch from errors to predictions during acquisition of predictive associations. *Nat. Commun.* **13**, 3294 (2022).

42. A. Clarke, J. Crivelli-Decker, C. Ranganath, Contextual expectations shape cortical reinstatement of sensory representations. *J. Neurosci.*, JN-RM-2045-21 (2022).

43. J. Crivelli-Decker, *et al.*, Goal-centered representations in the human hippocampus. *bioRxiv* (2021).

44. J. L. Thom, A. C. Nobre, F. Van Ede, D. Draschkow, Heading Direction Tracks Internally Directed Selective Attention in Visual Working Memory. *J. Cogn. Neurosci.* **35**, 856–868 (2023).

45. F. Van Ede, A. C. Nobre, Turning Attention Inside Out: How Working Memory Serves Behavior. *Annu. Rev. Psychol.* **74**, 137–165 (2023).

46. M. A. Cohen, T. L. Botch, C. E. Robertson, The limits of color awareness during active, real-world vision. *Proc. Natl. Acad. Sci.* **117**, 13821–13827 (2020).

47. A. J. Haskins, J. Mentch, T. L. Botch, C. E. Robertson, Active vision in immersive, 360° real-world environments. *Sci. Rep.* **10**, 14304 (2020).

Figures and Tables

A Study takes place entirely in immersive head-mounted virtual reality

B Training Phase (Day 1 and 2): Participants explore scenes with no task

occluder
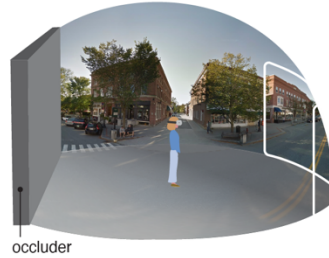
C Each scene has an open and closed view on opposite sides of the initial view

open   initial   closed

-180°   -90°   0°   90°   180°

D Trial timecourse for Priming Test (Day 2)

prime (300ms)

head turn

target (until response)

Time

"Open or closed?"
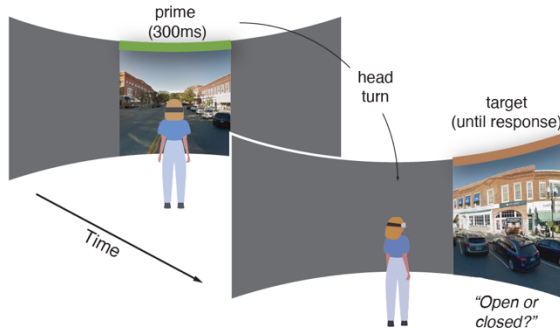
E Priming conditions (Experiments 1 and 2)

