

How undistorted spatial memories can produce distorted responses

Mark W. Schurgin · Jonathan I. Flombaum

Published online: 18 March 2014
© Psychonomic Society, Inc. 2014

Abstract Reproducing the location of an object from the contents of spatial working memory requires the translation of a noisy representation into an action at a single location—for instance, a mouse click or a mark with a writing utensil. In many studies, these kinds of actions result in biased responses that suggest distortions in spatial working memory. We sought to investigate the possibility of one mechanism by which distortions could arise, involving an interaction between undistorted memories and nonuniformities in attention. Specifically, the resolution of attention is finer below than above fixation, which led us to predict that bias could arise if participants tend to respond in locations below as opposed to above fixation. In Experiment 1 we found such a bias to respond below the true position of an object. Experiment 2 demonstrated with eye-tracking that fixations during response were unbiased and centered on the remembered object's true position. Experiment 3 further evidenced a dependency on attention relative to fixation, by shifting the effect horizontally when participants were required to tilt their heads. Together, these results highlight the complex pathway involved in translating probabilistic memories into discrete actions, and they present a new attentional mechanism by which undistorted spatial memories can lead to distorted reproduction responses.

Keywords Attention and memory · Eye movements · Visual attention · Spatial memory

Representations of space and position support cognition in a variety of domains and act as scaffolding for perceptual

representations in general. In the effort to characterize the form and content of spatial representations, spatial distortions—apparent misrepresentations of space and spatial relations—have played a key role. Understanding why and how distortions arise can supply a window into the underlying structure of mental space in many contexts.

The variety of known spatial distortions and their explanations is too large to review exhaustively. But a few cases are worth considering for insight into how the study of distortion reveals broader principles of cognition. For example, reproducing maps from memory results in systematic distortions to the sizes and alignments of territories, revealing heuristic and categorical features of map memory (Tversky, 1981). Classic distortions arising in illusions, including the Müller-Lyer illusion, reveal the cues employed for uncovering the three-dimensional structure of the world, given two-dimensional inputs (Marr, 1982). In addition, object positions can be systematically reproduced from memory as attracting one another (Sheth & Shimojo, 2001), repelling one another (Liverence & Scholl, 2011), and attracting or repelling landmarks and boundaries (Diedrichsen, Werner, Schmidt, & Trommershäuser, 2004; Huttenlocher, Hedges, & Duncan, 1991; Nelson & Chaiklin, 1980), revealing features of relational processing in memory, perception, and attention.

Surprisingly, even a single remembered object in a relatively sparse context is, at times, distorted systematically during reproduction. For example, Sheth and Shimojo (2001) asked participants to localize a single dot from memory, finding systematic attraction toward a fixation point (fixation was maintained during both presentation and response). Similarly, an influential study by Huttenlocher et al. (1991) showed systematic distortions in the reproduction of a single dot, dependent on the dot's position within a circular boundary. For example, reproductions of dots near but not abutting a circular boundary tended to be distorted away from the boundary.

M. W. Schurgin (✉) · J. I. Flombaum (✉)
Department of Psychological and Brain Sciences, Johns Hopkins University, Ames Hall, 3400 N. Charles Street, Baltimore, MD 21218, USA
e-mail: maschurgin@jhu.edu
e-mail: flombaum@jhu.edu

Importantly, their interpretation of these results, and their subsequent theory, are different from prevailing interpretations in related studies. Whether implicitly or explicitly, most studies endorse a view suggesting that emergent spatial distortions arise from distorted representations—usually distorted memories, but not exclusively (e.g., Liverence & Scholl, 2011; Sheth & Shimojo, 2001). In contrast, Huttenlocher et al. (1991) endorsed a theory of multiple *undistorted* representational formats, which must be reconciled in order to produce a single response. In their words, “item location is coded at two levels of detail, each of which is unbiased although inexact, and that bias arises in combining information from the two levels to produce an estimate” (p. 354). In the example from above, of reproduction within and near a circular boundary, the idea is that a probabilistic estimate of a dot’s coordinates contributes to reproduction along with a categorical representation (something like “near, but not touching the boundary”). The probabilistic representation may include positions very near or even touching the boundary, but the categorical representation leads to biased sampling and neglect of those positions. Biased sampling in turn *appears* as distorted memory. Thus, the main point made by Huttenlocher and colleagues is that truncated or biased sampling from unbiased distributions can lead to biased responses.

