Title: Systematic transition from boundary extension to contraction along an object-to-scene continuum

Running head: Transition from boundary extension to contraction

Authors: Jeongho Park¹, Emilie Josephs¹, & Talia Konkle¹

¹Department of Psychology, Harvard University

Corresponding author: Talia Konkle, tkonkle@fas.harvard.edu

Abstract

After viewing a picture of an environment, our memory of it typically extends beyond what was presented—a phenomenon called boundary extension. But sometimes, memory errors show the opposite pattern—boundary contraction—and the relationship between these phenomena is controversial. We constructed virtual 3D environments, and created a series of views at different distances, from object close-ups to wide-angle indoor views, and tested for memory errors along this object-to-scene continuum. Boundary extension was evident for close-scale views, and transitioned parametrically to boundary contraction for far-scale views. However, this transition point was not tied to a specific position in the environment; instead, it tracked with judgments of the best looking view. We propose that boundary extension and contraction are in fact integrated phenomena, and we offer an account where competition between object-based and scene-based affordances determine whether a view will extend or contract in memory.

Keywords: scene memory, canonical views, virtual environments, aesthetic preferences

Introduction

Boundary extension is a classic psychological phenomenon, where a participant’s memory of a photograph includes details that extend beyond the boundaries of the original image (Intraub & Richardson, 1989). This phenomenon is thought to arise from automatic amodal scene construction: during perception, a view of a scene automatically evokes predictions about the surrounding environment. The visual trace of the image is integrated with this extended scene schema during encoding, leading to an extended representation of the image at retrieval. Interestingly, however, the amount of boundary extension evoked by a scene depends on its depicted scale—boundary extension is stronger for close-scale views, and is reduced or even absent as the view increases in distance (Bertamini et al., 2005; Intraub et al., 1992; McDunn et al., 2016). Given that viewing distance is a continuous variable with nearly infinite range, this raises an intriguing question: how does memory for an environment change across an extensive range of distances?

This question is particularly important in light of recent controversies surrounding the boundary extension phenomenon and how to account for evidence of memory errors in the opposite direction (Bainbridge & Baker, 2020a, 2020b; Intraub, 2020). In particular, initial studies showed only extension errors for close-up views (e.g., Intraub & Richardson, 1989); when subsequent studies with different stimuli revealed cases of boundary contraction, these were explained by invoking an additional mechanism of normalization to the mean of the stimulus set (e.g., the two-process *extension-normalization model*, Intraub et al., 1996). And, of particular relevance here, having a range of depicted distances in the stimulus set is one factor that has been argued to lead to stronger normalization processes (Intraub, 2020). From this theoretical perspective, Intraub argues that empirical findings of boundary contraction are indicative of normalization, and do not inform the phenomenon of boundary extension.

However, recently Bainbridge & Baker (2020a) have challenged this position. They probed scene memory errors across a substantial sampling of natural images, and found clear evidence for both boundary extension and contraction (Bainbridge & Baker, 2020a), even in paradigms that are not thought to induce normalization (Bainbridge & Baker 2020b). They suggest that these bi-directional memory errors do not necessarily require distinct mechanisms. Of key importance here, in an exploratory analysis, the authors found a correlation between the depicted distance of image and the magnitude of the memory errors. These results thus provide converging indirect evidence that the depicted distance of a scene view is a key factor underlying whether memory for a scene will extend or contract.

Critically, to date, no study has directly probed for the direction of memory errors over a wide range of distances, parametrically, in tightly controlled images. In the present work, we do just this, constructing virtual 3-dimensional naturalistic indoor environments in order to capture a smooth continuum of views from close-scale object-focused views to wider-angle scene-focused views. This empirical work thus has the potential to directly probe for integrated links between boundary extension and boundary contraction, tied to the depicted distance in the view.

Memory Errors along an Object-Scene Continuum

In the first experiment, we tested for memory errors for scene views taken from an object-to-scene continuum, using the same paradigm as Bainbridge and Baker (2020a).

Method

Participants. 377 participants from Amazon Mechanical Turk (M-Turk) participated in the Memory Experiment. All participants gave informed consent and were compensated with \$0.25 for each Human Intelligence Task (HIT). All procedures were approved by the Harvard University Human Subjects Institutional Review Board. The sample size for each condition of interest was 900 trials (45 per image x 20 environments); a bootstrap resampling method (Strong & Alvarez, 2019; see Supplemental Material) showed >95% power with 40 trials per condition, indicating this design is extremely high powered.

Stimuli. Computer generated (CGI) environments were constructed using the Unity video game engine (Unity Technologies, Version 2017.3.0). Twenty indoor environments were constructed, reflecting a variety of semantic categories (e.g., kitchens, bedrooms, laboratories, cafeterias, etc.). All rooms had the same physical dimensions (4 width x 3 height x 6 depth arbitrary units (a.u.) in Unity), with an extended horizontal surface along the back wall, containing a centrally-positioned object. Each environment was additionally populated with the kinds of objects typically encountered in those locations, creating naturalistic, “object-rich” CGI environments.

