Boundary extension is constrained by naturalistic image properties

Marco Gandolfo, Hendrik Nägele, & Marius V. Peelen

Donders Institute for Brain, Cognition, and Behaviour,

Radboud University, 6525 HR, Nijmegen, NL

Keywords: Visual Perception; Boundary Extension; Scene memory; Natural Vision; Depth of Field; Visual Memory

Author Note

Correspondence concerning this article should be addressed to Marco Gandolfo; e-mail: marco.gandolfo@donders.ru.nl; Phone number: +3124361141

Preprint submitted to bioRxiv

Abstract

Boundary extension (BE) is a classical memory illusion in which observers remember more of a scene than was presented. According to predictive accounts, BE reflects the integration of visual input and expectations of what is beyond a scene's boundaries. Alternatively, according to normalization accounts, BE reflects one end of a normalization process towards a scene's typically-experienced viewing distance, such that BE and boundary contraction (BC) are equally common. Here, we show that BE and BC depend on depth-of-field (DOF), as determined by the aperture settings on a camera. Photographs with naturalistic DOF led to the strongest BE across a large stimulus set, while BC was primarily observed for unnaturalistic DOF. The relationship between DOF and BE was confirmed in three controlled experiments that isolated DOF from co-varying factors. In line with predictive accounts, we propose that BE is strongest for scene images that resemble day-to-day visual experience.

Statement of Relevance

In daily life, we experience a rich and continuous visual world in spite of the capacity limits of the visual system. Boundary extension (BE) is a memory illusion that sheds light on how observers compensate for such limits – that is, by filling-in the visual input with anticipatory representations of upcoming views, based on memory. However, not all images equally lead to BE. In this set of studies, we show that BE is strongest for images showing naturalistic depth-of-field, resembling human visual experience. Based on these findings, we propose that BE reflects a mechanism with adaptive value that is conditional to a scene being perceived as naturalistic. More generally, the strong reliance of a cognitive effect, such as BE, on naturalistic image properties indicates that it is imperative to use image sets that are ecologically-representative when studying the cognitive, computational, and neural mechanisms of natural vision.

Introduction

Boundary extension (BE) is a classical memory illusion in which observers remember more of a scene than was actually presented (Intraub et al., 1989). For example, when participants draw a previously observed scene, they typically include information not present in the photograph, expanding the scene beyond its boundaries (Intraub et al., 1989). Similarly, when rating the viewpoint of a second probe photograph relative to a remembered target photograph, participants rate the exact same view as closer than the original image, indicating that they remembered more of the scene than was presented (Bertamini et al., 2005; Gottesman & Intraub, 2002; Intraub & Dickinson, 2008; Dickinson and Intraub, 2008, Bainbridge & Baker, 2020).

Why does this illusion occur? One possibility is that it reflects a predictive mechanism, in which scene percepts do not correspond solely to the visual input but are rather *constructed* through integration of visual input and expectations of what is beyond a scene's boundaries (Intraub, 2010; Intraub & Dickinson, 2008; Maguire & Mullally, 2013). The discrete views sampled by the human eye may thus be supplemented with anticipatory representations of the scene, providing a continuous visual experience.

Alternatively, BE may reflect one end of a generic normalization process in memory (Bartlett, 1932), where the perceived scene viewpoint regresses towards a prototypical viewing distance: A view that is close-up compared to the observer's average viewing distance of that scene will be remembered as farther away (yielding BE), whereas wide-angle views will be remembered as closer-up, generating the opposite error, boundary contraction (BC). On this account, BE and BC are in principle equally common, only depending on whether the presented photographs are taken from a relatively close or far distance.

Previous work has shown that BE occurs more frequently than BC (Hubbard et al., 2010; Intraub, 2020). Furthermore, BE occurs for scenes but not for single, isolated objects (Gottesman & Intraub, 2002; Intraub et al., 1998), unlike normalization, which also occurs for isolated objects (Brady, et al., 2011; Konkle & Oliva, 2007; Intraub, et al., 1998). However, while these results appear to favor the predictive account, this conclusion relies on the strong assumption that the scene photographs used in previous work were representative of our daily-life visual experience, particularly with respect to viewing distance. If the photographs used in BE studies were not representative, for example by primarily showing views that were closer than usually experienced, results may similarly be explained by a normalization account.

Indeed, using a large and diverse stimulus set (1000 photographs), BE was recently shown to be equally common as BC (Bainbridge & Baker, 2020). Moreover, the direction of these memory errors was correlated with subjective image distance ratings: exactly as the normalization account would predict, subjectively close scenes tended to extend, and subjectively far scenes tended to contract. These results suggest that BE may be one side of a regression towards the typical scene-viewing distance.

The above discussion highlights that approximating daily-life visual experience is critical for both accounts: for the predictive account because it proposes that BE reflects expectations based on real-world experience; and for the normalization account because normalization is relative to our prior experience. If BE effects are explained by predictive

accounts, then BE should be observed (and be more common than BC) when using views representative of our visual experience. Importantly, this is not automatically achieved through large and diverse image sets. Even if these are representative of pictorial views of scenes (that is, of how scene photographs are taken), they are not necessarily representative of how an observer perceives their surroundings during day-to-day visual experience. A striking example is that a camera can produce pictures with a deeper depth-of-field (DOF) than the eye can (Middleton, 1958), leading to views that are potentially highly unnatural and never encountered in daily life.

Here, we show that BE and BC strongly depend on DOF, as determined by the aperture settings on a camera. Re-analyzing the 1000-image dataset (Bainbridge & Baker, 2020), we find that BC primarily occurred for photographs with unnaturalistic DOF. These photographs were characterized by an extremely deep DOF (i.e., everything is in focus), something that we cannot perceive in daily life due to physiological constraints of the visual system. Photographs with DOFs within the range of human vision overwhelmingly led to BE. Subsequently, in three experiments, we demonstrate the specific influence of DOF on BE while controlling for factors that may naturally co-vary with DOF in scene photographs, including viewing distance, number and type of objects, and image statistics. Again, BE was strongest for the shallow DOF that characterizes human vision. Altogether, our results are most in line with views of BE as a predictive mechanism, with this error being reliably observed for naturalistic views of scenes to which the brain has adapted.