We were motivated to build upon this theory, expanding the scope of mechanisms by which response distortions can arise from undistorted representations. Specifically, we were interested in the possibility that nonuniformities in the resolution of attention could produce distorted responses when estimating from undistorted memory. To explain this perspective, we return to the basic framework set up by Huttenlocher et al. (1991). They characterized the challenge of reproduction from memory as the translation of an inexact and probabilistic representation of position (see also Bays & Husain, 2008) into discrete responses (e.g., a mouse click). Thus, an observer “estimates” by choosing a single location from a representation of all the positions likely to have been occupied by the memory item. As a result, responses should deviate probabilistically around the true position of a memory item.

But a second level of complexity is not typically included in discussions of human performance, even in a task this simple: Specifically, if an observer intends to execute an action in a particular location, motor noise should result in outcome deviations from that intent, regardless of how accurate the intent was in the first place. To our knowledge, this point has not been made in the context of the majority of spatial-memory and -attention tasks, though it played an important role in a series of studies on actions toward spatial targets with different risk and reward schedules (e.g., Trommershäuser, Landy, & Maloney, 2006). The key insight in this work is that an observer should take all sources of noise into account when planning an action.

To make this idea concrete, we draw a distinction between the position of a *goal* and the position of an *action target*. By *goal*, we mean the actual position in the world where a response would count as a “hit,” where an agent could consider his or her action successful with respect to the task. In American football, for example, a quarterback seeks to pass the ball to a receiver who occupies or will come to occupy a specific position in the world. However, the quarterback does not know exactly what that position is, and thus selects an action target from a probability distribution of likely goal positions. Of course, the action target should be selected as a best guess about the goal’s position, so that the action will be likely to result in success. In this way, the selection of action targets from represented goal positions will contribute to variability in a quarterback’s outcomes (i.e., successful passes).

Another factor will contribute as well: deviations of an action’s end point from its intended action target. Returning again to football, a throw can be inaccurate because an action target deviates considerably from a receiver’s actual position (the goal), and also because the throw itself deviates from its action target. So, a quarterback should expect that his intent could be inaccurate and also that his throw—the execution of his intent—could be inaccurate. More broadly, if an observer possesses knowledge of his or her own variability—a common assumption in research that views perception and action probabilistically—as well as knowledge about the sources of that variability, then he or she should expect that the action targets are probabilistic estimates with respect to the goals, and also that the executed actions will deviate probabilistically from intentions.

Given this framework, how should action targets be selected? Of course, the first step is to sample from represented distributions of goal positions in a way that maximizes the likelihood of selecting a goal’s true position. But, within the space of generally likely positions, perhaps a second pressure could ultimately result in bias: Specifically, one should seek an action target that is easily discriminable from its neighbors; one should prefer precise positions for action targets. This may seem unintuitive, and so an analogy is perhaps useful. Imagine launching a missile. Code has been written to direct the missile to its intended coordinates as accurately as possible. You are in control of selecting the coordinates. With the expectation that the missile’s landing will deviate somewhat from the input coordinates, and the desire to avoid unintended consequences, you should prefer a target with a position that can be described very precisely, with as many significant figures as possible. So, for example, given multiple potential targets, you should prefer to aim for one whose position can be described in terms of a precise longitude and latitude, including multiple decimal places, as opposed to one that is known only to be someplace in North America.

