Perceived distance alters memory for scene boundaries

Alon Hafri^{a,b,*}, Shreya Wadhwa^a, & Michael F. Bonner^{a,*}

^aDepartment of Cognitive Science, Johns Hopkins University, Baltimore, MD 21218

^bDepartment of Psychological & Brain Sciences, Johns Hopkins University, Baltimore, MD 21218

Abstract

Memory often fills in what's not there. A striking example of this is boundary extension, whereby observers mistakenly recall a view that extends beyond what was seen. However, not all visual memories extend in this way—suggesting that this process depends on specific scene properties. What factors determine when visual memories will include details that go beyond perceptual experience? Here, seven experiments (n=1,100 adults) explored whether spatial scale—specifically, perceived viewing distance—drives boundary extension. We created “fake miniatures” by exploiting tilt-shift, a photographic effect which selectively reduces perceived distance while preserving other scene properties (e.g., making a distant railway appear like a model-train). Fake-miniaturization increased boundary extension for otherwise identical scenes: Participants performing a scene-memory task misremembered fake-miniaturized views as farther than they actually were. This effect went beyond low-level image changes and generalized to a completely different distance manipulation. Thus, visual memory is modulated by the spatial scale at which the environment is viewed.

Keywords: scene perception, boundary extension, boundary transformation, tilt-shift, memory distortion, visual memory, spatial memory

* To whom correspondence should be addressed.

Email addresses: alon@jhu.edu (Alon Hafri), mfbonner@jhu.edu (Michael F. Bonner)

Introduction

Remembering is a constructive process that is prone to systematic distortions. These distortions can reveal the mechanisms by which the visual environment is encoded in the mind. A foundational example of such memory distortion is *boundary extension*, whereby memories of visual scenes include information beyond the boundaries actually observed (Intraub & Richardson, 1989; Fig. 1a). This distortion is striking because it suggests that the mind adds information to visual memories that was not there in the first place. Moreover, it is empirically robust: boundary extension is observed whether participants are probed after short or long delay periods (Intraub et al., 1992), manifests in both image-recognition and drawing-reproduction tasks (Bainbridge & Baker, 2020a; Gottesman, 2011; Intraub et al., 1998), emerges early in infancy (Quinn & Intraub, 2007), and even occurs in non-visual modalities (Mullally et al., 2012).

The existence of boundary extension was originally taken to reflect a “filling-in” process, in which perceptual information is integrated with a mental schema about the likely contents of the environment beyond the immediate view (Intraub & Richardson, 1989; Intraub, 2010). However, it turns out that memory for image boundaries does not always systematically extend; instead, for many images, boundaries in memory do not extend at all, and may even *contract*, suggesting that boundaries may transform in both directions (Bainbridge & Baker, 2020a; Fig. 1a). This recent discovery has reinvigorated interest in the possible mechanisms and cognitive functions of boundary transformation (Intraub, 2020; Bainbridge & Baker, 2020b; Park et al., 2021), by revealing that boundary extension is not a universal phenomenon of image memory. Instead, the magnitude and directionality of memory distortions appear to be *scene-specific*: for example, an image of a seashell may consistently elicit boundary extension, while an image of a remote island may consistently elicit boundary contraction.

The scene-specific nature of these effects raises a question: what scene properties cause boundary transformation in the first place? The literature has speculated that *viewing distance* may play a prominent role in boundary transformation, whereby increasingly close views should result in more boundary extension (or less boundary contraction), and vice-versa for increasingly distant views (Bainbridge & Baker, 2020a; Bertamini et al., 2005; Intraub & Richardson, 1989). However, the degree to which viewing distance plays a causal role in boundary transformation remains unknown.

The Challenge: Isolating Scene Properties to Understand Visual Memory

How might one determine the independent contribution of scene properties like viewing distance to boundary transformation? At first glance, it would seem impossible, as viewing distance generally covaries with myriad other perceptual and semantic properties. For example, the images in Figure 1b differ in viewing distance, but so much more besides: the presence of certain objects (e.g., water vs. trees), spatial frequencies (low vs. high), navigational affordances (open vs. closed), and other factors. Even in careful attempts to parametrically vary distance in tightly controlled scenes, the problem remains: it is impossible to move the camera without changing something else, such as the retinal size of objects (Bertamini et al., 2005), or the presence of walls or furniture (Park et al., 2021). Nevertheless, discovering the underlying causes of boundary-transformation phenomena remains a crucial prerequisite for revealing the mechanisms and potential functional role of these memory distortions in cognition.

The importance of determining the causal factors of boundary transformation goes beyond merely understanding more about this specific phenomenon; it is critical for developing a broader theory of the relationship between perception and memory as a whole. In particular, one might expect that it is primarily the *contents* of an image that determine the manner in which perception and memory are integrated. Images vary in the objects, people, and settings in which they occur—e.g., when we observe an island in the distance, some friends strolling along the beach, or a seashell that has washed ashore—all of which might be expected to influence memory, based on prior knowledge about such contents (Bartlett, 1932; Hemmer & Steyvers, 2009). By contrast, if boundary transformation can be modulated by viewing distance alone, this suggests that visual memories may be deeply influenced by the general spatial contexts in which they are formed (here, the spatial scale of the environment).

Fig. 1. (a) Some scene images elicit *boundary extension*, in which observers' memory for the boundaries of a visual scene are more wide-angle than they actually were (as if they had viewed the scene from a farther vantage point than in actuality). Other scenes may elicit *boundary contraction*, in which memory for boundaries is more restricted than was actually observed (as if they viewed the scene from a closer vantage point). (b) Determining whether scene properties such as viewing distance cause boundary transformation has proved challenging, as many perceptual, spatial, and semantic properties of scenes covary.

The Present Experiments: “Shrinking” Visual Scenes

In the current study we confront the challenge of isolating scene properties head-on, and in doing so, provide the first direct test of the causal role of perceived distance in boundary-transformation phenomena. To accomplish this, we created “fake miniatures” from images of distant scenes by leveraging a photographic

technique called *tilt-shift*, which mimics the shallow depth-of-field inherent in optics at close-range (Held et al., 2010; Fig. 2a)—an effect which we simulated digitally (Fig. 2b). Crucially, this image manipulation selectively reduces perceived distance by altering the scale of the scene while preserving its semantic content, spatial structure, and most other perceptual properties. Indeed, it is phenomenologically quite striking: in Figure 2b, the distant railway bridge convincingly appears as a close-up model-train set. We embedded images of distant scenes and their “close-up” counterparts in a scene-memory task to ask the question: Does altering perceived viewing distance modulate memory for image boundaries? If memory distortions are dependent on how far away a scene *appears*—regardless of its other content—then this simple manipulation (fake-miniaturization) should increase boundary extension for an otherwise identical scene.

Our first experiment explored boundary extension for fake miniatures in the way just described. Several follow-up experiments showed that the observed boundary-extension effects went beyond low-level changes introduced by the image manipulation. A final set of studies asked whether the effects generalized to a completely different distance manipulation. To anticipate, our experiments provide evidence that perceived viewing distance plays a causal role in driving memory distortions for scene boundaries. Readers can experience the tasks for themselves at <https://cognitivestudies.online/shrunkenscenes>.

Experiment 1: “Fake-Miniaturization”

Does perceiving a scene as close-up cause boundary extension? To address this question, in Experiment 1, we embedded distant scenes and their close-up (fake-miniaturized) counterparts in a scene-memory task, and asked participants to report whether a probe image was a closer or farther view than their memory for a target image.