Fig. 1. Stimuli. (A) Example images along the 30-point continuum, from two different environments. (B) Schematic of the 3d environment in unity with views taken from a camera from the front to the back of the scene. The field of view (FOV) of the camera, and the rotation angle of the camera was gradually changed to interpolate between an object-centered view of a single central object and a wider field of view of the entire scene.

Images spanning a continuum of distances from the central object were captured from each environment, ranging from a close-up view of the object to a far-scale view, which included the whole room. Images were generated by systematically varying the location of the camera (hereafter “Position”) along 30 evenly-spaced points arrayed from the “front” to the “back” of the room (i.e., from right in front of the central object to across the room from it, Fig. 1A). Close-up views were captured with a smaller camera field of view (FOV), so that only the central object appeared in the frame, and the FOV increased logarithmically with each step away from the object. The camera angle was parallel to the floor plane for far-scale views, and was gradually adjusted downward for closer positions, so that the central object was always at the center of the image (Fig. 1B). These camera parameters were used for all 20 environments, yielding 600 unique stimuli (20 environments x 30 positions). See <https://bit.ly/3o5L8zc> for all stimuli.

Experimental Design. The experiment paradigm followed the procedures from Bainbridge & Baker (2020a). Each trial started with a central fixation cross for 1 sec. The first stimulus was presented for 250ms, followed by 5 mosaic-scrambled mask images (50ms x 5; random order). The second stimulus was then presented for 1 sec. Unbeknownst to participants, the second stimulus was always identical to the first. After the second stimulus, participants were asked: “Compared to the first image, the second image is...”, with two response options of “Closer” or “Farther” (Fig. 2A). Participants were instructed to answer within 3 seconds using a keyboard; if no response was entered within this time, a warning message was shown before moving to the next trial. Responses (closer/farther) and reaction times (RT) were recorded.

Stimuli and masks were presented at 350 x 467 pixels. To create the masks, 15 CGI scenes (not used in the experiment) were broken down into 25 x 25 mosaic grids, then reassembled by randomly sampling from this pool.

The 600 images were divided into sets, such that a single HIT contained images from each of the 30 positions, selected from 10 of the 20 environments (thus, each HIT had 3 images from the same environment, viewed from different positions). The presentation order of stimuli within each set was randomly determined for each HIT. This method of dividing the images into sets ensured that all participants were exposed to the full range of views (i.e., Position 1-30) within a HIT, and allowed for the full stimulus set to be tested with 20 different image sets. Across all HITs and subjects, 45 trials were collected for each image.

Data Pre-processing. Trials with reaction times exceeding the allowable response time window were excluded (3 sec from stimulus onset; 2% of trials). Participants were excluded if (1) their average RT across all trials was faster than 3 standard deviations from the mean over all participants, (2) more than 50% of their trials exceeded the allowable response time window, or (3) their average RT across all trials was faster than 2 standard deviations from the mean and answered more than 90% of trials with the same response key, leading to the exclusion of 1.6% of the data. Reaction times were log transformed prior to trimming to account for the right skewness of RT distributions (Palmer et al., 2011; Ratcliff, 1979).

Calculating memory distortion scores. “Closer” responses were assigned a score of -1 and “Farther” responses were assigned a score of +1. These scores were averaged to obtain a *memory distortion score* for each image. Memory distortion scores thus fall in a range between -1 and 1,

where negative scores indicate boundary extension (remembering the initial scene as further away than it was), positive scores indicate boundary contraction (remembering the scene closer than it was), and 0 indicates accurate position memory, with no distortion or bias one way or the other.

Linear mixed-effects logistic regression. Given the hypothesis that the size and direction of memory distortions would vary continuously as a factor of depicted distance, we modelled memory distortion scores as a factor of Position using linear mixed-effects logistic regression, implemented with the R code package lme4 (Bates et al., 2015). For the dichotomous response outcomes in this study (Closer or Farther), a binomial logit link function was used in a generalized linear mixed model. The model included *Position* as a fixed-effect term and *Environment* as a random intercept.

To test the significance of the fixed effect of Position, we compared our model against a null model, which had the same random effects structure but without the fixed effect. If the full model performed significantly better than the null model at predicting the data, then Position significantly contributed to the degree of memory distortion. P-values were estimated using the code package lmerTest with Satterthwaite approximation (Kuznetsova et al., 2017).

If the direction of the memory distortion varied continuously as a function of depicted distance to the central object, then there should be a Position where there is no distortion or bias (i.e., average memory distortion score = 0). These *transition points* were estimated for each environment by analytically solving the logistic regression function (Fischer & Whitney, 2014). Specifically, we used mixed-effects logistic regression of the form $y = 1/(1 + e^{-((b_0 + a_i) + b_1x)})$, where y is the proportion of ‘farther’ responses, x is the Position, b_0 is the fixed intercept, a_i ($i=1:20$) is the random intercept of the individual environment, and b_1 is the fixed slope. The transition point was defined as the x value (Position) at which $y = 0.5$ (equally likely to respond to Farther or Closer, i.e. average memory distortion score = 0). This analysis yielded a transition point for each of the 20 environments, and we report the *transition range* along the continuum falling 1 standard deviation around the mean transition point, across environments.