In human vision, several factors can contribute to the precision with which positions can be discriminated from one another. One is eccentricity. Imagine several laterally arranged dartboards and one dart. Hitting the bull's eye on any of the boards will result in a large reward. But there is a catch: You must fixate in one predetermined location during your throw. It should be obvious that the dartboard you should target is the one closest to fixation.

The present study

Our interest in the present study was not in effects related to eccentricity, but in a similar kind of effect that could arise because of nonuniformities in the resolution of attention. Specifically, resolution is more precise below fixation than above it (Intriligator & Cavanagh, 2001).¹ Accordingly, when making a reproduction response from memory, an observer should first fixate as closely as possible to the object's remembered position, and then she should select target positions for action *below fixation*, since those positions should appear more discriminable from one another—that is, able to be targeted more precisely—than positions above fixation. As a result, we hypothesized that observers should produce responses that are systematically biased below their fixation position when translating probabilistic knowledge into a response at a single location.

To investigate these issues, we sought the simplest memory task in which a spatial distortion could arise. In Experiment 1, participants saw a display with a randomly placed dot that was then masked, and at test the task was to reproduce the dot's position via a mouse click. To foreshadow, the responses were biased downward, as expected, and Experiments 2 and 3 were designed to interrogate this bias, demonstrating that it arises because of nonuniformities in the resolution of attention that are dependent on fixation.

Experiment 1: A downward response bias in a simple task

We designed the simplest experiment that we could think of in which a downward spatial bias could arise. In it, participants saw a dot that was then masked, and they used the mouse to localize the dot to the position they remembered it being in. As we noted already, many previous studies have revealed spatial distortions in tasks that involve localizing just one or a few dots. Those designs have to varying degrees been more complicated than ours, involving, among other things, selective

¹ These effects have been previously described as effects of *attentional resolution*, and so we utilize this characterization throughout the article. However, the exact causes of the effect are unknown, to our knowledge, and one might prefer a more theoretically neutral term, *spatial precision*. For the present purposes, these terms are interchangeable.

attention to some objects and not others, enforced fixation during presentation and test, and/or multiple items in the scene (Dent & Smyth, 2006; Liverence & Scholl, 2011; Sheth & Shimojo, 2001; Suzuki & Cavanagh, 1997; Zimmermann, Fink, & Cavanagh, 2013). We intentionally avoided such manipulations in the present design, on the presumption that distortions can arise due to a variety of competing factors, which could make it difficult to isolate the causes of overall effects—in this case, those having to do with upward or downward biases relative to fixation when making a response. For example, if fixation were enforced during the presentation of the stimulus, then distortions in localization known to occur toward fixation (e.g., Tsal & Bareket, 2005) would be expected to interact with response fixation effects (if present). This would make it more difficult to identify the response-driven effects, especially if distortion magnitude varies as a function of cause. Similarly, previous work has suggested that visual space becomes distorted when one makes saccades across a display (Ross, Morrone, & Burr, 1997). By allowing free viewing in the present experiment, and by randomizing dot placement in the display, we hoped that any such effects would cancel each other out through the accumulation of trials.

Experiment 1 did not include eye-tracking so that we could first determine whether any biases would arise in our paradigm. We predicted that observers should report dot positions, on average, below their true locations. This would be consistent with undistorted memories being followed by responses biased to land below fixation.

Method

Participants A group of 23 Johns Hopkins University undergraduates participated in Experiment 1, and the results from two participants were excluded due to noncompliance with the instructions. All participants reported normal or corrected-to-normal visual acuity. Participation was voluntary and done in exchange for extra credit in related courses. The experimental protocol was approved by the Johns Hopkins University institutional review board (IRB).

Apparatus Experiment 1 took place in a dimly lit sound-attenuated room. Stimuli were presented on a 29.5 × 22.5 cm Macintosh iMac computer with a refresh rate of 60 Hz. The viewing distance was 60 cm, so that the display subtended 39.43° × 24.76° of visual angle.