Fig. 2. We created “fake miniatures” from images of distant scenes by applying a technique called *tilt-shift*. (a) Tilt-shift simulates the shallow depth-of-field that is inherent in optics at close-range. When focusing on scene elements that are far away, elements within a wide range of distances may be in focus at once; in

contrast, when focusing on things close-up, only a narrow range of scene elements may be in focus. (b) We simulated tilt-shift digitally by applying a vertical blur gradient and saturation to distant scene images (in blue), which creates fake miniatures (in orange). This manipulation selectively alters perceived viewing distance while preserving most other image properties, allowing us to test the causal role of perceived viewing distance on boundary transformation in memory.

Methods

Open Science Practices

All data, code, analyses, stimuli, and pre-registrations (for this experiment and all others reported here) are available at <https://cognitivestudies.online/shrunkenscenes>. This webpage also includes demos of each experiment, so that readers can experience these tasks as participants did. The sample sizes and analysis plans (as well as other details) for all experiments were pre-registered.

Participants

100 adult participants were recruited from the online platform Prolific (<https://prolific.co>). (For a discussion of the reliability of this subject pool, see Peer et al., 2017). This was chosen based on the sample size of a recent study with a similar population and experimental paradigm (Bainbridge & Baker, 2020a). Sample sizes were pre-registered for this and all other experiments. All studies were approved by the Johns Hopkins University Institutional Review Board.

Stimuli

We collected 32 natural scene images from the internet that varied in setting (e.g., construction site, farmhouse, airport). These were distant, high-angle views of the environment, which are the kind of images that are ideal for producing fake miniatures. To achieve the fake-miniaturization effect, we digitally simulated a photographic technique called “tilt-shift”. In tilt-shift, the orientation and position of the lens is changed such that a vertically oriented blur gradient is introduced (i.e., blurred regions appear at the top and bottom of the image). This mimics the shallow depth-of-field (DoF) that results from the optics that naturally arise when viewing a close-up scene (Held et al., 2010; Watt et al., 2005; Vishwanath & Blaser, 2010; Figure 2a). To simulate tilt-shift digitally, we applied a vertically oriented blur gradient to the image that makes the top and bottom of the image blurry while keeping the central part sharp (Fig. 2b). We also increased the saturation and contrast of the image, which mimics the bright colors of a painted diorama or miniature model. (Although these latter additions are generally not considered to be necessary for the change in apparent scale, they enhance the impression of a miniature.)

In a few cases, we were able to obtain ready-made fake-miniature versions of scenes. For most images, we used a publicly available tool for this purpose, <https://tiltshiftmaker.com>. Since the effectiveness of the miniaturization depends on blurring the image in a way consistent with the changing distances of the image content, the vertical location and extent of the blur gradient were specified manually for each image. These location and width values are available on the project’s OSF repository. Example image pairs can be viewed in Figure 4a. We also created mask images, which were 32 extra fake-miniaturized images, each box-scrambled in a 13x20 grid.

All stimuli were displayed at 800x520 pixels in the participant’s Web browser. Due to the nature of online studies, we cannot know the exact viewing distance, screen size, luminance (etc.) of these stimuli as they

appeared to participants. However, any distortions introduced by a given participant's viewing distance or monitor settings would have been equated across all stimuli and conditions.

Design and Procedure

To test the hypothesis that perceived viewing distance drives memory distortions for scene boundaries, participants performed a rapid scene-recognition task commonly used to assess boundary-extension effects (Bainbridge & Baker, 2020a; Intraub & Dickinson, 2008; Mullally et al., 2012; depicted in Fig. 3a). Participants were instructed that on each trial, they would view a briefly displayed target image and would have to decide whether a probe image appearing a moment later was “zoomed-in” or “zoomed-out” relative to the target. To initiate a trial, participants pressed the spacebar. The trial sequence was as follows: a fixation cross (250ms), the target scene image (250ms), a dynamic mask (5 random mask images each for 50ms), a fixation cross (350ms), and finally the probe image. The probe image remained on screen until the participant responded. Participants pressed one key for closer and another for farther (v and m keys, with the mapping randomized across participants). The target and probe images were always exactly the same in all respects, although crucially, participants were not informed of this. A post-experiment questionnaire confirmed that the majority of participants were not aware that the zoom did not vary between the target and probe images.

Participants viewed each scene identity exactly once, in one of two Image Type conditions: Original or Fake-Miniature. Each participant was assigned to one of two stimulus lists (counterbalanced across participants). The two lists differed in which condition each scene identity was assigned to. Image order was randomized for each participant.

In addition to test trials, four catch trials were randomly interspersed among the test trials, for two reasons: to ensure focus and to exclude participants who failed to adequately engage in the task. On catch trials, instead of being asked about image zoom, participants were asked to respond whether the probe image was exactly the same as, or completely different from, the target image. Half the catch trials required a positive response (same exact image) and half a negative response (completely different images, e.g., a forest and a city). Before the main experiment started, participants were also given five practice trials (one as a catch trial), which included feedback. These were the same trial structure as the main trials, but included real (and obvious) changes in zoom levels between the target and probe images, to convince participants that trials in the main study would likewise have different zooms (even if more subtle).

Logic and Predictions

The experimental logic of our design is depicted in Figure 3b. As described above, on test trials, the target and probe images were always exactly the same image. Thus, if no boundary transformation occurs on the target image, participants should respond “closer” about as often as responding “farther”. However, if boundary extension occurs, then the probe (which remained on screen) should appear *closer* than the participants' memory for the target, since this memory includes visual information beyond the image boundaries; thus, the participant would respond “closer” more often than “farther”. In contrast, if boundary contraction occurs, the probe should appear *farther* than the participants' memory for the target, since this memory *excludes* visual information at the edge of the image; thus, the participant would respond “farther” more often than “closer”. Our key prediction was that decreasing the perceived viewing distance of scenes (via fake-miniaturization) should increase boundary extension for those very same scenes, relative to the original, non-manipulated versions of those scenes. In our paradigm, this would manifest as a significant increase in how often participants reported the probe image as being a closer view than the target image.

Fig. 3. (a) Scene-memory task: On each trial, participants viewed a target image, which was dynamically masked, and then they were presented with a probe image. Participants were asked to respond whether the probe appeared to be a closer or farther view than the target. The probe image was always exactly the same as the target, but participants were not informed of this. (b) Logic and predictions: Fake-miniaturized images were predicted to elicit boundary extension, in which case the probe would appear closer than the memory for the target, so participants would respond “Closer”. (If no boundary transformation were to occur, participants would be expected to choose both responses at similar rates.)

Exclusions

In accordance with our pre-registered analysis plan, we excluded participants for the following reasons: if they did not contribute a complete dataset, if they had low accuracy on catch trials (<75% correct), if they exhibited a large number of unreasonably fast reaction times (>10% of RTs <200ms), or if they gave the same response on more than 95% of trials. Practice and catch trials were excluded from analyses. We also planned to exclude any remaining trials with RTs less than 200ms, although no trials met this criterion. Seven participants were excluded, although none of the results reported here or in subsequent experiments were dependent on the exclusion criteria.

Analyses

We tested our predictions with mixed-effects logistic regression on trial-level data. Mixed-effects models allow for generalization of statistical inferences simultaneously across participants and items (Barr et al., 2013). The dependent variable was “closer” responses. The primary independent variable was Image Type (Original or Fake-Miniature, sum-coded as -0.5 and 0.5 , respectively), which we predicted would show significance as a main effect. A main effect of Trial Number (centered) was always included in the baseline model to account for general order effects. An interaction of Image Type and Trial Number was also tested, in case the effect of Image Type changed over the course of the study.