Experiment 1

The goal of Experiment 1 was to test for a relationship between the perceived DOF of a photograph and its boundary transformation scores (BE or BC) in the large image set recently used by Bainbridge & Baker (2020). To this end, we collected aperture ratings (as judged from DOF) for each of the 1000 scenes, and related these ratings to the boundary transformation scores made available by the original study. Based on the predictive account, we hypothesized that the more an image would be rated to be shot with large aperture – resembling natural vision – the more strongly it would elicit boundary extension.

Methods

Participants

Twenty-six participants (22 females, mean age = 24.5 ± 5.84) signed up for the online scene rating experiment. Participants were recruited through the participant panel of Radboud University (SONA) in return of course credits and through social media. The experiment was advertised for people who knew the basics of photography. To further ensure participants' expertise, prior to the scene ratings, participants went through a photography quiz. First, participants were asked to define their photography experience with a multiple-choice question (i.e. How would you define your experience with photography? 1) Only took pictures with my smartphone; 2) Some knowledge, but no hands-on experience with Single-Lens-Reflex (SLR) cameras; 3) Some experience using SLR cameras; 4) (Semi-) Professional photographer). Second, participants had to complete a quiz with 5 multiple-choice questions. Four of these questions asked how a photograph would look like if shot with large aperture (i.e. shallow depth of field, lots of light entering the camera) and small aperture (i.e. deep depth of field, little light entering the camera). Finally, one question asked

whether a lens aperture of f/5.6 was larger or smaller than f/ 22.0, testing their knowledge of f-stop values. Based on the above questions, six participants were excluded because they either declared to have shot pictures only with their smartphone (2 participants) and/or failed 3 or more questions out of 5 in the photography quiz (4 participants). The final sample consisted of 20 participants. This and the following studies were approved by the Faculty's Ethics Committee (ECSW2017-2306).

Stimuli and apparatus

Stimuli were taken from Bainbridge and Baker (2020), corresponding to the 1000 images of their rapid serial visual presentation (RSVP) experiments. These were downloaded from the OSF link provided in their paper**.** Further, for the practice trials, 24 images were downloaded from Shutterdial (shutterdial – search images by camera settings). This website allows to search for photos according to their metadata (including aperture) from the Flickr API. We selected twelve pictures with large aperture *f* <= 4 and twelve pictures with small aperture *f* >= 22.0 so that the difference between shallow and deep depth of field was salient. This and the following experiments were online, coded in Javascript using the JSPsych toolbox 6.2.0 (de Leeuw, 2015). The data were saved on the Pavlovia.org servers (www.pavlovia.org).

Procedure

Following the photography quiz, participants were shown the task instructions. Here, the concept of aperture was briefly re-explained. Afterwards, each participant was presented with 6 randomly selected practice images (3 with large, and 3 with small aperture). In the practice, participants simply had to respond whether an image was shot with large or small aperture using the mouse, and received feedback after each trial. When the practice was completed, more detailed example images were presented, this time stressing the continuity of the aperture effect over the depth of field of the pictures. This was achieved by showing images of the same photograph shot at different levels of aperture (i.e. large – middle – small). Finally, participants were informed that the upcoming ratings would be performed on a slider from 0 (large) to 100 (small) aperture. Each rating trial began with a fixation cross with a random duration of 1, 1.25 or 1.5 seconds. Following fixation, participants saw an image appear on screen concurrently with the slider positioned in the middle (aperture = 50). To continue to the next trial, participants had to touch the slider on the scale and press the button "continue" placed below the image and the slider. Every 50 ratings, participants were invited to take a break.

Design

Each participant completed 1 of 4 versions of the experiment**.** Each version included 250 images from Bainbridge & Baker (2020), equally divided between the SUN database (Xiao et al., 2016) and the Google Open Images (Kuznetsova et al., 2020) database, where the images were originally from. The images were presented in random order to each participant. We obtained 5 aperture ratings for each image of the 1000 stimuli of Bainbridge & Baker (2020), matching the number of subjective distance ratings collected by Bainbridge & Baker (2020).

Analysis

All the analyses were conducted using R (R core team, version 4.0.1). We calculated the mean aperture rating for each of the 1000 images. We also split these ratings (from 0-100) in

5 categorical bins of 20 (XL 0-20; L 21-40; M 41-60; S 61-80; XS 81-100). Afterwards, we downloaded the raw data coming from the 1000-image RSVP Experiments 1 and 2 (Bainbridge & Baker, 2020). We re-scored the boundary transformation scores to -1 ("closer"), 0 ("same", only for Experiment 2), and 1 ("farther") and calculated the mean boundary transformation score per image. With these two measures we 1) ran the correlation between mean aperture ratings and the boundary transformation scores and 2) computed the average boundary rating across images in each categorical aperture rating bin (see Figure 1A-B).