Stimuli and procedure Stimuli were generated using MATLAB and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). All stimuli were presented within the full display frame of 39.43° × 24.76°. At the start of each trial, participants were shown a screen instructing them to press any button on the keyboard when they were ready to begin. A target dot

(0.43° in diameter) then appeared in a randomly selected position any place on the screen. After 200 ms, it was followed by a mask for 900 ms. We used 72 unique masks, each comprising between 40 and 80 randomly placed dots that were identical to the stimulus dot.

After the mask, a mouse cursor appeared on a blank gray screen. Participants were instructed to click on the location of the remembered dot. Figure 1 depicts the procedure.

The mouse cursor was not present during the sample stimulus or the mask display, appearing only during the response display. The cursor's onset position was unpredictable, selected randomly to fall on the perimeter of a circle centered on the memory dot's true position, and with a radius length of 1.5° , 3° , 4.5° , or 6° . When the cursor appeared, it was identical in shape and color to the memory target, so that reproduction of the test display could be thought of in terms of moving the dot to its prior location. Each participant completed 432 trials. Eye movements were not monitored, and no instructions were given concerning eye movements.

Results

The main question of interest was whether response distributions would tend to center below a memory object's true location. Additionally, we investigated the presence of any left/right biases, the possibility of effects that were mediated by the position of the cursor when it first appeared, and whether the target appeared in the upper or lower visual field.

Response distributions Figure 2 displays all of the responses made by participants in Experiment 1, standardized so that the true position of a memory item was 0,0. As is clear in the figure, a greater proportion of responses fell below the true




Fig. 2 Standardized plot of all mouse responses in Experiment 1. 0, 0 designates the true position of each memory item in each trial. A group of 21 participants produced a total of 8,857 responses. The accompanying histograms show the distribution of standardized X and Y response coordinates

position than above. This effect is summarized by the histogram to the right of the response plot. For comparison, we found no horizontal bias, as is shown by the histogram below the plot. These results can be summarized as follows: Response distributions were centered reliably 0.29° below the memory dot's true position, $t(20) = 6.913$, $p < .001$, $d = 1.51$. This effect did not vary significantly as a function of cursor distance, $F(3, 60) = 1.682$, $p = .18$, $\eta_G^2 = .008$, nor of the cursor's relative starting position to the left/right or above/below the true position, $t(20) = 1.034$, $p = .314$, $t(20) = 0.209$,




Fig. 1 Procedure of Experiment 1

$p = .837$, respectively. The effect also did not vary as a function of whether the dot's true position was in the upper or lower visual field, $t(20) = 0.74$, $p = .34$, $d = 0.27$, and no statistically significant horizontal bias was apparent, $M = 0.07^\circ$ to the left, $t(20) = 1.81$, $p = .09$, $d = 0.40$. On a participant-by-participant basis, the extent to which responses appeared biased to the left was driven primarily by two participants. When they were excluded from the analysis, the effect was smaller, $M = 0.03^\circ$ left, $t(18) = 1.05$, $p = .31$, $d = 0.24$.

Discussion

Experiment 1 accomplished its goal, identifying a downward response bias in a simple task. This response bias did not vary as a result of where the dot appeared on the screen (i.e., upper or lower visual field) or as a function of the position of the mouse cursor around the target's true position. In general, an intuitive explanation for effects like these might appeal to kinematic effects, such as biases to move as little as possible or to move an object away or toward one's body as much or as little as possible. But note that in this experiment, the target positions were selected randomly on a trial-by-trial basis, and the mouse was not repositioned in any systematic way after each trial. In other words, no systematic relationship should have emerged between the amount or direction of movement applied to the mouse and the position of a response. Thus, the effect seems likely to have arisen from causes that were at least not exclusively motoric.