We tested for significance of variables by using likelihood ratio tests on the Chi-square values from nested model comparisons with the same random effects structure. We started with the maximal random effects structure: correlated random intercepts and random slopes for Image Type by participant and by item (scene identity). If models did not converge, we simplified the random effects structure by first using uncorrelated intercepts and slopes, and followed that by dropping random intercepts and slopes until convergence, starting with those that accounted for the least variance. Random effects structures for all mixed-effects models reported in this paper are available in this project's OSF repository.

For all experiments, marginal and conditional pseudo- R^2 statistics are reported as effect sizes for the best-fitting model; these denote the proportion of variance explained by the fixed effects only (marginal, or R_m^2) and by the full model with both fixed and random effects (conditional, or R_c^2) (Nakagawa et al., 2017). For experiments using mixed-effects logistic regression, we also report effect sizes in terms of odds ratios (ORs), derived from β (the logit-transformed fixed effect coefficient). For the fixed effect of Category, the OR represents the increased likelihood of a “closer” response for one image category vs. another. For example, an OR of 2.0 would mean that a “closer” response is two times more likely for, e.g., the Fake-Miniaturized vs. Original image condition.

Mixed-effect modeling and statistical analyses for all experiments proceeded as described above, unless noted otherwise.

Fig. 4. Image manipulation and results for Experiment 1 (scene-memory task). (a) Examples of distant scenes (in blue) that were made into fake-miniaturized scenes (in orange). (b) There were significantly more "closer" responses for close-up images (fake-miniaturized, in orange) relative to their far-away counterparts (original images, in blue), which is evidence for more boundary extension for fake-miniaturized images. The two colored circles are means across items (scene identities), error bars are 95% within-item confidence intervals, and light-gray points and connecting lines represent data for individual items in each condition. *** $p < .001$.

Results

Results can be seen in Figure 4b. As predicted, participants reported that the probe image appeared closer than the target more often for the Fake-Miniature condition ($M=69.3\%$) as compared to the Original ($M=51.7\%$). The robustness of this effect was also evident non-parametrically: 73 of 93 participants (78%) and 30 of 32 scene identities (94%) went in the direction of this effect. As a reminder, probe and target images were exactly identical. Thus, reporting more often that the probe image is closer (more “zoomed-in”) than the target image is evidence that boundary extension occurred on the memory of the target image. These results indicate that making a distant scene appear as a close-up “miniature” causes memory for its boundaries to be more extended than it would otherwise have been. In other words, a simple image manipulation, fake-miniaturization (“tilt-shift”), was sufficient to dramatically increase boundary-extension effects for that very same scene.

These results were confirmed in mixed-effects model comparisons. The best-fitting model was one that included a main effect of Image Type ($\beta = 0.84$ [95% CI 0.63 to 1.05], $z = 7.86$, $p < .001$, $OR = 2.32$ [95% CI 1.88 to 2.86]). This model was a better fit than one that did not include this factor, $\chi^2(1) = 38.44$, $p < .001$, $R_m^2 = 0.05$, $R^2 = 0.17$. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(1) = 1.02$, $p = .312$), suggesting that the effect of Image Type was stable across trials.

Experiments 2a and 2b: Mere Reduction in Image Details?

Although fake-miniaturization preserves almost all perceptual, spatial, and semantic properties of images, there are necessarily *some* perceptual differences introduced by the manipulation. Notice that fake-miniaturization reduces fine-grained detail in the blurred regions, and may even decrease the number of resolvable objects (effectively turning a scene with many objects into one with fewer). This could also have downstream effects on the salience of certain scene regions and on the distribution of attention across the scenes. Might a reduction in “image resolvability” lead to the observed increase in boundary extension?

We addressed this possibility by reducing resolvability to the extreme: we occluded each fake-miniaturized scene where its blurred regions were located (Fig. 5a). If image resolvability were the cause of increased boundary extension (rather than changes in perceived distance), then fully eliminating resolvability via image occlusion should greatly *increase* boundary extension. We predicted instead that resolvability would “break” the change in perceived distance, and thus would diminish the boundary-extension effects of fake-miniaturization. We tested these predictions in Experiment 2a (distance ratings) and 2b (scene-memory task).

Methods

Participants

Two groups of participants were recruited from Prolific. We chose a sample size of 500 for Experiment 2a based on power analyses of pilot data suggesting that we would need this sample size to detect distance-rating differences.¹ We chose a sample size of 100 for Experiment 2b to match the sample size of Experiment 1 (the scene-memory task).

¹ A pilot version of Experiment 2a suggested that 500 participants would achieve 79.8% power to detect a distance-rating difference between the fake-miniaturization and occlusion control conditions; we thus pre-registered this sample size for Experiment 2a. As is clear from the results, however, the effect in Experiment 2a was remarkably robust, such that 500

Stimuli

The 32 scene images from Experiment 1 were used here. To create the “occlusion control” condition, we added grayscale occluders to the fake-miniaturized images in the region(s) where the blur gradient was, spanning the whole horizontal range of the images. Occlusions matched the mean luminance of each image. Examples can be seen in Figure 5a. Masks were the same as in Experiment 1. All stimuli were displayed at 654x425 pixels in the participant’s Web browser for Experiment 2a and 800x520 pixels for Experiment 2b.

Design and Procedure

In Experiment 2a, participants rated how far away the main object or thing in each image appeared. We designed the stimulus-presentation procedure to match the viewing conditions of the scene-memory task used in Experiment 1 (and later in Experiments 2b, 3b, and 4b): the target scene image was displayed for 250ms and then dynamically masked. This was followed by a radio-button rating scale. The scale was based on one used in several previous studies (Bainbridge & Baker, 2020a; Lescroart et al., 2015). Participants could choose among five distance options: (1) far away, > 100 ft (30 m); (2) short walk/some effort, < 100 ft (30 m); (3) nearby/same room, < 20 ft (6 m); (4) arm’s length, ~3–4 ft (0.9–1.2 m); and (5) extreme closeup, ~0–2 ft (0–0.6 m). The left-right order of these choices was randomly assigned across participants. In addition to distance-rating trials, four catch trials were randomly interspersed among the test trials, to ensure focus and to exclude participants for failing to engage with the task. On such trials, instead of a scene image being displayed, an image with the text “Catch Trial!” was displayed and masked. Participants were instructed to watch for these no-scene trials, and to select a sixth “catch” radio button instead of one of the five distance options.

In Experiment 2b, the experimental procedure was identical to Experiment 1: Participants reported whether a probe image depicted a closer or farther view than what they remembered (although in actuality the probe image was always exactly the same as the target image).

In both Experiments 2a and 2b, participants viewed each scene identity exactly once, in just one of the three Image Type conditions: Original, Fake-Miniature, or Occlusion Control. Each participant was assigned to one of six stimulus lists (counterbalanced across participants). Each list differed in which condition each scene identity was assigned to.