Results

Do memory representations show systematic boundary extension across the entire object-scene continuum? Or, do a portion of the responses reflect boundary contraction? We found systematic evidence for the latter possibility (Fig. 2B). That is, close-scale views elicited strong boundary extension effects (i.e., negative memory distortion scores; 44.8% of images), while far-scale views elicited strong boundary contraction effects (i.e., positive memory distortion scores; 51% of images; 4.1% of the images showed no boundary effect).

Fig. 2. Memory Experiment Paradigm and Results. (A) Procedure of memory experiment. The first image was shown for 250ms, followed by 250ms of mosaic-scrambled masks. The second image was presented for 1 sec, and participants were asked to answer whether the second image was Closer or Farther compared to the first image. Unbeknownst to participants, the second image was always identical to the first one. (B) Memory distortion scores averaged across environments, by each position. The error bars represent the standard error of the mean. The negative score indicates boundary extension, and the positive score indicates boundary contraction. Overall, there was a smooth transition from boundary extension to contraction as it changed from object-centered to scene-centered images.

The magnitude of the memory distortion scores varied continuously between these extremes, (Spearman rank correlation between position and memory distortion: $\rho=0.66$, $p=8.76 \times 10^{-77}$, Supplemental Fig. 3A, confirmed with mixed-effects logistic regression, estimated odds ratio: 1.03, $p < 0.001$). Interestingly, across all environments there was a relatively consistent transition point (mean position = 14.3/30, s.d.=3.0), depicted as a vertical gradient zone in Fig. 2B, which reflects the positions within 1 standard deviation of the mean (see Supplementary Fig. 1 for all environments' transition points). Given this consistency, these data indicate that the systematic transition between extension and contraction in memory likely depends on something that all scenes had in common: namely the same virtual layout of walls, with a back counter or large object surface and a small central object.

Overall, these results demonstrate that the direction and magnitude of memory distortions for these views is well-predicted by the distance depicted in an image, with a systematic point at which the memory distortion shifts from extension to contraction. We next explored two possible explanations for the cognitive significance of this transitional range.

Perceived Reachability and Canonical Views

One possibility is that this transition range is related to whether the depicted space appears subjectively within-reach versus subjectively out-of-reach. Indeed, visual inspection of the images falling in the transition range reveals a set of views where the structures at the front of the room (i.e. tables or counters) appeared just reachable with the hand. Previous work has shown that the distinction between peripersonal and extrapersonal space is very salient in visuo-spatial encoding (Cléry et al., 2015), and neuropsychological studies have shown separable attentional mechanisms for near and far space (e.g. double dissociations in visual hemi-neglect, Cowey et al., 1994, 1998;

Halligan & Marshall, 1991). Additionally, we have recently shown that the human visual system is sensitive to the distinction between views of intermediate-scale “reachspaces” and navigable-scale scenes (Josephs & Konkle, 2019, 2020). Thus, it is possible that the subjective reachability of an environment is a behaviorally-relevant signal toward which perceptual representations are biased. To explore this possibility, we asked participants to rate the subjective reachability for every view along the continuum.

Another possibility is that the transitional region reflects a canonical viewing distance along the continuum—that is, a view between the extremes which has some perceptual and representational privilege. This hypothesis draws on object representation research, where objects have canonical orientations and visual sizes at which observers prefer to view them, and these settings serve as anchoring points in visual memory experiments (Blanz et al., 1999; Konkle & Oliva, 2007, 2011; Palmer et al., 1981). Thus, it is possible that a similar phenomenon occurs in scene representation, where environments may be represented with respect to a canonical view. We explore this possibility by obtaining “goodness-of-view” judgements using a preference-judgment task similar to previous studies on canonical representations (Palmer et al., 2013).

Method

Participants. 213 participants were recruited for the Reachability task and 257 participants were recruited for the Goodness-of-View task (8 participants completed both tasks; 18 completed both reachability and the above memory distortion task; 23 completed both goodness-of-view and the memory distortion tasks. Twenty-three had also completed the Memory Distortion task; 8 had also completed Reachability). All participants were recruited on M-Turk, gave informed consent, and were compensated with \$0.25 per Human Intelligence Task (HIT). All procedures were approved by the Harvard University Human Subjects Institutional Review Board. The sample size of both tasks was matched to the Memory Experiment.

Experimental Design. In each trial, an image (500 x 667 pixels) was displayed, and participants made an untimed judgment. For the Reachability task, the instructions were to “Judge whether the view was *within reach* or *out of reach*”, and participants chose among those two options in a forced choice task. For the Goodness-of-view task, the instructions were to “Judge how good the current view is. Let’s say you are a photographer. Given this view, which would give you a better view?” Participants chose among three response options: taking one step *forward*, taking one step *backward*, and taking *no steps*. 45 ratings were obtained for each image, for both tasks. The same stimuli and counterbalancing method used in the Memory Experiment was used in both rating tasks. The presentation order of stimuli within each set was randomly determined for each HIT.