Experiment 2: Eye-tracking

Our hypothesis was that the response bias observed in Experiment 1 would emerge because of differences in the resolution of attention above and below fixation. Where observers fixated when making a response was critical to support this reasoning. Thus, in Experiment 2 we sought to replicate the effect discovered in Experiment 1, while also tracking observers' fixations and saccades. We emphasize that the measure of interest was fixation position *during response*. As in Experiment 1, participants were allowed to move their eyes freely throughout a trial.

Method

Participants A new group of ten Johns Hopkins University undergraduates participated in Experiment 1. The results from two participants were excluded due to erratic fixations that likely arose from poor calibration (all of the main effects remained the same with these two participants included). All participants reported normal or corrected-to-normal visual acuity, and participation was voluntary and done in exchange

for extra credit in related courses. The experimental protocol was approved by the Johns Hopkins University IRB.

Apparatus Experiment 2 took place in a dimly lit room. A Mac Mini computer equipped with MATLAB software and the Psychophysics Toolbox was used to present stimuli on a 36×27.5 cm Dell P991 CRT monitor with a refresh rate of 60 Hz. The viewing distance was 63 cm, so that the display subtended $31.9^\circ \times 24.6^\circ$ of visual angle. Eye movements were recorded using an EyeLink 1000 eye-tracker (SR Research, 0.15° resolution). A chinrest was used to minimize head movements.

Stimuli and procedure The procedure was identical to that of Experiment 1, except that (a) participants completed six practice trials before eye-tracker calibration, to ensure that they understood the instructions, and (b) a standard calibration and validation procedure was conducted before the start of the experimental session. No instructions were given as to how or where to fixate; participants were simply told that we were interested in measuring eye movements during a memory task.

Results

Response distributions Experiment 2 replicated the behavioral results observed in Experiment 1. Response distributions were, on average, centered 0.46° below the memory dot's true position, $t(7) = 7.547$, $p < .001$, $d = 2.67$. This effect did not vary as a function of cursor distance, $F(3, 21) = 0.917$, $p = .450$, $\eta_G^2 = .02$, or of whether the dot appeared in the upper or lower visual field, $t(7) = 0.56$, $p = .61$, $d = 0.17$. Also, we found no overall bias to respond to either the left or the right of a target, $M = 0.05^\circ$ left, $t(7) = 1.041$, $p = .332$, $d = 0.37$.

Eye-tracking data Eye-tracking was used for two reasons. (1) We hypothesized that participant fixations could reveal unbiased knowledge of the memory item's position. (2) We expected that responses should tend to fall below fixations. We analyzed the results as follows: For each participant, the last fixation of each trial was compared to the position of the memory dot. Figure 3a displays all final fixations made by participants, standardized so that the true position of the memory item is 0,0 (i.e., allowing for a comparison with click responses, shown in Fig. 3b). The final fixations were, on average, unbiased in any direction, and distributed normally around a memory item's true position. Statistical analyses confirmed the lack of an upward or downward bias, $M = 0.17^\circ$ above, $t(7) = 0.829$, $p = .435$, $d = 0.29$, as well as the lack of a rightward or leftward bias, $M = 0.23^\circ$ left, $t(7) = 1.216$, $p = .263$, $d = 0.43$.

One might be concerned that fixations did not reveal a bias because they included more variance, generally, than mouse