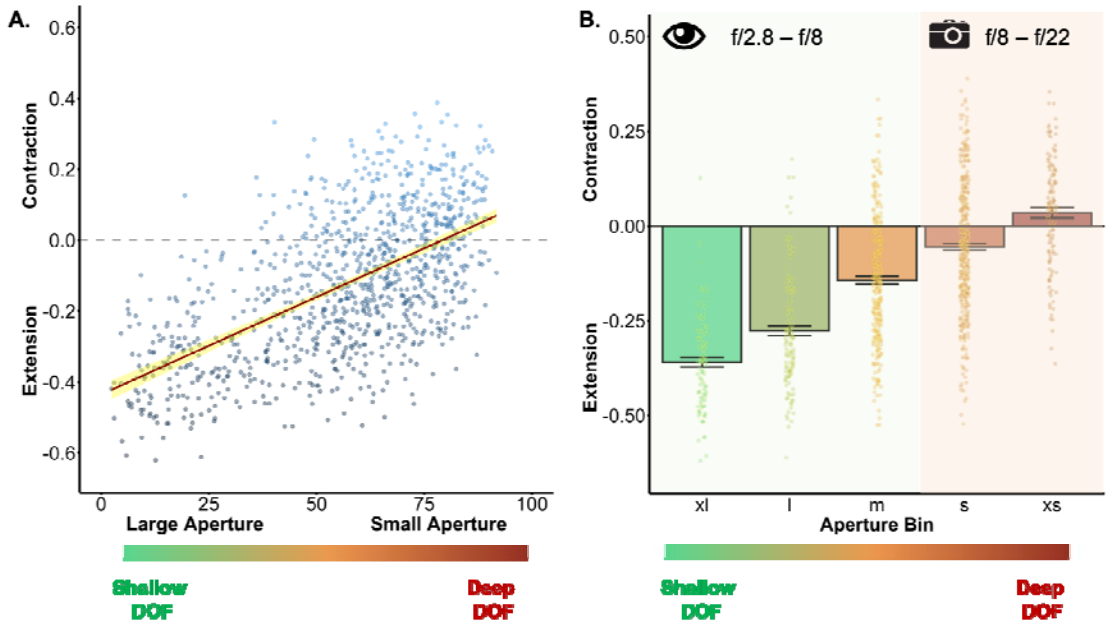

Figure 1. (A) Scatterplot with a linear regression line (in red), and 95% CI (in yellow). This shows the linear relation between aperture ratings and boundary transformation scores (combined across the two 1000 images experiments of Bainbridge and Baker (2020). Images rated to be shot with large aperture led to strong boundary extension. Images rated to be shot with smaller aperture led to less boundary extension or even contraction. **(B)** Bar plot averaging boundary transformation scores across images for each of 5 aperture bins. Images in the S and XS bins were judged by professional photographers to be shot with apertures between f/8 and f/22, as computed from comparisons with aperture judgements of pictures with known aperture levels (Supplementary Materials). These apertures are not ecologically representative. Data points here represent mean boundary transformation value per image. Error bars represent the SE of the mean.

Results

There was a strong and reliable correlation between rated aperture and boundary transformation score (Figure 1A, ρ = 0.58, $p < 0.001$). Images with a shallow DOF (large aperture) led to the largest extension effects, while images with deep DOF (small aperture) showed less extension or even contraction. To further inspect this relationship, we created 5 bins, each representing an interval of 20 on the aperture rating scale (XL – 0-20; L – 21-40; M – 41-60; S – 61-80; XS – 81-100). Compellingly, the bar plot in Figure 1B shows that contraction was reliable, on average, only for images rated to be shot with a small aperture, (t-test vs 0 on the values of the XS bin - t (121) =2.56, $p = 0.01$, d = 0.23, 95% CI = [0.01, 0.06]) – values of aperture higher than *f*/8, reachable by a camera but not by the human eye (Middleton, 1958; Hecht, 1987; see supplementary materials for how we estimated the *f*/stop value corresponding to each of the 5 bins shown in the figure).

Bainbridge & Baker (2020) showed that subjective ratings of viewing distance were correlated with boundary transformation scores, in line with other studies (Intraub et al., 1992; Hafri et al., 2021; Park et al., 2021). Images rated to have a close view led to larger extension than images rated to have a far-away view. It is possible that the relationship between rated DOF and boundary scores observed here is entirely absorbed by the relationship between subjective distance and boundary scores. To address this, we partialed out ratings of subjective distance. The correlation between aperture and boundary scores was still reliable ($\rho = 0.29$, $p < 0.001$), even when the shared variance between DOF ratings and subjective distance ratings was regressed out.

Experiment 2

In Experiment 1 we showed that boundary transformation scores and perceived DOF were related to each other across a large and diverse set of photographs, thus revealing that a previously unexplored image property (aperture) affects memory for scene boundaries. However, similar to viewing distance, the DOF of an image may covary with numerous other visual and semantic properties. For example, images with deep DOF are more likely to show an outdoor environment. The correlational approach used in Experiment 1 is thus unable to unambiguously ascribe the effects to DOF (or, indeed, viewing distance). Therefore, to firmly establish an effect of DOF on BE, we need to manipulate DOF while keeping all other properties equal. To do so, we shot the same 30 pictures of outdoor environments with three different camera apertures. This factor was crossed with the camera's focus distance, focusing either on a nearby or a far-away object in the scene (Figure 2). In Experiment 2, we used these images in the same paradigm as used in Experiment 1 (Bainbridge & Baker; 2020). If the size of memory distortions depends on how the DOF of an image resembles natural viewing conditions, then images with shallow DOF should show larger extension than images with deep DOF, independently of focus distance.

Methods

Participants

Seventy-eight participants signed up for the online experiment (65 females, mean age = 20.62 ± 4.19 years). Participants were students recruited via Radboud University's participant panel (SONA) in return for course credits. Because we had no a-priori hypothesis on the size of the effect this sample size was determined by the student availability on the participant panel of the University in a period of four weeks.