Exclusions

In accordance with our pre-registered analysis plan, we excluded participants from Experiment 2a for the following reasons: if they did not contribute a complete dataset, if they had low accuracy on catch trials (<75% correct), if they had too many “catch” responses (false-alarms) on non-catch trials (>10%), if they exhibited a large number of unreasonably fast reaction times (>10% of RTs <400ms), or if they gave the same response on more than 95% of trials. Individual trials were also excluded if the RT was less than 400ms or if the trial was a “catch” response. Twenty-nine participants and 97 remaining trials were excluded from Experiment 2a. Exclusion criteria for Experiment 2b were the same as in Experiment 1; seven participants and two remaining trials were excluded.

participants turned out to be far more than was required to detect the relevant effect (with $n = 500$, observed power was 99.9%; and for $n = 100$, observed power was 97.8%). Thus, for the remaining experiments, we pre-registered sample sizes of 100, to match Experiment 1’s sample size.

Analysis

For Experiment 2a, we tested whether occluding the blur gradient in the image would break the effect of the fake-miniaturization on perceived distance by running linear mixed-effects regression on trial-level data, predicting distance ratings (higher ratings closer, lower ratings farther). For Experiment 2b, we tested whether occlusion would diminish boundary-extension effects by running mixed-effects logistic regression, predicting “closer” responses (as in Experiment 1). For each experiment, the primary independent variable was Image Type (Original, Fake-Miniature, or Occlusion Control, treatment-coded with Original as the baseline condition). In addition to testing for differences among all three Image Types, we also compared them pairwise (i.e., we separately compared Original vs. Fake-Miniature, Original vs. Occlusion Control, and Occlusion Control vs. Fake-Miniature, all sum-coded as -0.5 and 0.5 , respectively, within each pair).

Results

Results can be seen in Figure 5. Inspection of this figure suggests several findings. First, as predicted, distance ratings in Experiment 2a were dependent on Image Type: Fake-Miniature images were rated as closer ($M=2.02$) than both the Original images ($M=1.86$) and Occlusion Control ($M=1.84$), while the latter two image types were not different from one another (Fig. 5b). These distance-rating results confirm that the Fake-Miniature manipulation indeed had the intended effect of reducing perceived distance, and that the Occlusion Control condition “broke” this perceived-distance effect.² Crucially, this same occlusion manipulation appeared to reduce boundary-extension effects: In Experiment 2b, participants reported that the probe image appeared closer than the target more often for Fake-Miniature images ($M=71.0\%$) than for both the Original ($M=56.1\%$) and Occlusion Control ($M=56.5\%$) images, which did not differ from one another (Fig. 5c). These results also replicate the results from Experiment 1, in which Fake-Miniature images showed greater boundary-extension effects than the Original images.

These observations were confirmed statistically. For Experiment 2a (distance ratings), a mixed-effect model with the main effect of Image Type was a significantly better fit than a model without it, $\chi^2(2) = 30.54, p < .001$, $R_m^2 = 0.02$, $R_e^2 = 0.42$. Although adding an interaction of Image Type and Trial Number did improve model fit ($\chi^2(2) = 20.01, p < .001$) (suggesting that the effect of Image Type changed across trials), the main effect of Image Type was still significant in both the first quarter and last quarter of trials (first: $\chi^2(2) = 30.75, p < .001$; last: $\chi^2(2) = 9.81, p = .007$). Follow-up analyses comparing image types pairwise showed a significant main effect of Image Type for both the Fake-Miniature vs. Original comparison ($\chi^2(1) = 26.58, p < .001$, $R_m^2 = 0.02$, $R_e^2 = 0.43$; $\beta = 0.16$ [95% CI 0.11 to 0.21], $t = 5.91$), and the Fake-Miniature vs. Occlusion Control comparison ($\chi^2(1) = 20.04, p < .001$, $R_m^2 = 0.02$, $R_e^2 = 0.40$; $\beta = 0.19$ [95% CI 0.12 to 0.26], $t = 5.13$). By contrast, for the

² The effect of fake-miniaturization on perceived distance may at first seem subtle based on these mean distance ratings. However, there are various reasons that the raw magnitude of this difference should be treated with caution. First, naive observers may be quite imprecise in evaluating metric distances (vs. relative distances; Held et al., 2010; Watt et al., 2005). Thus, it is unclear exactly what magnitude change one might expect to see here. Additionally, despite having instructed them otherwise, we suspect that observers may have sometimes considered their prior knowledge about the usual viewing distances of object and scene categories to make their judgments (e.g., they know how far away a train on a railway bridge generally appears in the real world), which would bias distance ratings of fake-miniaturized scenes to be farther than they appeared in actuality. Indeed, in post-experiment questionnaires, some participants stated as much: “I judged the images by my past memories of seeing similar settings”; “[I just made] an estimate of how far things seemed to be away from me using life experience”.

Regardless of these issues, however, the effect of fake-miniaturization on perceived-distance ratings was extremely robust: the standardized effect size was very large (Cohen’s $d = 1.12$). Furthermore, 27 of 32 (84%) of the scene identities exhibited this pattern, meaning that the vast majority of fake-miniaturized scenes are perceived as closer than their non-manipulated counterparts. This point similarly applies to the other distance-rating experiments reported in this manuscript.

Occlusion Control vs. Original comparison, a model with Image Type was *not* a significantly better fit, $\chi^2(1) = 0.74$, $p = .390$, $R_m^2 = 0.02$, $R_e^2 = 0.43$; $\beta = -0.03$ [95% CI -0.10 to 0.04], $t = -0.85$.

Fig. 5. Image manipulations and results from Experiments 2a and 2b. (a) In addition to the original scenes (blue), a control occlusion manipulation (gray) was applied to fake-miniaturized scenes (orange). This manipulation placed grayscale occlusions where the blur regions were located in the fake-miniaturized scenes. (b) In Experiment 2a (distance ratings), this occlusion eliminated the perceived-distance effect of the fake-miniaturization manipulation (higher indicates closer ratings). (c) In Experiment 2b, we ran the same scene-memory task as in Experiment 1 with the addition of these occlusion control images; it was found that the occlusions eliminated the effect of fake-miniaturization on boundary extension, suggesting that mere reduction in image details or number of resolvable objects was not the source of fake-miniaturization's effects on boundary extension. In plots, colored circles are means across items (scene identities), error bars are 95% within-item confidence intervals, and light-gray points and connecting lines represent data for individual items in each condition. *** $p < .001$.

Similar results were found for Experiment 2b (scene-memory task), where a mixed-effect model with the main effect of Image Type was a better fit than a model without it, $\chi^2(2) = 47.54$, $p < .001$, $R_m^2 = 0.03$, $R_e^2 = 0.23$. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(2) = 1.17$, $p = .556$), suggesting that the effect of Image Type was stable across trials. Follow-up analyses comparing image types pairwise showed a significant main effect of Image Type for both the Fake-Miniature vs. Original comparison ($\chi^2(1) = 27.30$, $p < .001$, $R_m^2 = 0.04$, $R_e^2 = 0.27$; $\beta = 0.80$ [95% CI 0.55 to 1.04], $z = 6.45$, $p < .001$, OR = 2.22 [95% CI 1.74 to 2.82]) and the Fake-Miniature vs. Occlusion Control comparison ($\chi^2(1) = 37.19$, $p < .001$, $R_m^2 = 0.04$, $R_e^2 = 0.20$; $\beta = 0.75$ [95% CI 0.53 to 0.98], $z = 6.54$, $p < .001$, OR = 2.13 [95% CI 1.70 to 2.66]). By contrast, the Occlusion Control vs. Original comparison yielded no significant effect of Image Type, $\chi^2(1) = 0.01$, $p = .929$, $R_m^2 < 0.001$, $R_e^2 = 0.19$; $\beta = 0.01$ [95% CI -0.21 to 0.23], $z = 0.09$, $p = .929$, OR = 1.01 [95% CI 0.81 to 1.26]).