Data Pre-processing. Any HITs with average reaction times (RTs) falling outside of 3 standard deviations from the mean were removed. As above, all RTs were log transformed prior to trimming. This procedure was performed separately for Reachability dataset (1.3% excluded) and Goodness-of-view dataset (1% excluded).

Calculating reachability score and goodness-of-view score. For the Reachability task, “Out of Reach” responses were assigned a score of +1 and “Within Reach” responses were assigned -1. These scores were averaged to obtain a *reachability score* for each image. Here, a positive score indicates a reachable view, and a negative score indicates a non-reachable view, in a range of $[-1, 1]$. For the Goodness-of-view task, we assigned -1 for “Backward”, 0 for “No Move”, and +1 for “Forward”, and averaged the values for each image. A negative Goodness-of-view score indicates a preference to move backward, and a positive score indicates a preference to move forward, within a range of $[-1, 1]$.

Linear mixed-effects logistic regression. A linear mixed-effects logistic regression was used to fit the data from the reachability and goodness-of-view tasks, using lme4 (Bates et al., 2015) implemented in R (Team, 2016). For the Reachability dataset, the model included *Position* as a fixed-effect term and *Environment* as a random intercept and predicted a binary outcome (within-reach or out-of-reach). For the Goodness-of-view dataset, we excluded *no move* response trials for the logistic regression and modeled the data only with two responses (move forward or move backward); the model included *Position* as a fixed-effect term and *Environment* as a random intercept. The significance tests and estimation of the transition points were performed following the same methods as in the Memory Experiment.

Results

Using subjective reachability and Goodness-of-view judgments, we examined whether these behaviorally-relevant signatures had transition points in similar positions as the Memory Experiment. The Reachability ratings had a transition point closer to the near-scale views of the environment (positions 8.9 - 12.2 out of 30, mean=10.5, s.d.=1.6; Fig. 3B). Additionally, this reachability range overlapped with the transition between boundary extension and boundary contraction memory errors (positions 11.4–17.3). This correspondence is consistent with the idea that the memory distortions of images vary with depicted distance, and that the direction of distortions might depend on whether the space is within or out of reach.

However, the results of the Goodness-of-view ratings also showed a similar transition range. That is, participants chose to step backward for near-scale images but chose to step forward for far-scale images. The transition point between these response options, i.e., the preferred or canonical view, *also* overlapped with the transition point for memory distortions (canonical view 10.4-13.4 out of 30, mean=11.9, s.d.=1.5; Fig. 3C; see also Supplemental Material). Additionally, the ‘no move’ responses peaked around Position 10-11 (Supplemental Fig. 4A), further supporting the existence of a consistent canonical view in this region along the continuum¹. These results are thus also consistent with the possibility that the memory distortion effect is anchored around a canonical view of the environment.

¹Furthermore, in a pilot experiment limited to sampling Positions 1-20 (rather than 1-30), the preferred view was approximately the same (median at position 10.0).

Fig. 3. Transition Range Examples and Results from Reachability and Goodness-of-View Judgments. (A) Examples of the memory transition points from four different environments (see Supplemental Fig. 1 for all environment transition points). (B) The reachability judgment score averaged by each position, and an example image of the mean transition point. The negative score represents within-reach, the positive represents out-of-reach judgments, and the zero represents the transition point of subjective reachability. (C) The goodness-of-view judgment score averaged by each position, and an example image of the mean transition point. The negative score represents preference to step backward, the positive represents to step forward, and the zero represents the most preferred (“looks good”) view. For both (B) and (C), a vertical gradient bar indicates a transitional range (mean $\pm$ 1 s.d. across environments), and the horizontal bar on top of the plot shows its overlap with the transition range from memory distortions (light blue).

It is interesting that the point of subjective reachability and the preferred view both overlap with the transition range of the memory effect in these object-rich, naturalistic environments. However, we hypothesized that it may be possible to modify the virtual environments to dissociate Reachability and Goodness-of-view judgments. Intuitively, Reachability judgments should be relatively insensitive to the number and position of the surrounding objects in the environment, since such judgments are a factor of the length of the viewer’s arm, and the presence of a target relative to which they can measure distance. Goodness-of-view judgments rely more on aesthetic preferences, which may be relatively more sensitive to the overall distribution of image content, including objects (Leyssen et al., 2012). Thus, in the following experiments, we created “low-object environments” and repeated the Memory, Reachability, and Goodness-of-view experiments.

Low-Object Environments

A second stimulus set of visual environments was created from the original environments, but with all of the small and manipulable objects removed (e.g., cups or books), leaving only the immovable surfaces and furniture-like structures (e.g., kitchen countertops or large tables); see Fig. 4. Thus, these low-object environments shared the same spatial layout and background content as the original object-rich environments, but substantially differed in their object content. Next, images along the continuum were generated using the same parameters as the previous experiments, resulting in 600 unique stimuli (20 environments x 30 positions).