Stimuli and apparatus

The main stimulus set consisted of 6 photographs of 30 unique outdoor scenes (180 images). Each scene featured one or more objects in the foreground at a distance of approximately two meters, and one or more objects in the background at a distance of at least 5 meters. For each scene, 6 photographs were taken using a DSLR (digital single-lens reflex camera) with a 50mm fixed focal length lens placed on a tripod. Photos were shot at 3 aperture levels (*f*/5.6 (large); *f*/11.0 (medium); *f*/22 (small)), each one either manually placing focus on the foreground object(s) or on the background object(s) (See Figure 3). Further, the ISO (light sensitivity of the sensor) was set manually in each scene (but kept constant across aperture and focus variations of that scene) and the shutter speed was automated to

achieve similar lighting for the 6 photographs. Three additional scenes, including photographs for each of the aperture and focus conditions, were used as catch trials. For these pictures, we generated a zoomed-in version enlarging wider-angle views to 24% of the surface area of the image (in line with the procedure of Intraub and Dickinson (2008)) and then resized them to their original size. All pictures were then grayscaled and resized to 750 x 500 pixels. In total, the stimulus set included 216 photographs (180 for the test trials, and 36 for the catch trials). Further, 8 different masks were generated by computing the average image of all the photos in the stimulus set and iterating image scrambling in 20x20 pixel blocks (randomly scrambling on the x and then on the y axis, or vice versa). For this and the following experiments, image processing was performed using the R package “Imager” (Barthelme & Tschumperle, 2019).

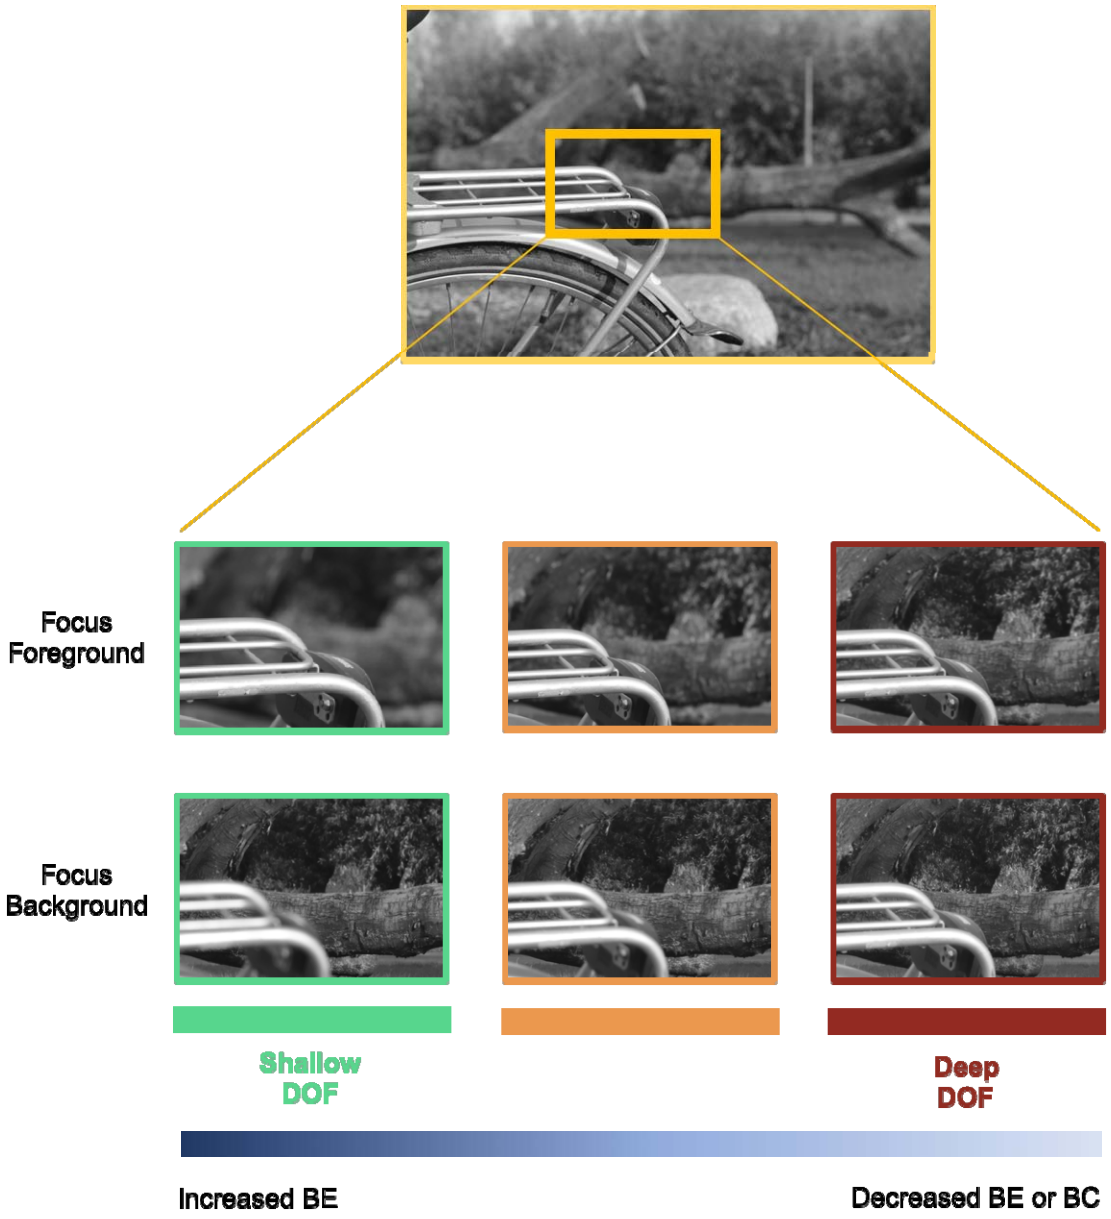

Figure 2. Example of the experimental conditions showed on a zoomed portion of one of the scene photographs. Each outdoor scene image was shot focusing either on the foreground object or on the background, at 3 levels of aperture – f/5.6 (shallow DOF), f/11, f/22 (deep DOF). As a result, the number of elements in focus (either in the background or in the foreground) increased as a function of DOF. The full stimulus set is publicly available to download at the OSF repository (Link), including the images equated for spatial frequency and histogram used in Experiment 4.