To summarize these results, occluding the blurred regions in the fake-miniaturized images eliminated the effects of fake-miniaturization on both perceived distance and boundary extension. This suggests that other alternative explanations related to image resolvability were not the source of the boundary-extension effects elicited by fake-miniaturization. Instead, these results are consistent with the conclusion that reduction in perceived distance was the cause.

Experiments 3a and 3b: Perceived Distance Per Se?

Experiments 2a and 2b ruled out image resolvability as an alternative explanation for fake-miniaturization's effects on boundary extension, by making such details fully unresolvable (and in fact, entirely *absent*) via occlusion. However, it is still possible that the *presence* of the low-level perceptual differences introduced by the manipulation (i.e., the blur gradient and saturation) were the cause, rather than the change in perceived distance induced by the manipulation. We addressed this possibility by introducing one simple change to the fake-miniaturization manipulation: we rotated the blur gradient by 90 degrees. This misaligned the blur gradient (now side-to-side) and the gradient of viewing-distances in the scene (bottom-to-top) with one another (Fig. 6a). As in Experiments 2a and 2b, the present experiments tested whether this “rotated control” condition would reduce or eliminate the changes in perceived distance and boundary extension introduced by the image manipulation, and thus rule out lower-level explanations for such differences.

Methods

Participants

Two groups of 100 participants each were recruited from Prolific for Experiments 3a and 3b (i.e., 200 participants total). We chose sample sizes of 100 in both experiments to match the sample sizes used in Experiments 1 and 2b.

Stimuli

The same 32 scene images from the previous experiments were used here. These images were cropped to be square. Cropping was largely symmetrical, except in cases where prominent scene content (e.g., a house) would be cut off. In addition to the original and fake-miniaturized versions of each scene image, we introduced a “rotated control” manipulation. To create this control condition, we used <https://tiltshiftmaker.com> as in Experiment 1 to add saturation and blur, applying the horizontal (rotated) blur gradient rather than the default

vertical blur gradient (i.e., for each image, the blur gradient was rotated clockwise by 90 degrees such that it was horizontal). This rotation was expected to disrupt the percept of a shallow depth-of-field introduced by the usual vertical blur gradient, thereby reducing or eliminating the change in perceived distance (Held et al., 2010). Examples can be seen in Figure 6a. Masks were the same as in Experiment 1, except that the images used to create the masks were scrambled in a 13x13 grid. All stimuli were displayed at 425x425 pixels in the participant's Web browser for Experiment 3a and 500x500 pixels for Experiment 3b.

Design, Procedure, and Analyses

The design, procedure, exclusion criteria, and analyses of Experiment 3a and 3b were the same as in Experiments 2a and 2b, respectively (with the Rotated Control condition used here instead of the Occlusion Control condition). Nine participants and 14 remaining trials were excluded from Experiment 3a, and 12 participants and three remaining trials were excluded from Experiment 3b.

Results

Results can be seen in Figure 6. Inspection of this figure suggests several findings. First, as predicted, distance ratings in Experiment 3a were dependent on Image Type: Fake-Miniature images were rated as closer ($M=2.12$) than both the Original images ($M=1.91$) and Rotated Control ($M=1.92$), while the latter two image types were not rated differently from one another (Fig. 6b). These distance-rating results again confirm that the Fake-Miniature manipulation had the intended effect of reducing perceived distance, while the closely matched Rotated Control condition “broke” this perceived-distance effect. Crucially, this same rotation manipulation appeared to reduce boundary-extension effects: In Experiment 3b, participants reported that the probe image appeared closer than the target more often for Fake-Miniature images ($M=70.9\%$) than for both the Original ($M=57.6\%$) and Rotated Control ($M=63.5\%$) images (Fig. 6c).

These observations were confirmed statistically. For Experiment 3a (distance ratings), a mixed-effect model with the main effect of Image Type was a significantly better fit than one without, $\chi^2(2) = 40.87, p < .001, R_m^2 = 0.02, R_e^2 = 0.39$. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(2) = 0.22, p = .896$), suggesting that the effect of Image Type was stable across trials. Follow-up analyses comparing image types pairwise showed a significant main effect of Image Type for both the Fake-Miniature vs. Original comparison ($\chi^2(1) = 12.84, p < .001, R_m^2 = 0.02, R_e^2 = 0.44; \beta = 0.21$ [95% CI 0.10 to 0.32], $t = 3.78$), and the Fake-Miniature vs. Rotated Control comparison ($\chi^2(1) = 24.98, p < .001, R_m^2 = 0.01, R_e^2 = 0.41; \beta = 0.18$ [95% CI 0.11 to 0.26], $t = 5.01$). By contrast, for the Rotated Control vs. Original comparison, a model with Image Type was *not* a significantly better fit, $\chi^2(1) = 0.25, p = .618, R_m^2 = 0.01, R_e^2 = 0.39; \beta = 0.01$ [95% CI -0.06 to 0.10], $t = 0.50$).

Fig. 6. Image manipulations and results from Experiments 3a and 3b. (a) In addition to fake-miniaturization (orange), a rotated control manipulation (gray) was applied to distant scenes (blue). This manipulation rotated the blur gradient of the fake-miniaturization manipulation by 90 degrees, which misaligned the blur gradient and the gradient of scene viewing-distances. (b) In Experiment 3a, this rotation eliminated the perceived distance effect of the image manipulation, as shown by distance ratings (higher indicates closer ratings). (c) In Experiment 3b, we ran the same scene-memory task as in Experiment 1 with the addition of these rotated control images; it was found that the rotation diminished the effect of the image manipulation on boundary extension. In plots, colored circles are means across items (scene identities), error bars are 95% within-item confidence intervals, and light-gray points and connecting lines represent data for individual items in each condition. *** $p < .001$, ** $p < .01$, * $p < .05$.

Similar results were found for Experiment 3b (the scene-memory task), where a mixed-effects model with the main effect of Image Type was a significantly better fit than one without, $\chi^2(2) = 26.17$, $p < .001$, $R_m^2 = 0.02$, $R_e^2 = 0.14$. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(2) = 1.10$, $p = .578$), suggesting that the effect of Image Type was stable across trials. Follow-up analyses comparing image types pairwise showed a significant main effect of Image Type for the Fake-Miniature vs. Original comparison ($\chi^2(1) = 24.66$, $p < .001$, $R_m^2 = 0.02$, $R_e^2 = 0.16$; $\beta = 0.62$ [95% CI 0.39 to 0.85], $\zeta = 5.33$, $p < .001$, OR = 1.86 [95% CI 1.48 to 2.34]). Interestingly, a significant main effect of Image Type was also found for the Rotated Control vs. Original comparison ($\chi^2(1) = 4.31$, $p = .038$, $R_m^2 = 0.01$, $R_e^2 = 0.15$; $\beta = 0.27$ [95% CI 0.02 to 0.51], $\zeta = 2.14$, $p = .032$, OR = 1.31 [95% CI 1.02 to 1.67]), suggesting at least some low-level influence of the rotated image manipulation on boundary-extension effects. Perhaps the gradual blur along the sides of the scene images induced some boundary extension, which would be consistent with work by Gagnier and colleagues (2013) showing that cropping one side of an object (e.g., a basketball) at the edge of an image leads to more boundary extension in that direction. Crucially, however, these effects were significantly greater in the Fake-Miniature condition compared with the Rotated Control condition: a mixed-effects model with the Image Type factor was a significantly better fit than one without ($\chi^2(1) = 7.98$, $p = .005$, $R_m^2 = 0.01$, $R_e^2 = 0.13$; $\beta = 0.34$ [95% CI 0.12 to 0.57], $\zeta = 3.03$, $p = .002$, OR = 1.42 [95% CI 1.13 to 1.77]), demonstrating that the rotation of the blur gradient indeed diminished the effects of the image manipulation on boundary transformation.