Same-scene

prime   target

Neutral

Different-scene

Figure 1. Experimental paradigm. A. Participants were fitted with virtual reality (VR) headsets which were used throughout study. B. On both days of the Training Phase, participants studied each real-world scene in immersive VR for 20 seconds, instructed to simply "Look around like you normally would". C. Each scene had one open and one closed view on opposing sides of the initial facing-direction in the scene, spaced 180° apart on the horizontal axis. D. On Day 2, participants performed a Priming Test in which they viewed a 300ms prime image before turning left or right toward a target image (open or closed view) to make a speeded perceptual (open/closed) judgment. Targets were presented until a response was made. E. Prime images in Experiments 1-2 either depicted a continuous view adjacent to the target from the initial direction (Same-scene), a blank gray rectangle (Neutral) or an image taken from the initial direction of a different scene (Different-scene).
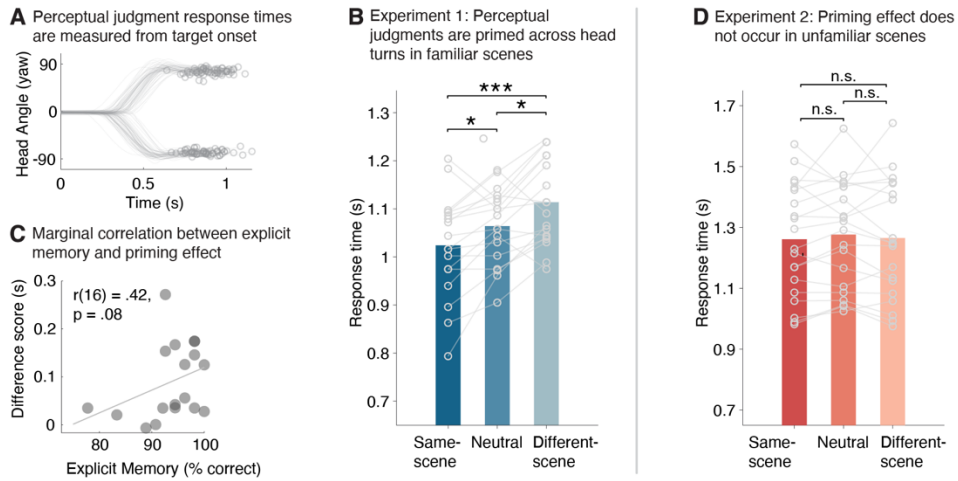
Figure 2. Same-scene primes facilitate faster perceptual judgments in familiar (but not unfamiliar) scenes. A. Example participant's data depicting head angle on the horizontal plane across a trial. Each line depicts head angle on the horizontal plane (yaw) for one trial from target onset. Each open circle marks the (time, head angle) of the participant at the time of the perceptual judgment. B. We compared median response times for perceptual judgments across priming conditions in Experiment 1 (N=18). Compared to Neutral primes, Same-scene primes quickened response times and Different-scene primes slowed response times. Connected points represent the same participant across conditions. C. There was a marginal correlation between extent of priming (difference between median response time for Different-scene and Same-scene primes) and accuracy in the Explicit Memory Test. Each point represents one participant. D. There was no difference in median response times for perceptual judgments across priming conditions in Experiment 2 (N=20), which contained unfamiliar scenes. Connected points represent the same participant across conditions. * p <= .05, ** p <=.01, *** p <= .001 difference between median response time for indicated priming conditions (Bonferroni-corrected paired t-tests).
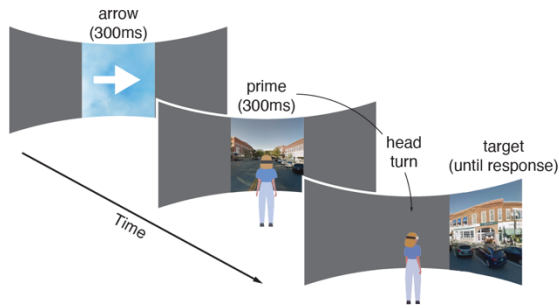
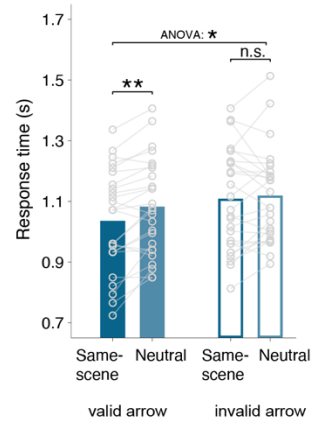**Figure 3.** Priming is biased in the direction of intended head turns. A. On each trial of Experiment 3, participants saw an arrow indicating a direction to plan a head turn in, followed by a prime and target. Arrows were either valid (i.e. correctly indicated the location of the upcoming target) or invalid. B. Median response times for perceptual judgments across priming conditions in Experiment 3 (N=26) were only faster in Same-scene- compared to Neutral primes when preceded by a valid arrow. Connected points represent the same participant across conditions. * p <= .05, ** p <=.01, *** p <= .001 difference between median response time for indicated priming conditions (Bonferroni-corrected paired t-tests unless otherwise noted).
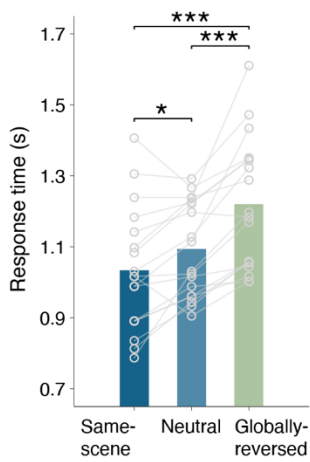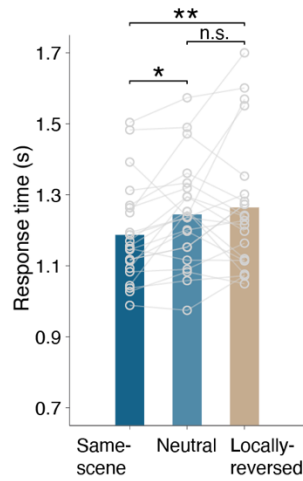
A Experiment 4: Globally-reversed primes
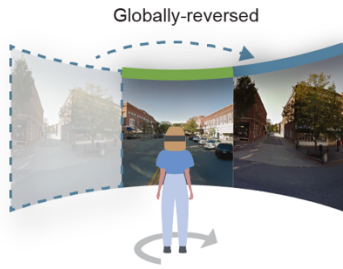
Globally-reversed