Procedure

The experiment mimicked the procedures of the 1000-image RSVP Experiment 2 of Bainbridge and Baker (2020), a paradigm commonly used to measure BE effects (Park & Konkle, 2021, Intraub & Dickinson, 2008; Figure 3A). In each trial, participants saw a scene presented for 250 ms, followed by a 250 ms dynamic mask (5 masks randomly selected from the 8 available masks for each participant, presented for 50 ms each), and then the same scene was shown again. After 1 second, a response screen appeared for 3 seconds or until response. The response screen displayed text asking whether the second image was closer (scored as -1, “J” key), the same (scored as 0, “k” key) or farther (scored as +1, “l” key) than the first image.

In the instructions, participants were informed that the study concerned scene memory. They were also informed that the picture to be judged could have been either taken from the same, a closer-up, or a farther away view. Importantly, in the trials of interest (90% of the total number of trials) the image to be judged was the same as the first image. In the catch trials (10% of the total number of trials), the second image was either an obviously zoomed-in or a zoomed-out version of the first photo. The catch-trials allowed to 1) have attentive participants to reinforce their belief that there were actual changes in the images and 2) have an accuracy measure that could be taken as an index of data quality (i.e. a suitable attention check for this task). Participants were invited to take a break at the end of each block (every 33 trials).

Design

The experiment consisted of 6 blocks of 33 trials. In each block, we ensured that the same scene was not presented more than once. In every block, there were 6 scene images for each of the experimental conditions (i.e. 2 focus levels, foreground/background x 3 apertures, large/medium/small). Further, there were three catch trials (with the second image being actually zoomed in or out). These were selected randomly in each participant from a pool of 36 possible trials (i.e. 2 focus levels x 3 apertures x 2 zoom types x 3 photos), making sure that no catch trial image was repeated within a block. The full design thus included 198 trials, 180 experimental trials and 18 catch trials.

Analyses

Based on the catch-trials performance we excluded participants who were below or at chance accuracy in the catch trials (mean accuracy <= 33%), the participants that responded more than 90% of trials “same” (excluding the catch trials), and the participants who did not respond (letting the trial to time out) in more than 50% of the total number of experimental trials. Based on these criteria we excluded 12 participants (6 for low-catch trial performance; 5 for not responding in more than 50% trials; 1 participant for responding “same” in > 90% of the experimental trials). The analyses thus included data from 66 participants. We then

scored “closer” responses to -1, “same” responses to 0 and “farther” responses to 1. Finally, we computed the mean responses per participant in each condition and ran a 3 (aperture levels) x 2 (focus on foreground/background) repeated measures ANOVA.

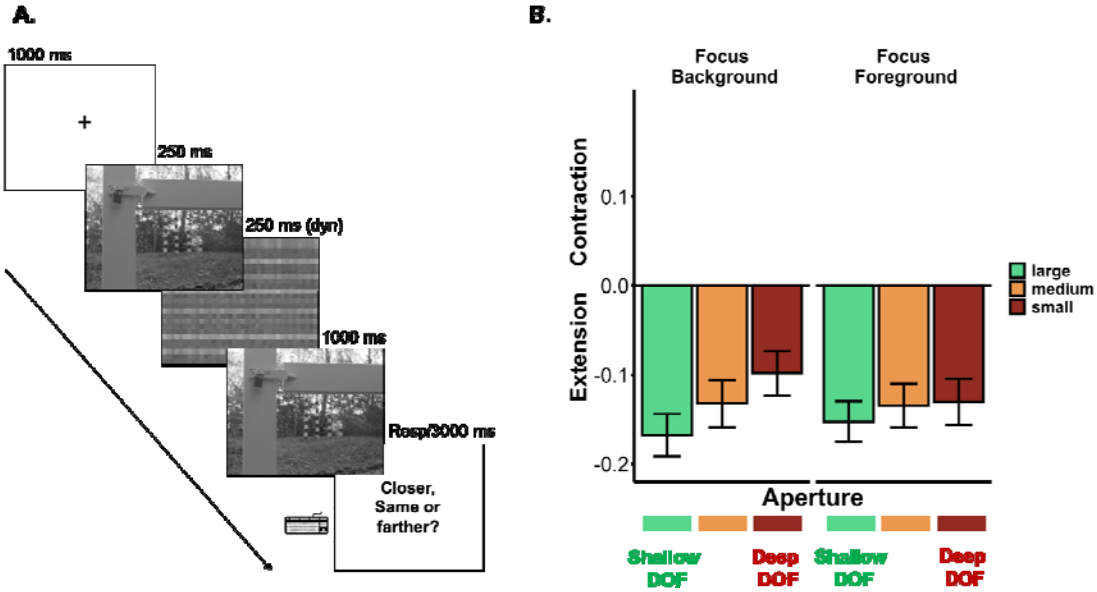

Figure 3. (A) Graphical representation of Experiment 1. The procedure matches Bainbridge and Baker’s (2020) 1000 images Experiment 2, thus including the “same” response option. **(B)** Bar plot representing the main effect of aperture in Experiment 1, F(2,130) = 5.27, *p* = 0.006, occurring across both focus distance levels. Images with large aperture led to the largest extension. Error bars represent the SE of the mean.

Results

Confirming our hypothesis, we found a main effect of lens aperture, F (2, 130) = 5.27, *p* = 0.006, peta^2 = 0.08), showing stronger BE for images shot with large aperture, f/5.6 (M = -0.16, SD = 0.18) than images shot with medium aperture, f/11 (M = -0.13, SD = 0.19; t(65) = -2.31, p= 0.02, d = -0.28, 95% CI = [-0.53, -0.04]) and images shot with small aperture, f/22 (M = -0.11, SD = 0.18; t(65) = -2.79, p= 0.007, d = -0.34, 95% CI = [-0.59, -0.09]). There was no significant difference between medium and small aperture stop (t(65) = -1.27, p= 0.21, d = -0.16, 95% CI = [-0.05, 0.01]). Thus, while boundary extension was found in all conditions, it was reduced when images showed deeper DOF (Figure 3B). The effect of aperture did not interact with focus distance (F (2, 130) = 0.64, *p* = 0.53) and there was no main effect of focus distance (F(1,65) = 0.99, *p* = 0.32).