To summarize these results, simply rotating the blur gradient so that it was misaligned with the distance gradient in the image dramatically reduced the manipulation's effects on both perceived distance and boundary extension. This suggests that changes in perceived distance per se—rather than in low-level image properties alone—are the source of fake-miniaturization's effect on boundary extension.

Experiments 4a and 4b: Generalization to a Different Distance Manipulation

The previous experiments demonstrated that fake-miniaturization increased boundary extension for otherwise identical scene images by reducing perceived distance. In a final set of studies, we asked whether this phenomenon would generalize to a completely different distance manipulation.

We were inspired by Ritchie and Van Buren (2020), who highlighted the perceptual principles at play in the surrealist René Magritte's paintings, many of which depict silhouettes that themselves contain a scene (e.g., a rustic evening scene inside a silhouette of a man). These paintings are bistable: When the silhouette's interpretation is as perceptual-ground, it appears as a "window" onto a distant scene with 3D structure; when its interpretation is as perceptual-figure, the silhouette appears to have a 2D "scene-texture" draped upon it.

To induce a similar perceptual switch on real-world scenes, we used a technique that we call "spherization", in which an image is warped such that it appears wrapped upon a spherical object (Fig. 7a). This distorts and minimizes the rich perspectival cues to distance inherent in spatial scenes, but preserves other visual and semantic properties. We predicted that spherization would make distant scenes appear closer by minimizing perspectival cues to distance, and would thus induce boundary extension on otherwise identical images. We tested these predictions in Experiments 4a (distance ratings) and 4b (scene-memory task).

Methods

Participants

Two groups of 100 participants each were recruited from Prolific for both Experiments 4a and 4b (i.e., 200 participants total). We chose sample sizes of 100 in both experiments to match the sample size used in most of the previous experiments reported here.

Stimuli

To maximize our chances of finding effects of our image manipulation on perceived distance and boundary extension, we chose the 40 scene images from Bainbridge & Baker (2020a) that showed the highest boundary-contraction effects in their studies (32 for the main experiment, and eight additional scenes used as example or practice items). We cropped these images with circular masks. We then created “spherized” versions of the original scenes by applying the spherization effect in GIMP open-source photo-editing software (with a curvature setting of 0.75); this effect warps the pixels in an image such that the image appears to be wrapped around a sphere. Finally, we applied a radial gradient (also spherized) as a shadow to enhance the 3D appearance of the spherization, as if the sphere were lit from the front. Together, these manipulations made the scenes appear as textures atop a convex dome or sphere, rather than as a scene with rich 3D structure and depth.³ Example images can be viewed in Figure 7a. Rectangular masks were created by collecting 32 additional images from the Bainbridge & Baker (2020a) stimulus set and box-scrambling them in a 14x14 grid. All stimuli were displayed at 425x425 pixels in the participant’s Web browser for Experiment 4a and 500x500 pixels for Experiment 4b.

Design, Procedure, and Analyses

The design, procedure, exclusion criteria, and analyses of Experiment 4a (distance ratings) was the same as in Experiments 2a and 3a, except that there were only two conditions here (Original and Spherized), rather than three. Thus, there were two stimulus lists rather than six (counterbalanced across participants), and mixed-effects models only had two conditions for the Image Type factor (Original and Spherized, sum-coded as -0.5 and 0.5 , respectively). The design, procedure, exclusion criteria, and analyses of Experiment 4b (scene-memory task) were the same as in Experiment 1. Four participants and 19 remaining trials were excluded from Experiment 4a, and 12 participants and one remaining trial were excluded from Experiment 4b.

³ We note that we also considered examining the effects of concave rather than convex spherized images. However, after testing a variety of ways to generate concave examples (including with different shadowing effects), they often continued to appear convex. This may reflect a general convexity bias in perception (e.g., Kleffner & Ramachandran, 1992). We therefore leave the question of how concavity manipulations influence boundary transformation for future work.

Fig. 7. Image manipulation and results from Experiment 3. (a) A spherization distortion was applied to images of distant scenes (blue), creating a situation where perspective cues to distance were distorted or diminished (orange). (b) A similar perceived-distance effect was obtained here as in the previous experiments

using fake-miniaturization: Spherized images were perceived as closer-up than their original counterparts. (c) In the scene-memory task, participants showed more boundary extension for spherized images than the originals. In plots, the colored circles are item means, and error bars are 95% within-item confidence intervals. Light-gray points and connecting lines represent data for individual items (scene identities). *** $p < .001$.

Results

Results can be seen in Figure 7. Inspection of this figure suggests that, as predicted, distance ratings were dependent on Image Type: Spherized images were rated as closer ($M=2.45$) than the Original images ($M=2.34$) (Fig. 7b). These distance-rating results show that spherization had the intended effect of reducing perceived distance. Crucially, spherization also appeared to increase boundary extension: in Experiment 4b, participants reported that the probe image appeared closer than the target more often for the Spherized condition ($M=60.9\%$) as compared to the Original condition ($M=50.3\%$) (Fig. 7c).

These observations were confirmed statistically. For Experiment 4a (distance ratings), a mixed-effects model with the main effect of Image Type was a significantly better fit than one without, $\chi^2(1) = 13.59, p < .001, R_m^2 = 0.004, R_e^2 = 0.61; \beta = 0.11$ [95% CI 0.06 to 0.17], $t = 4.06$. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(1) = 0.20, p = .653$), suggesting that the effect of Image Type was stable across trials. Likewise for Experiment 4b (scene-memory task), a mixed-effects model with the main effect of Image Type was a significantly better fit than one without, $\chi^2(1) = 17.84, p < .001, R_m^2 = 0.02, R_e^2 = 0.16; \beta = 0.50$ [95% CI 0.29 to 0.71], $\chi = 4.74, p < .001, OR = 1.65$ [95% CI 1.34 to 2.03]. Adding an interaction of Image Type and Trial Number did not improve model fit ($\chi^2(1) = 2.08, p = .150$), suggesting that the effect of Image Type was stable across trials.

To summarize these results, the spherization manipulation increased how often participants reported a probe image as closer (more “zoomed-in”) than a target image, which is an indication that boundary extension occurred on the memory of the target image. Together with the distance ratings on these images, these results indicate that the effect of perceived distance on boundary extension is quite general in nature, extending beyond fake-miniaturization to manipulations like spherization that distort and minimize perspectival cues to distance.

General Discussion

What drives distortions of memory for visual scenes? Visual memories are certainly influenced by the contents of what we perceive: the particular objects, settings, events, and so on. But beyond such content, we demonstrated across seven experiments that perceived distance itself plays a crucial role in determining how scene boundaries are reshaped in memory.

Our results are consistent with previous work showing an association between viewing distance and boundary transformation (Bainbridge & Baker, 2020a; Bertamini et al., 2005; Park et al., 2021). These previous studies were careful and systematic in their investigations of this association; yet despite such efforts, these studies were unable to conclude whether viewing distance alone has a causal role in boundary transformation, as viewing distance almost always covaries with other low- and high-level image properties. Thus, our results go beyond these previous findings in a crucial way: By using subtle image manipulations to selectively alter the perceived

viewing distance of scenes while holding other perceptual and semantic properties constant, we provided strong support for the hypothesis that perceived distance is independently sufficient to modulate the effects of boundary transformation.