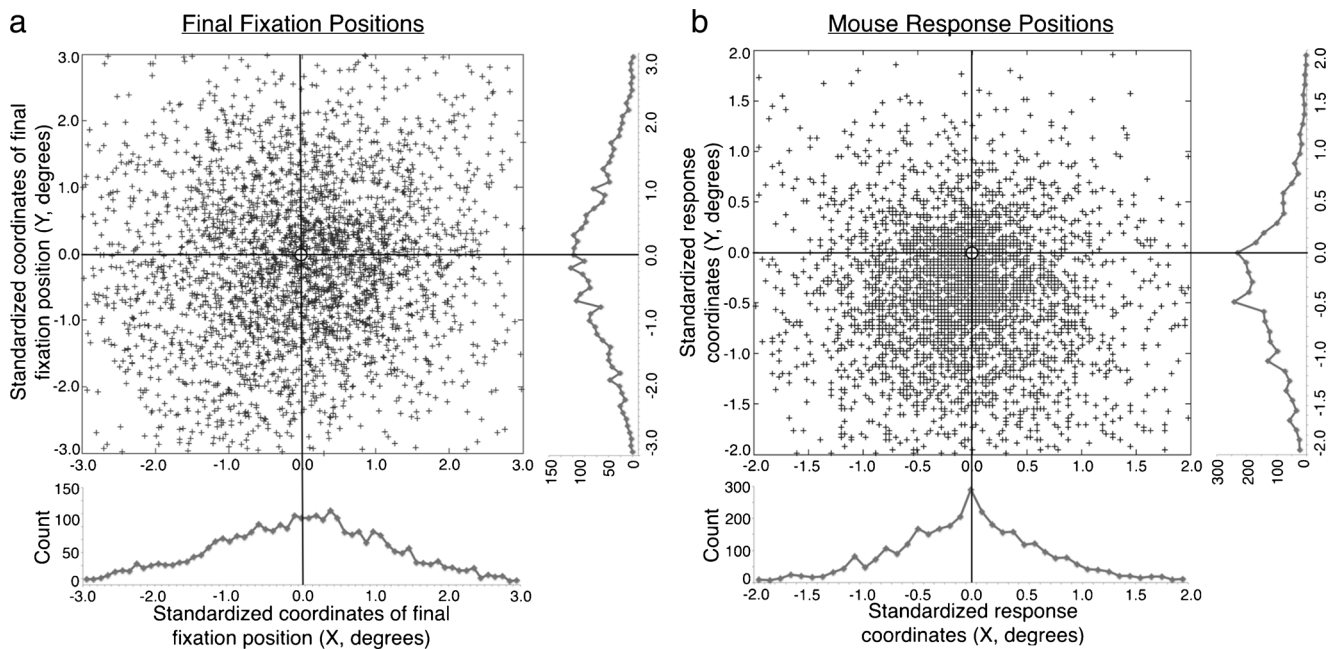


Fig. 3 Standardized plots of all final fixations (a) and mouse responses (b) in Experiment 2. 0, 0 designates the true position of each memory item in each trial. Eight participants produced a total of 3,453 responses. The

accompanying histograms show the distribution of standardized X and Y fixation and response coordinates

clicks, resulting in an underpowered tool. It is likely that this variance had more to do with the limited spatial resolution of the eyetracker (0.15°) versus the pixel-by-pixel resolution of mouse clicks. Moreover, the non-significant tendency in fixation positions was *above* a dot's actual position, making it unlikely that the actual tendency was downward, but hidden and underpowered.

In addition, we performed a follow-up analysis to investigate the relationship between individual fixations and paired responses for each trial. Specifically, we calculated the difference between each last-fixation position and its respective mouse response on a trial-to-trial basis. On average, the last fixation Y position was significantly above the response Y position by 0.63° , $t(7) = 0.99$, $p = .01$, $d = 1.20$. In contrast, we observed no significant difference between the last fixation and response X positions, $M = 0.19$, $t(7) = 0.70$, $p = .40$, $d = 0.32$. In other words, each individual mouse response tended to fall below the corresponding fixation at that moment in time. Thus, the difference between fixations and click responses was not just a difference that arose over the course of an experiment, on average, but also one that occurred in individual trials.