Experiments 3 and 4

Experiment 2 measured BE by asking participants whether a second image *looks* closer or farther than a previously observed image. In line with previous work, our findings show that participants tend to perceive the second image as closer, even though the two images are identical (Bainbridge and Baker, 2020; Park et al., 2021; Intraub and Dickinson, 2008; Hale et al., 2016). While this paradigm has proven effective in revealing BE, the ratings reflect a subjective impression rather than objective performance (Liu et al., 2016). To corroborate and generalize the findings of Experiment 2, in Experiments 3-4 we developed a new paradigm that quantified BE using a change detection task. In this task (Figure 4A), all trials

showed an actual change of view from the first to the second image, either changing from a wide to a close view or vice versa. Participants indicated whether the second image was closer or farther than the first image in a 2-alternative forced choice task and rated their confidence on a slider scale. If the boundaries of the first image are extended in memory, participants should be worse in detecting changes from close to wide (vs. changes from wide to close) because the first (close) image is remembered to be wider than it was. This new task gives us another opportunity to test for the effects of aperture and focus distance on boundary extension.

Methods

Participants

For experiments 3 and 4, the sample size was determined through a power analysis based on the data of Experiment 2. This analysis revealed that a sample size of 35 participants was required to obtain 80% power to detect a difference between large and small aperture conditions at least as large as the one observed in Experiment 2, on the basis of a paired samples t contrast. In Experiment 3, thirty-nine participants were recruited for the experiment (16 females, mean age = 24.66 ± 4.10) to arrive at a final sample of N = 35 after exclusions based on missing data or low accuracy (see Analysis). In Experiment 4, thirty-six participants were recruited for the experiment (11 females, mean age = 24.36 ± 4.66). Due to server issues, out of 36 we obtained only 32 complete datafiles in experiment 4 (see Analysis). Participants were recruited via Prolific (www.prolific.co) in return for monetary compensation at an hourly rate of 6.31 £/hr.

Stimuli and apparatus

In Experiments 3 and 4 the stimuli were the same as those used in Experiment 2, with the exception that two of the images that served as catch-trials were now used as part of the image set. Further, to simplify the design, in these experiments we did not use the medium (*f* 11.0) aperture level. Additionally, we generated “close view” scenes through resizing (without resampling) the original pictures. These were enlarged using a random percentage value between 17% to 24% (in steps of 0.5%) of the surface area in each scene, making sure that a given scene would be enlarged by the very same percentage for each aperture and focus distance level. Finally, in Experiment 4 we matched the stimuli for low level properties. Specifically, we used the SHINE toolbox (Willenbockel et al., 2010) running on Octave 5.1.0 (Eaton et al., 1997). Using the functions *sfMatch* and *HistMatch* we matched the spatial frequency and histograms of each scene across its focus and aperture levels. Therefore, for each scene four images were matched for their low-level properties – the small and large aperture images in both near and far focus levels. This procedure ensured that any difference between conditions in the experiment could not be due to covarying low-level properties emerging from the changes in focus distance and lens aperture. For these experiments, there were 128 images – 32 for each combination of aperture and focus distance. Further, we generated additional 128 close view images belonging to the same conditions through resizing. Finally, the same 8 masks of Experiment 2 were used.

Procedure

In each trial (Figure 4A), after a fixation cross with a random duration between 1 and 2 seconds, participants viewed a scene for 250 ms, followed by a 1s dynamic mask (changing every 250ms, 5 masks randomly selected from the 8 available masks in each trial), and then

a closer or wider version of the same scene was shown again for 250 ms. Participants were asked to respond as fast and as accurately as possible pressing the “A” key if the second image looked closer or the “S” key if the image looked farther compared to the first image. Responses were given with the index (S) and middle (A) finger of the left hand. Right after response, or after 3.5 seconds, participants had to rate how sure they were of their response on a slider scale from 0 (“Not Sure”) to 100 (“Sure”). In the instructions, participants were told that the study concerned scene memory. Further, they were shown a self-paced example of the trial and performed a 10-trial practice block in which they were given feedback about their performance and got used to the task speed. Every 32 trials there was a self-paced break. The above procedure applies to both Experiments 3 and 4.