The Construction of Image Memories

Although boundary extension has been investigated for over 30 years (starting with Intraub & Richardson, 1989), there is still debate about what kind of memory process it reflects, and its functional role, if any, in cognition (c.f. Bainbridge & Baker, 2020a, 2020b; Intraub, 2020). The original proposal was that boundary extension is an anticipatory “scene-construction” error: Scenes automatically activate a prediction of likely information outside the immediate view—perhaps for facilitating recognition in subsequent fixations—and this predicted information is incorporated into memory as if it were veridical (Intraub & Richardson, 1989; for review, see Hubbard et al., 2010; Intraub, 2010). More recent work has argued for unifying accounts of boundary extension and contraction (Bainbridge & Baker, 2020a, 2020b; Park et al., 2021; for an alternative view, see Intraub, 2020). One such proposal is that memory distortions for scene boundaries are prior-based, such that scene memories are normalized toward prototypes of their category (e.g., bathrooms are usually viewed close-up; Bainbridge & Baker, 2020a)—which may serve to minimize encoding errors from noisy or incomplete inputs (Bartlett, 1932; Hemmer & Steyvers, 2009). Our results do not settle this debate, but they have implications for both types of proposals. Our findings suggest that scene-construction theories need to account for the fact that simply perceiving something as close is sufficient to drive the construction process, and our findings suggest that prior-based theories need to explain generic effects of viewing distance on boundary transformation that are not tied to category-specific priors (e.g., how similar effects of viewing distance emerge across diverse categories like *beach* or *city-street*).

The present results add to the literature investigating the ways that perceiving and remembering interact in the service of constructing visual memories. For example, a central result in visual cognition is that the mind “plays forward” the movements of objects in memory—a phenomenon called *representational momentum* (Freyd, 1983; Freyd & Finke, 1984; Hubbard, 2005)—and such effects have recently been extended to surprisingly complex types of change, including to the physical states of objects (e.g., melting ice; Hafri et al., in press). Recent work has even shown that the mind adds vividness to scene images, such that the world is remembered in “higher definition” than it actually appears—a phenomenon called *vividness extension* (Rivera-Aparicio et al., 2021). Together with these findings, the current results broaden the view of what kinds of contents the mind may add to visual memories, and under what conditions it may do so. In our case, we have found that beyond scene content itself, visual memories are biased by the spatial contexts in which they are formed.

Open Questions and Possible Extensions

In our experiments using fake-miniaturization, we altered perceived distance by manipulating the apparent scale of the space; however, this manipulation also alters the apparent real-world size of the objects in the scenes. On one hand, there is reason to think that these scene and object factors may be dissociable, and that viewing distance itself may be a sufficient factor for modulating boundary transformation: Recent work has demonstrated that even for sparse spatial scenes without rich object cues, changes in viewing distance are associated with changes in boundary-transformation effects (Park et al., 2021). On the other hand, it may be highly challenging to demonstrate possible effects of object size that are independent of viewing distance, given

that an object that is large in real-world size also, by definition, occupies a larger-scale portion of space. Thus, we leave this question for future work.

Although our studies support the hypothesis that changes in perceived distance (caused by subtle image manipulations) modulate boundary transformation, they were not designed to determine the precise psychophysical relationship between these properties. However, other studies suggest that there is a continuous change from boundary-extension to boundary-contraction effects that is parametrically related to viewing distance (Bainbridge & Baker, 2020a; Park et al., 2021). Future work may capitalize on the image manipulations we used to investigate this relationship further while retaining precise control over other image factors.

Beyond perceived distance, there are almost surely other properties that contribute to visual memory distortions, and future work may take a similar approach to ours—using targeted image manipulations to investigate the causal contribution of these properties to such distortions. Some properties have already been explored in the context of boundary transformation. For example, highly emotional images are associated with greater boundary contraction (Takarangi et al., 2016); its causal role could be tested by eliciting adaptation aftereffects (e.g., Rutherford et al., 2008) or changing the observer's emotional state (e.g., Green et al., 2019). The number and prominence of objects may also contribute: Park et al. (2021) found that fewer in-scene objects are associated with greater boundary extension (also Bainbridge & Baker, 2020a). To test the causal role of this property, researchers might take advantage of numerosity illusions (DeWind et al., 2020; Franconeri et al., 2009) or ambiguous figure-ground images with different perceived numerosities (Wagemans et al., 2012).

Finally, the image manipulations we used may be valuable for other areas of visual cognition research beyond scene memory, as they decouple perceived distance from other image content. For example, these manipulations might be used to reveal how computations of size constancy differ between perceiving and acting on objects in scenes (e.g., for detection vs. grasping tasks; Chen et al., 2018). Other studies might use them to ask about the conditions under which cognitively distinct scene- and object-processing mechanisms are engaged (e.g., Bryan et al., 2016; Cheng et al., 2021; Park et al., 2011; Poltoratski & Tong, 2014); perhaps fake-miniaturized or spherized images would preferentially engage object-processing mechanisms.

Concluding Remarks

How does memory reconstruct scenes that are no longer visible, and how are such representations distorted? The initial discovery of boundary extension revealed a remarkable aspect of scene memory—that the mind adds information that was not there in the first place (Intraub & Richardson, 1989)—and inspired a wealth of work exploring its robustness to display-, timing-, and response-related factors, and its possible functional role in cognition (Hubbard et al., 2010; Intraub, 2010). Nevertheless, determining the influence of scene properties themselves has proved elusive, due to the inevitable covariance of such properties with one another. Here we overcame this challenge by adopting a causal approach with subtle and targeted image manipulations. In doing so, we demonstrated that perceived viewing distance plays a causal role in driving boundary extension; more broadly, we showed how such image manipulations can be used to reveal the rich ways that visual perception and memory interact.

Author Contributions

Alon Hafri: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Software, Visualization, Writing - original draft, and Writing - review & editing.
Shreya Wadhwa: Conceptualization, Formal analysis, Resources, and Writing - review & editing.
Michael F. Bonner: Conceptualization, Funding acquisition, Project administration, Resources, Supervision, and Writing - review & editing.

Acknowledgments

For helpful discussion and/or comments on previous drafts, the authors thank Chaz Firestone and Caterina Magri, and members of the JHU Cognitive Neuroscience & Machine Learning Lab. This work was partially supported by NSF SBE Postdoctoral Research Fellowship SMA-2105228 awarded to A.H.