Fig. 4. Low-Object Environments. All of the small and manipulable objects were removed from the rich-object environments, leaving the immovable surfaces and furniture-like structure present such as kitchen countertops or cabinets. As results, two sets of environments shared the same spatial layout but substantially differed in their object content.

We completed the same three experiments reported above, but using the new low-object stimulus set. All experimental details were identical. Data was collected from $n=173$ participants in the Memory Experiment, $n=213$ participants in the Reachability task, and $n=153$ participants in the Goodness-of-view task. The data trimming procedures led to the exclusion of 2.8% of trials from the Memory Experiment dataset, 1.5% of trials from the Reachability dataset, and 0.7% of trials from the Goodness-of-view dataset. As these paradigms were identical to the first set of experiments, the sample size was determined based on a power analysis on the rich-object environments data. The results are shown in Fig. 5A.

In these low-object environments, the Reachability and Goodness-of-view judgments were quite different. Specifically, as expected, the Reachability ratings did not change between the low-object and rich-object environments (transition range = positions 11.4-14.9, bottom row in Fig. 5A). However, the Goodness-of-view ratings were shifted towards a farther view of the scene (transition range = positions 19.6-24.2, middle row in Fig. 5A). Example views of the Reachability and the Goodness-of-view transition are shown in Fig. 5B.

Critically, the memory distortions for these low-object stimuli again showed boundary contraction and boundary extension, but with a transition range shifted substantially further back (positions 17.7-27.1). This tracked closely with the Goodness-of-view transition range (mean position = 22 in both cases; paired t -test, $t(19)=0.55$, $p=0.6$). These data again indicate that the distance depicted in images affects the direction of memory distortions, but clarifies that the precise distance at which memory distortions transition from boundary extension to contraction is sensitive to the object content in the images. The point of transition was not linked to the perceived reachability of the scene, but was instead linked to properties of the image or environment that also manifest in preference judgements.

Additionally, the memory distortions in these low-environment scenes further replicate several patterns evident in the object-rich environments. That is, there is both boundary extension and boundary contraction along the continuum (63% boundary extension, 32.5% boundary

contraction, 4.5% no difference; Supplemental Fig. 3B), and the degree of distortion is systematically stronger with more distance from the transition point.

A Results (Low Object Environments)

B Transition Point Views

Fig. 5. Low-Object Environment Results and transition points. (A) The mean score by position from each paradigm. Positions around zero (y-axis) indicate the mean transition point, and a vertical gradient bar represents transitional region (mean $\pm$ 1 s.d. across environments). Compared to the rich-object environments (Fig. 2, 3), the transition region shifted much farther in the memory distortion and goodness-of-view judgements. However, the transition region of subjective reachability remained in similar positions regardless of the stimulus set. Critically, only the transitional region of goodness-of-view closely tracked that of memory distortion, supporting the canonical view account. (B) Examples of transition points from each paradigm (see Supplemental Fig. 2 for all environment transition points from memory distortions).

General Discussion

The goal of the current study was to probe memory distortions for scene views by varying the depicted scale of view from a full-scale indoor scene to close-up object views, covering a much more extensive and systematically-sampled range of distances than previous studies. We found (1) consistent boundary extension and boundary contraction; (2) the magnitude of the effect varied highly systematically with the depicted scale of the view; (3) the point in the environment at which there was no memory distortion was not linked to a specific viewing distance, like the point of

reachability; (4) instead, this point depended on the presence of objects and was linked to judgments of the view that “looked best.” Broadly, these results provide support for a framework in which boundary extension and boundary contraction are meaningfully linked phenomena, worthy of integrated consideration when examining how observers encode and remember the visual world.

The present work relates deeply to some of the earliest findings on boundary extension (Intraub & Richardson, 1989). In this original work, boundary extension was found only for close views, with no memory distortions for wider-angle views. A theory of asymmetric errors was envisioned to account for this discrepancy. However, we suggest an alternate interpretation of those early findings. Perhaps initial observations of boundary extension showed no distortion at the widest-angle view as a consequence of a restricted range of depicted depth in the stimulus sets; had a larger range of distances been tested originally, we would predict based on the current work that boundary contraction would have been found. Further, the current results cannot be easily explained away by a secondary normalization mechanism. In particular, the low- and rich-object environments showed different degrees of extension and contraction, even though the range of depicted depths was the same; a simple normalization to the range of depicted depth cannot account for these results. Instead, our results suggest that boundary extension and contraction are meaningfully linked phenomena. Considering these effects together allows new (or modified) theoretical frameworks to be explored.