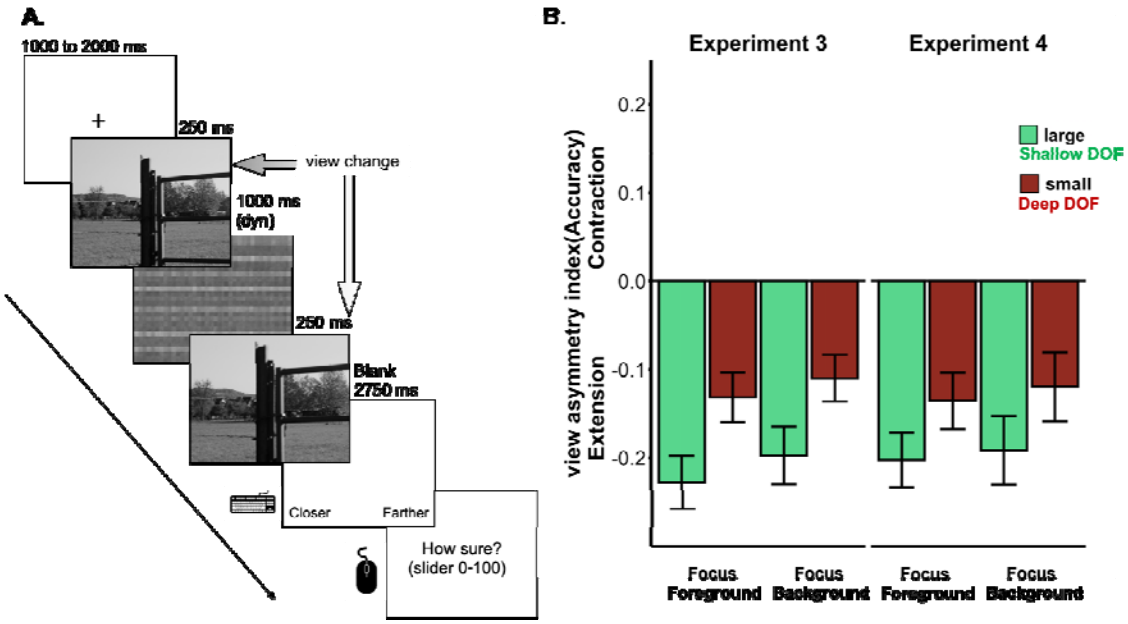

Figure 4. **(A)** Graphical representation of the procedure of Experiments 3 and 4. The example shows a wide to close trial of experiment 3. The probe picture shows a closer view than the first picture. **(B)** Results from Experiments 3 and 4. The View asymmetry index is computed as the difference in change detection accuracy for view changes (close-to-wide – wide-to-close). Negative values represent higher accuracy in detecting a change from wide to close, indicative of boundary extension. In both experiments there was an effect of apertures across both levels of focus distances. Error bars represent the SE of the mean.

Design

The experiment consisted of 8 blocks of 32 trials. In each block, we ensured that the same scene was not presented more than once. In every block, there were 4 scene images for each of the experimental conditions (i.e. 2 views, close to wide/wide to close; 2 focus levels, far/near; 2 apertures, large/small). The full design thus included 256 trials in total.

Analyses

In both experiments we excluded participants who did not have a complete datafile (1 participant for Experiment 3; 4 for Experiment 4) or who did not respond on more than 50% of trials (1 participant in Experiment 3). Further, participants were considered outliers when their response time (RT) or accuracy was 2.5 SDs above (RT) or below (accuracy) the group mean across conditions (2 participants in Experiment 3; 0 in Experiment 4). The analyses for Experiment 3 thus included data from 35 participants, while those for Experiment 4 included

data from 32 participants. Finally, for each experiment, we ran a 2 (View, close-to-wide/ wide-to-close view) x 2 (Aperture, large/small) x 2 (focus on foreground/background) repeated measures ANOVA on accuracy – correct detection of the second picture being closer or wider. For the purpose of visualization, we then computed a "view asymmetry index" subtracting accuracy on wide-to-close trials from accuracy on close-to-wide trials (Figure 4B). For this index, positive values indicate boundary contraction while negative values indicate boundary extension.

Results

In this paradigm, a better performance in detecting wide to close views than close to wide views is indicative of BE, the converse is indicative of BC. In both experiments, we found a main effect of view (experiment 3 - $F(1,34) = 44.20$, $p < 0.001$, $peta^2 = 0.57$; and experiment 4 - $F(1,31) = 26.69$, $p < 0.001$, $peta^2 = 0.46$). View changes from wide to close (Experiment 3 - M = 0.80, SE = 0.02; Experiment 4 - M = 0.81, SE = 0.02) were detected more accurately than view changes from close to wide (Experiment 3 - M = 0.63, SE = 0.02; Experiment 4 - M = 0.65, SE = 0.02). This effect indicates that boundaries were extended in memory across all conditions.

Most importantly, in both experiments we also found a significant View x Aperture interaction (Experiment 3 - $F (1, 34) = 35.67$, $p < 0.001$, $peta^2 = 0.51$; Experiment 4 - $F(1,31) = 26.69$, $p < 0.001$, $peta^2 = 0.46$). Photographs with large apertures led to larger BE than pictures with small apertures (Experiment 3 - $t(1,34) = 5.97$, $p < 0.001$, d = 1.01, 95% CI = [0.06, 0.12]; Experiment 4 - $t(1,31) = 3.38$, $p = 0.002$, d = 0.60, 95% CI = [0.03, 0.11]). View change did not interact with focus distance (Figure 4B, Experiment 3 - $F(1,34) = 2.24$, $p = 0.14$; Experiment 4 – $F(1,31) = 0.48$, $p = 0.49$), and there was no three way interaction (Experiment 3 – $F(1,34) = 0.05$, $p = 0.82$; Experiment 4 – $F(1,31) = 0.02$, $p = 0.88$). A similar pattern of results was observed for inverse efficiency scores (RT/% correct), confirming that results were not due to speed accuracy trade-offs (Supplementary Materials). Finally, the confidence ratings mirrored the accuracy results (Supplementary Materials).

General Discussion

In four experiments, we found that BE depends on a photograph's DOF. Across a large and variable image set, DOF was a strong predictor of BE, with BE being largest for images with shallow DOF, resembling human vision. By contrast, BC was only reliably observed for images with deep, unnaturalistic DOF. Three controlled experiments showed that DOF modulated BE even when keeping other properties of the scene constant across conditions. Altogether, these results demonstrate that BE is reliably observed for ecologically-representative stimuli.

Because DOF may covary with perceived distance across photographs, we ensured that the relationship between DOF and BE observed here could not be explained by differences in perceived distance. First, in Experiment 1, the relationship between DOF and BE remained reliable after regressing out subjective distance ratings. Second, in Experiments 2-4, focus distance did not significantly affect BE and did not interact with DOF, even though for large aperture focus distance was clearly visible (Figure 2). Therefore, DOF affects BE independently of perceived distance. However, previous work has shown that BE is larger for scenes that show closer compared to wider views (Intraub et al., 1989; Intraub et

al., 1992; Intraub & Richardson, 2008; Bertamini et al., 2005; Bainbridge & Baker, 2020; Park et al., 2021; Hafri et al., 2021). These results may suggest an additional role for normalization processes. This more generic normalization process (also observed for objects) could exist independently of the more scene-specific predictive mechanism driving BE (Intraub et al., 1996; 1998).