References

- Bainbridge, W. A., & Baker, C. I. (2020a). Boundaries Extend and Contract in Scene Memory Depending on Image Properties. *Current Biology*, 30(3), 537-543.e3.
- Bainbridge, W. A., & Baker, C. I. (2020b). Reply to Intraub. *Current Biology*, 30(24), R1465–R1466.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press.
- Bertamini, M., Jones, L. A., Spooner, A., & Hecht, H. (2005). Boundary Extension: The Role of Magnification, Object Size, Context, and Binocular Information. *Journal of Experimental Psychology: Human Perception and Performance*, 31(6), 1288–1307. <https://doi.org/10.1037/0096-1523.31.6.1288>
- Bryan, P.B., Julian, J.B., & Epstein, R.A. (2016) Rectilinear edge selectivity is insufficient to explain the category selectivity of the Parahippocampal Place Area. *Frontiers in Human Neuroscience*, 10(137).
- Chen, J., Sperandio, I., & Goodale, M. A. (2018). Proprioceptive distance cues restore perfect size constancy in grasping, but not perception, when vision is limited. *Current Biology*, 28(6), 927-932.
- Cheng, A., Walther, D. B., Park, S., & Dilks, D. D. (2021). Concavity as a diagnostic feature of visual scenes. *NeuroImage*, 232, 117920
- DeWind, N. K., Bonner, M. F., & Brannon, E. M. (2020). Similarly oriented objects appear more numerous. *Journal of Vision*, 20(4), 4. <https://doi.org/10.1167/jov.20.4.4>
- Franconeri, S. L., Bemis, D. K., & Alvarez, G. A. (2009). Number estimation relies on a set of segmented objects. *Cognition*, 113(1), 1–13. <https://doi.org/10.1016/j.cognition.2009.07.002>

- Freyd, J. J. (1983). The mental representation of movement when static stimuli are viewed. *Perception & Psychophysics*, 33(6), 575–581. <https://doi.org/10.3758/BF03202940>
- Freyd, J. J., & Finke, R. A. (1984). Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 126–132. <https://doi.org/10.1037/0278-7393.10.1.126>
- Gagnier, K. M., Dickinson, C. A., & Intraub, H. (2013). Fixating picture boundaries does not eliminate boundary extension: Implications for scene representation. *The Quarterly Journal of Experimental Psychology*, 66(11), 2161–2186. <https://doi.org/10.1080/17470218.2013.775595>
- Gottesman, C.V. (2011). Mental layout extrapolations prime spatial processing of scenes. *J. Exp. Psychol. Hum. Percept. Perform.* 37, 382–395.
- Green, D. M., Wilcock, J. A., & Takarangi, M. K. T. (2019). The role of arousal in boundary judgement errors. *Memory & Cognition*. <https://doi.org/10.3758/s13421-019-00914-8>
- Hafri, A., Boger, T., & Firestone, C. (in press). Melting ice with your mind: Representational momentum for physical states. *Psychological Science*.
- Held, R. T., Cooper, E. A., O'Brien, J. F., & Banks, M. S. (2010). Using blur to affect perceived distance and size. *ACM Transactions on Graphics*, 29(2), 1–16. <https://doi.org/10.1145/1731047.1731057>
- Hemmer, P., & Steyvers, M. (2009). A Bayesian Account of Reconstructive Memory. *Topics in Cognitive Science*, 1(1), 189–202. <https://doi.org/10.1111/j.1756-8765.2008.01010.x>
- Hubbard, T. L. (2005). Representational momentum and related displacements in spatial memory: A review of the findings. *Psychonomic Bulletin & Review*, 12(5), 822–851. <https://doi.org/10.3758/BF03196775>
- Hubbard, T. L., Hutchison, J. L., & Courtney, J. R. (2010). Boundary extension: Findings and theories. *Quarterly Journal of Experimental Psychology*, 63(8), 1467–1494.
- Intraub, H. (2010). Rethinking Scene Perception. In *Psychology of Learning and Motivation* (Vol. 52, pp. 231–264). Elsevier. [https://doi.org/10.1016/S0079-7421\(10\)52006-1](https://doi.org/10.1016/S0079-7421(10)52006-1)
- Intraub, H. (2020). Searching for boundary extension. *Current Biology*, 30(24), R1463–R1464.
- Intraub, H., & Dickinson, C. A. (2008). False Memory 1/20th of a Second Later: What the Early Onset of Boundary Extension Reveals About Perception. *Psychological Science*, 19(10), 1007–1014. <https://doi.org/10.1111/j.1467-9280.2008.02192.x>
- Intraub, H., & Richardson, M. (1989). Wide-Angle Memories of Close-Up Scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(2), 179–187.
- Intraub, H., Bender, R.S., and Mangels, J.A. (1992). Looking at pictures but remembering scenes. *J. Exp. Psychol. Learn. Mem. Cogn.* 18, 180–191.
- Intraub, H., Gottesman, C.V., & Bills, A.J. (1998). Effects of perceiving and imagining scenes on memory for pictures. *J. Exp. Psychol. Learn. Mem. Cogn.* 24, 186–201.

- Kleffner, D. A., & Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics*, 52(1), 18-36.
- Lescroart, M. D., Stansbury, D. E., & Gallant, J. L. (2015). Fourier power, subjective distance, and object categories all provide plausible models of BOLD responses in scene-selective visual areas. *Frontiers in Computational Neuroscience*, 9. <https://doi.org/10.3389/fncom.2015.00135>
- Mullally, S. L., Intraub, H., & Maguire, E. A. (2012). Attenuated Boundary Extension Produces a Paradoxical Memory Advantage in Amnesic Patients. *Current Biology*, 22(4), 261–268.
- Nakagawa, S., Johnson, P. C. D., & Schielzeth, H. (2017). The coefficient of determination R^2 and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of The Royal Society Interface*, 14(134), 20170213. <https://doi.org/gddpnq>
- Park, S., Brady, T. F., Greene, M. R., & Oliva, A. (2011). Disentangling scene content from spatial boundary: Complementary roles for the parahippocampal place area and lateral occipital complex in representing real-world scenes. *Journal of Neuroscience*, 31(4), 1333-1340.
- Park, J., Josephs, E. L., & Konkle, T. (2021, January 26). Systematic transition from boundary extension to contraction along an object-to-scene continuum. *PsyArXiv*, <https://doi.org/10.31234/osf.io/84exs>
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163. <https://doi.org/10.1016/j.jesp.2017.01.006>
- Quinn, P. C., & Intraub, H. (2007). Perceiving “outside the box” occurs early in development: Evidence for boundary extension in three- to seven-month-old infants. *Child Development*, 78(1), 324-334.
- Ritchie, J. B., & van Buren, B. (2020). When Scenes Look Like Materials: René Magritte’s Reversible Figure–Ground Motif. *Art and Perception*, 1–12. <https://doi.org/10.1163/22134913-bja10002>
- Rivera-Aparicio, J., Yu, Q., & Firestone, C. (2021). Hi-def memories of Lo-def scenes. *Psychonomic Bulletin & Review*, 28, 928–936.
- Rutherford, M. D., Chattha, H.M., Krysko, K.M. (2008). The use of aftereffects in the study of relationships among emotion categories. *Journal of Experimental Psychology: Human Perception and Performance*, 34(1), 27-40.
- Selya, A. S., Rose, J. S., Dierker, L. C., Hedeker, D., & Mermelstein, R. J. (2012). A practical guide to calculating Cohen’s f^2 , a measure of local effect size, from PROC MIXED. *Frontiers in Psychology*, 3, 1– 6.
- Takarangi, M. K. T., Oulton, J. M., Green, D. M., & Strange, D. (2016). Boundary Restriction for Negative Emotional Images Is an Example of Memory Amplification. *Clinical Psychological Science*, 4(1), 82–95. <https://doi.org/10.1177/2167702615569912>
- Vishwanath, D., & Blaser, E. (2010). Retinal blur and the perception of egocentric distance. *Journal of Vision*, 10(10): 26. <https://doi.org/10.1167/10.10.26>

- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, 138(6), 1172–1217. <https://doi.org/10.1037/a0029333>
- Watt, S. J., Akeley, K., Ernst, M. O., & Banks, M. S. (2005). Focus cues affect perceived depth. *Journal of Vision*, 5(10), 7. <https://doi.org/10.1167/5.10.7>