One possible *prior-based* account draws heavily on the reconstructive memory literature, in which perceptual traces are encoded with respect to a structured prior distribution (e.g., Barlett, 1932; Hemmer & Steyvers, 2009b; Huttenlocher et al., 2000). In this Bayesian framework, systematic memory biases are signatures of optimal encoding of noisy perceptual input (e.g., Hemmer & Steyvers, 2009a). On such an account, each view of a scene is encoded with respect to some scene prior, which includes information about a canonical vantage point within the space, biasing memory toward that point and yielding parametric extension or contraction. Since judgments of the best-looking views are also thought to draw on this internal scene prior (Gardner & Palmer, 2010), this account can also explain why the Goodness-of-view ratings and memory errors closely tracked each other. However, this account alone does not articulate what scene properties are relevant for the prior; e.g., it does not make *a priori* predictions that the canonical view would shift back in low-object environments. Thus, this account requires further specification; currently it can likely account for any pattern of data.

A more perceptual, *dynamic-tension* account of these data—one that we favor—is that for any view, a balance of two oppositional forces determines whether subsequent memory will be extended or contracted. On this account, each view is encoded in parallel with scene-based and object-based mechanisms. The scene-based mechanism is sensitive to global scene properties, such as concavity, navigability, openness and mean depth (Bonner & Epstein, 2017; Cheng et al., 2019; Greene & Oliva, 2009; Park & Park, 2020), and is biased to construct a broader view, perhaps via the amodal scene construction originally proposed by Intraub (2012). At the same time, the object-based mechanism exerts pressure to process objects in more detail or with higher fidelity, perhaps via object-based attention (Beighley et al., 2019), resulting in down-weighting or loss of peripheral information.

On this *dynamic-tension* proposal, the view that “looks best” is one that balances these opposing affordances (rather than corresponding to the peak of an internal distribution). This account also naturally lends an explanation to some key patterns of our data: rich-object environments, with their prominent object content, may have closer transition points (and preferred views) to balance the needs of object recognition, while low-object environments have further transition points to favor spatial analysis of the environment. Indeed, the view that looks best for the low-object environments is just about where you start to see the side walls (Fig. 5), which provide a stronger sense of 3D boundary and concavity to a viewer. Overall, this account generally promotes thinking of the interplay between competing object-processing and spatial-layout processing mechanisms (Mullin & Steeves, 2013), e.g. predicting sensitivity to attentional manipulations and task demands, beyond static image-based information.

Finally, it is important to note that in order to obtain a smooth continuum between scene-centered and object-centered views, we varied the field-of-view (FOV) as well as distance, coming into a close crop of the object for the closest views to be compatible with previous studies. However, during naturalistic viewing, the vergence of the eyes on an object changes the human FOV far less dramatically than implied here, and visual content is not cut off with an image border. Future research could explore memory distortions in full-field viewing, allowing far-peripheral visual processing to be engaged in situating the content and operations of central vision. Given the importance of central and peripheral information for visual system organization (Arcaro & Livingstone, 2017; Hasson et al., 2003; Knapen, 2021), this may be a productive line of inquiry, moving the study of visual scene memory beyond the picture-plane into agent-directed movements through an environment.

References

- Arcaro, M. J., & Livingstone, M. S. (2017). Retinotopic organization of scene areas in macaque inferior temporal cortex. *Journal of Neuroscience*, 37(31), 7373–7389. <https://doi.org/10.1523/JNEUROSCI.0569-17.2017>
- 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. <https://doi.org/10.1016/j.cub.2019.12.004>
- Bainbridge, W. A., & Baker, C. I. (2020b). Reply to Intraub. *Current Biology*, 30(24), R1465–R1466. <https://doi.org/10.1016/j.cub.2020.10.032>
- Barlett, H. B. (1932). Occurrence and Properties of Crystalline Alumina in Silicate Melts. *Journal of the American Ceramic Society*, 15(7), 361–361. <https://doi.org/10.1111/j.1151-2916.1932.tb13944.x>
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Beighley, S., Sacco, G. R., Bauer, L., Hayes, A. M., & Intraub, H. (2019). Remembering: Does the emotional content of a photograph affect boundary extension? *Emotion*, 19(4), 699–714. <https://doi.org/10.1037/emo0000477>
- 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>
- Blanz, V., Tarr, M. J., & Bülthoff, H. H. (1999). What object attributes determine canonical views? *Perception*, 28(5), 575–597. <https://doi.org/10.1068/p2897>
- Bonner, M. F., & Epstein, R. A. (2017). Coding of navigational affordances in the human visual system. *Proceedings of the National Academy of Sciences of the United States of America*, 114(18), 4793–4798. <https://doi.org/10.1073/pnas.1618228114>
- Cheng, A., Walther, D. B., Park, S., & Dilks, D. D. (2019). What is a scene? Concavity as an intrinsic property of a scene. *Journal of Vision*. <https://doi.org/10.1167/19.10.128c>
- Cléry, J., Guipponi, O., Wardak, C., & Ben Hamed, S. (2015). Neuronal bases of peripersonal and extrapersonal spaces, their plasticity and their dynamics: Knowns and unknowns. *Neuropsychologia*, 70, 313–326. <https://doi.org/10.1016/j.neuropsychologia.2014.10.022>
- Cowey, A., Small, M., & Ellis, S. (1994). Left visuo-spatial neglect can be worse in far than in near space. *Neuropsychologia*, 32(9), 1059–1066. [https://doi.org/10.1016/0028-3932\(94\)90152-X](https://doi.org/10.1016/0028-3932(94)90152-X)
- Cowey, A., Small, M., & Ellis, S. (1998). No abrupt change in visual hemineglect from near to far space. *Neuropsychologia*, 37(1), 1–6. [https://doi.org/10.1016/S0028-3932\(98\)00071-2](https://doi.org/10.1016/S0028-3932(98)00071-2)
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743. <https://doi.org/10.1038/nn.3689>
- Gardner, J. S., & Palmer, S. E. (2010). Representational Fit in Position and Perspective: A Unified Aesthetic Account. *Journal of Vision*. <https://doi.org/10.1167/10.7.1232>
- Greene, M. R., & Oliva, A. (2009). Recognition of natural scenes from global properties: Seeing the forest without representing the trees. *Cognitive Psychology*, 58(2), 137–176. <https://doi.org/10.1016/j.cogpsych.2008.06.001>
- Halligan, P. W., & Marshall, J. C. (1991). Left neglect for near but not far space in man. *Nature*,