When integrating our results with previous findings, we see the largest BE for photographs with a shallow DOF, taken from close distance, with the main object seen from a central vantage point (Gagnier et al., 2011). What do all these aspects have in common? One possibility is that images shot under these conditions best resemble how we naturally sample the visual world. Indeed, images with these characteristics show a more partial view of the scene than images with deep DOF, wide view and lateral vantage point. These partial views are more likely to require integration of visual input: When the image is less complete, observers may need to rely more on top-down expectations of scene layout, drawing on sources other than the visual input, to complete the percept.

Based on our results, we propose that BE reflects a constructive mechanism with adaptive value that is *conditional* to a scene being perceived as naturalistic. An image shown at a plausible distance from the observer, with a DOF resembling the day-to-day perceptual experience will likely lead to extrapolation of scene layout from memory, and therefore BE. Where one or multiple of these properties are removed then BE is attenuated, and at the extremes, contraction may occur. Future research is needed to identify those conditions of the external visual input that allow the mental representation of scenes, stored in memory, to complete external percepts.

Finally, our results have broad implications for disciplines aiming to study natural vision. Specifically, many recent psychological, neuroscientific, and computer vision studies implement the use of large stimuli sets to achieve higher ecological validity (e.g., Hebart et al., 2020; Bainbridge & Baker, 2020; Chang et al., 2020; Mehrer et al., 2021). The current findings indicate that such large image sets are not necessarily representative of how observers perceive their visuo-spatial world. By showing that DOF can drastically change the influence of top-down knowledge on visual processing, our results imply that it is imperative to use images with naturalistic DOF in future work.

Conclusion

To conclude, our findings indicate that the DOF of an image strongly affects boundary errors for scenes, independently of other image properties. Further, BE is much more common than BC when considering photographs with naturalistic DOF, resembling our visual experience. As such, our findings are most in line with predictive accounts of BE, in which visual input is integrated with top-down expectations. These predictive mechanisms operate most strongly when the visual input is representative of the visuo-spatial world to which the human visual system has adapted.

References

Bainbridge, W. A., & Baker, C. I. (2020). Boundaries extend and contract in scene memory depending on image properties. *Current Biology*, *30*(3), 537-543.

Barthelmé, S., & Tschumperlé, D. (2019). imager: an R package for image processing based on CImg. *Journal of Open Source Software*, *4*(38), 1012.

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.

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of vision*, *11*(5), 4-4.

Chang, N., Pyles, J. A., Marcus, A., Gupta, A., Tarr, M. J., & Aminoff, E. M. (2019). BOLD5000, a public fMRI dataset while viewing 5000 visual images. *Scientific data*, *6*(1), 1-18.

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods*, *47*(1), 1-12.

Dickinson, C. A., & Intraub, H. (2008). Transsaccadic representation of layout: What is the time course of boundary extension?. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(3), 543.

Eaton, J. W., Bateman, D., & Hauberg, S. (1997). *Gnu octave* (p. 42). London: Network theory.

Gagnier, K. M., Intraub, H., Oliva, A., & Wolfe, J. M. (2011). Why does vantage point affect boundary extension?. *Visual Cognition*, *19*(2), 234-257.

Gottesman, C. V., & Intraub, H. (2002). Surface construal and the mental representation of scenes. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(3), 589.

Hafri, A., Wadhwa, S., & Bonner, M. F. (2021, September 7). "Honey, I shrunk the scene": Decoupling the impact of distance from content on memory for scene boundaries. https://doi.org/10.31234/osf.io/hy3qs

Hale, R. G., Brown, J. M., & McDunn, B. A. (2016). Increasing task demand by obstructing object recognition increases boundary extension. *Psychonomic bulletin & review*, *23*(5), 1497-1503.

Hebart, M. N., Zheng, C. Y., Pereira, F., & Baker, C. I. (2020). Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature human behaviour*, *4*(11), 1173-1185.

Hecht, E. (1987). Optics 2nd edition. *Optics 2nd edition by Eugene Hecht Reading*.

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: A multisource model. In *Psychology of learning and motivation* (Vol. 52, pp. 231-264). Academic Press.

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.

Intraub, H., & Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 179.

Intraub, H., Gottesman, C. V., & Bills, A. J. (1998). Effects of perceiving and imagining scenes on memory for pictures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(1), 186.

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.

Intraub, H. (2020). Searching for boundary extension. *Current Biology, 30*(24), R1463-R1464.

Konkle, T., & Oliva, A. (2007). Normative representation of objects: Evidence for an ecological bias in object perception and memory. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 29, No. 29).

Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., ... & Ferrari, V. (2020). The open images dataset v4. *International Journal of Computer Vision*, *128*(7), 1956-1981.

Liu, Z., Yang, X., & Intraub, H. (2016). Boundary extension: Insights from signal detection theory. *Journal of vision*, *16*(8), 7-7.

Maguire, E. A., & Mullally, S. L. (2013). The hippocampus: a manifesto for change. *Journal of Experimental Psychology: General*, *142*(4), 1180.

Mehrer, J., Spoerer, C. J., Jones, E. C., Kriegeskorte, N., & Kietzmann, T. C. (2021). An ecologically motivated image dataset for deep learning yields better models of human vision. *Proceedings of the National Academy of Sciences*, *118*(8).

Middleton, W. E. K. (1957). Vision through the atmosphere. In *Geophysik II/Geophysics II* (pp. 254-287). Springer, Berlin, Heidelberg.

Park, J., Josephs, E. L., & Konkle, T. (2021, January 26). Systematic transition from boundary extension to contraction along an object-to-scene continuum. https://doi.org/10.31234/osf.io/84exs

R Core Team (2021). R: A language and environment for statistical computing. R Foundation for statistical computing, Vienna, Austria. URL https://www.R-project.org/.

Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: the SHINE toolbox. *Behavior research methods*, *42*(3), 671-684.

Xiao, J., Ehinger, K. A., Hays, J., Torralba, A., & Oliva, A. (2016). Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, *119*(1), 3-22.