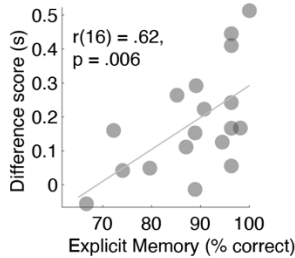**D** Experiment 5: Locally-reversed primes

Locally-reversed

B Priming obeys the global spatial structure of a scene

E Priming obeys the local spatial structure of a scene

C Correlation between explicit memory and priming effect

$r(16) = .62,$
$p = .006$

F No correlation between explicit memory and priming effect
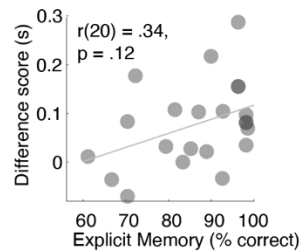
$r(20) = .34,$
$p = .12$

Figure 4. Priming obeys the global and local spatial structure of scenes. A. In Globally-reversed trials, targets were presented opposite their learned location in 360º space. B. Median response times for perceptual judgments in Experiment 4 (N=18) were quickened by Same-scene primes compared to Neutral while Globally-reversed primes slowed median response times compared to Neutral. Connected points represent the same participant across conditions. C. There was a correlation between extent of priming (difference between median response time for Globally-reversed and Same-scene primes) and accuracy in the Explicit Memory Test in Experiment 4. Each point represents one participant. D. Locally-reversed trials contained prime and target images from the same scene, but presented the target image – in its original location – mirror-reversed across the vertical axis. E. Median response times for perceptual judgments in Experiment 5 (N=22) were

quickened by Same-scene primes compared to Neutral while response times for Locally-reversed primes did not differ from Neutral. Connected points represent the same participant across conditions. F. There was no significant correlation between extent of priming (difference between median response time for Locally-reversed and Same-scene primes) and accuracy in the Explicit Memory Test in Experiment 5. Each point represents one participant. * p <= .05, ** p <=.01, *** p <= .001 difference between median response time for indicated priming conditions (Bonferroni-corrected paired t-tests).

Methods Table 1. Overview of Experiments 1-5.

	# Days	Scenes	Priming Test Conditions	Condition split (%)	# Priming Test Trials	Neutral Type	Explicit Memory Task	# Participants collected	# Participants included
Experiment 1	2	Dartmouth	Same-scene	50	144	Blank grey	yes	26	18
			Neutral	25					
			Different-scene	25					
Experiment 2	1	Unfamiliar locations	Same-scene	50	144	Blank grey	no	24	20
			Neutral	25					
			Different-scene	25					
Experiment 3	2	Dartmouth	Same-scene valid arrow	50	216	Scrambled	no	46	26
			Neutral valid arrow	16.7					
			Same-scene invalid arrow	16.7					
			Neutral invalid arrow	16.7					
Experiment 4	2	Dartmouth	Same-scene	50	144	Blank grey	yes	23	18
			Neutral	25					
			Globally-reversed	25					
Experiment 5	2	Dartmouth	Same-scene	50	144	Scrambled	yes	31	22
			Neutral	25					
			Locally-reversed	25					