- 350(6318), 498–500. <https://doi.org/10.1038/350498a0>
- Hasson, U., Harel, M., Levy, I., & Malach, R. (2003). Large-scale mirror-symmetry organization of human occipito-temporal object areas. *Neuron*, 37(6), 1027–1041. [https://doi.org/10.1016/S0896-6273\(03\)00144-2](https://doi.org/10.1016/S0896-6273(03)00144-2)
- Hemmer, P., & Steyvers, M. (2009a). 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>
- Hemmer, P., & Steyvers, M. (2009b). Integrating episodic memories and prior knowledge at multiple levels of abstraction. *Psychonomic Bulletin and Review*, 16(1), 80–87. <https://doi.org/10.3758/PBR.16.1.80>
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129(2), 220–241. <https://doi.org/10.1037/0096-3445.129.2.220>
- Intraub, H. (2012). Rethinking visual scene perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(1), 117–127. <https://doi.org/10.1002/wcs.149>
- Intraub, H. (2020). Searching for boundary extension. *Current Biology*, 30(24), R1463–R1464. <https://doi.org/10.1016/j.cub.2020.10.031>
- Intraub, H., Bender, R. S., & Mangels, J. A. (1992). Looking at Pictures But Remembering Scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(1), 180–191. <https://doi.org/10.1037/0278-7393.18.1.180>
- Intraub, H., Gottesman, C. V., Willey, E. V., & Zuk, I. J. (1996). Boundary extension for briefly glimpsed photographs: Do common perceptual processes result in unexpected memory distortions? *Journal of Memory and Language*, 35(2), 118–134. <https://doi.org/10.1006/jmla.1996.0007>
- Intraub, H., & Richardson, M. (1989). Wide-Angle Memories of Close-Up Scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(2), 179–187. <https://doi.org/10.1037/0278-7393.15.2.179>
- Josephs, E. L., & Konkle, T. (2019). Perceptual dissociations among views of objects, scenes, and reachable spaces. *Journal of Experimental Psychology: Human Perception and Performance*, 45(6), 715–728. <https://doi.org/10.1037/xhp0000626>
- Josephs, E. L., & Konkle, T. (2020). Large-scale dissociations between views of objects, scenes, and reachable-scale environments in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 117(47), 29354–29362. <https://doi.org/10.1073/pnas.1912333117>
- Knapen, T. (2021). Topographic connectivity reveals task-dependent retinotopic processing throughout the human brain. *Proceedings of the National Academy of Sciences*, 118(2), e2017032118. <https://doi.org/10.1073/pnas.2017032118>
- Konkle, T., & Oliva, A. (2007). Normative Representation of Objects: Evidence for an Ecological Bias in Object Perception and Memory. *Proceedings of the Twenty-Nine Annual Conference of the Cognitive Science Society*, 407–412.
- Konkle, T., & Oliva, A. (2011). Canonical Visual Size for Real-World Objects. *Journal of Experimental Psychology: Human Perception and Performance*, 37(1), 23–37. <https://doi.org/10.1037/a0020413>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(13). <https://doi.org/10.18637/jss.v082.i13>
- Leyssen, M. H. R., Linsen, S., Sammartino, J., & Palmer, S. E. (2012). Aesthetic preference for

- spatial composition in multiobject pictures. *I-Perception*, 3(1), 25–49.
<https://doi.org/10.1068/i0458aap>
- McDunn, B. A., Brown, J. M., Hale, R. G., & Siddiqui, A. P. (2016). Disentangling boundary extension and normalization of view memory for scenes. *Visual Cognition*, 24(5–6), 356–368. <https://doi.org/10.1080/13506285.2016.1274810>
- Mullin, C. R., & Steeves, J. K. E. (2013). Consecutive TMS-fMRI reveals an inverse relationship in BOLD signal between object and scene processing. *Journal of Neuroscience*, 33(49), 19243–19249. <https://doi.org/10.1523/JNEUROSCI.2537-13.2013>
- Palmer, E. M., Horowitz, T. S., Torralba, A., & Wolfe, J. M. (2011). What Are the Shapes of Response Time Distributions in Visual Search? *Journal of Experimental Psychology: Human Perception and Performance*, 37(1), 58–71. <https://doi.org/10.1037/a0020747>
- Palmer, S. E., Schloss, K. B., & Sammartino, J. (2013). Visual aesthetics and human preference. *Annual Review of Psychology*, 64, 77–107. <https://doi.org/10.1146/annurev-psych-120710-100504>
- Palmer, S., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In *Attention and Performance IX*.
- Park, J., & Park, S. (2020). Coding of navigational distance and functional constraint of boundaries in the human scene-selective cortex. *Journal of Neuroscience*, 40(18), 3621–3630. <https://doi.org/10.1523/JNEUROSCI.1991-19.2020>
- Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*. <https://doi.org/10.1037/0033-2909.86.3.446>
- Strong, R., & Alvarez, G. (2019). *Using simulation and resampling to improve the statistical power and reproducibility of psychological research*. <https://doi.org/10.31234/osf.io/2bt6q>
- Team, R. C. (2016). R: A Language and Environment for Statistical Computing. In *R Foundation for Statistical Computing*.

Supplemental Material

Memory Distortion Transition Views across all rich-object environments

Supplemental Fig. 1. Memory Transition Views (rich-object environments). The transition point of each environment was estimated using the mixed-effects logistic regression.

Memory Distortion Transition Views across all low-object environments

Supplemental Fig. 2. Memory Transition Views (low-object environments). The transition point of each environment was estimated using the mixed-effects logistic regression.

Supplemental Fig. 3. Memory Distortion Scores by Each Image. (A) Left: rich-object environment stimuli showed a similar portion of boundary extension (44.7%) and contraction (51%). Right: there was a high correlation between the memory score and Position (Spearman rank correlation: $\rho=0.66$, $p=8.76 \times 10^{-77}$). (B) Left: low-object environment stimuli also showed a significant proportion of boundary contraction (32.5%) in addition to the boundary extension (63%). Right: there was a significant correlation between the memory score and Position (Spearman rank correlation: $\rho=0.45$, $p=9.96 \times 10^{-32}$).

Supplemental Fig. 4. Proportion of each response from Goodness-of-View Judgments. A height of each color bar represents the proportion of each response type, for each given position. Thus, the three color bars in each position sums up to 1. The proportion of No Move peaks around the estimated transition range then gradually decreases as it gets further from the transition range.

Bootstrap resampling method (power analysis)

Given the sample size of our experiments, we estimated the power of our experiments using a bootstrap resampling method (Strong & Alvarez, 2019). As all our analyses were performed after aggregating the data across individual M-Turk workers, we examined the sample size of each experimental condition, rather than the number of workers. The sample size of the first memory experiment (rich-object environments) was 900 trials per Position, or 45 trials per each Environment at each Position (20 environments per Position). Using a simulation approach, we examined how the power varies depending on the sample size. The specific steps of the simulation approach are as follows.

1) we set a range of sample sizes to test, from 2 to 65 trials. 2) we created a hypothetical dataset with a sampling with replacement method. For example, when testing a sample size of 10, we randomly sampled 10 trials out of 45 trials from each Environment at each Position; the same trial was allowed to be sampled more than one time. 3) we computed the mean score for each Position, by averaging across Environments in the hypothetical dataset. Then, a Spearman correlation was measured between the *mean score* and the *distance* (Pos1 to Pos30). The resulting *p*-value was recorded. 4) we repeated steps 2-3 for 1,000 times, which resulted in 1,000 *p*-values. We then computed how many *p*-values were under the significance criterion (0.05). This ratio was considered as Power of the testing sample size. 5) steps 2-4 were repeated for each sample size in the range set up at the step 1.

We tested both cases where we change the sample size of trials per Position or trials per Environment at each Position. Both showed >95% power with 40 trials per Position, indicating extremely high power of our design.

Transition points estimation for Goodness-of-view

We noticed a slight overestimation of transition points for Goodness-of-view dataset when using a logit as a link function, potentially due to the 3 alternative-forced-choice (AFC) paradigm instead of 2AFC (Brockhoff & Christensen, 2010). To address this, we confirmed the results with another approach. We estimated transition points by interpolating between two positions that has the last negative mean score and the first positive mean score, which gave the transition range of 8.5 – 12.6 Position (mean 10.6, s.d.=2) that overlaps with both Memory transition range and the Reachability transition region.

References

- Brockhoff, P. B., & Christensen, R. H. B. (2010). Thurstonian models for sensory discrimination tests as generalized linear models. *Food Quality and Preference*, 21(3), 330–338.
<https://doi.org/10.1016/j.foodqual.2009.04.003>
- Strong, R., & Alvarez, G. (2019). *Using simulation and resampling to improve the statistical power and reproducibility of psychological research*. <https://doi.org/10.31234/osf.io/2bt6q>