Discussion

The main result of Experiment 1 was a downward response bias in a memory reproduction task for the position of a single dot. Experiment 2 demonstrated that in the same task, fixations during response were unbiased, centered appropriately, and distributed normally around a memory item's true

position. Accordingly, *at least* the mechanisms guiding eye movements possess unbiased memory for the object's position. It is possible that the mechanisms guiding mouse responses utilize different memory than do eye movement mechanisms, but it is simpler to assume that they rely on the same memory. Thus, we propose that click responses become biased because positions below fixation are more attractive as targets for response. This is because the resolution of attention is better below a fixation position than above it, an effect that has been shown previously by, for example, asking participants to perform an object-tracking task while fixating either above or below the display (with better performance for fixation positions above the display; Intriligator & Cavanagh, 2001).

According to our hypothesis, then, the events culminating in the execution of a reproduction response unfold more or less as follows: Before the appearance of the memory stimulus, a participant does not know where the memory item will appear. When the memory item appears, a participant stores an unbiased but probabilistic estimate of its position on the screen. The memory item disappears after 100 ms—not enough time to complete a saccade to it. However, in preparation for making a response, a participant will move his or her gaze to fixate a position drawn from the memory estimate of the target item's position, ultimately producing fixation patterns that are normally distributed around the item's position. This is consistent with the fact that, in general, human observers direct fixations toward the targets of their actions. Finally, the observer must select a position to target an action toward. He or she should know (implicitly) that the current

fixation position is only an approximate estimate of the target's true position, and also, that the action will deviate probabilistically from his or her intent. As a result, the participant would tend to select positions to target below the fixation, since these positions appear more discriminable from one another—more precisely targetable—than positions above.

According to this interpretation, the implication is that biased responses do not always reflect biased memory, an outcome that is consistent with prior work on prototype and boundary effects in spatial memory (Huttenlocher et al., 1991). These experiments add to this literature by demonstrating that the demand to select a target position from a probabilistic representation depends on attentional mechanisms. Experiment 3 was designed to further support these conclusions.

Experiment 3: A lateral bias yoked to a lateral head tilt

Up and *down* are absolute terms; at least in common usage, we reference them with respect to the sky and the ground. The claim with respect to Experiments 1 and 2 was that the bias observed was not absolute—though it was downward for all participants—but egocentric instead. In egocentric terms, our theory is that the observed response bias is caused by better attentional resolution in the “nose-ward” direction, with respect to the axis adjoining the two eyes. Experiment 3 was designed to constitute a very direct test of this interpretation. It was identical to Experiment 1, except that on each trial a participant tilted his or her head to either the left or the right (resting it on a stack of books). We predicted that response biases should materialize in the horizontal direction opposite the direction of head tilt in a given trial.

Method

Participants A new group of ten undergraduates took part in Experiment 3. One additional participant was tested in a follow-up experiment, as reported below. All participants reported normal or corrected-to-normal visual acuity; half of them were summer interns in the Department of Psychological and Brain Sciences who volunteered to participate, and the remaining half were paid volunteers from among the summer session students throughout the university. The experimental protocol was approved by the Johns Hopkins University IRB.

Apparatus The apparatus was the same as that used in Experiment 1, with the following exception. A stack of books was placed in front of the screen for the entire experiment (Fig. 4).

Stimuli and procedure The stimuli and procedure were identical to those of Experiment 1, with the following exceptions. Participants tilted and rested their heads on a stack of books

that was covered with a pillow for comfort. For half of the experiment, they performed the task with their heads tilted in one direction, and during the other half of the experiment, they performed the task with their heads tilted in the other direction. For clarity, we will refer to rightward tilts as meaning that a participant's right ear moved toward the right shoulder, and leftward tilts will designate the opposite situation. The order of head tilt directions was counterbalanced across participants.

Results

The prediction for this experiment was straightforward: Leftward head tilts should result in rightward response biases, and vice versa. That is, bias should be in the opposite direction from head tilt. Consistent with this prediction, we observed a rightward bias when participants tilted their heads to the left ($M = 0.21^\circ$ right) and a leftward bias when they tilted their heads to the right ($M = 0.18^\circ$ left). No significant difference was apparent between the biases observed as a function of tilt direction, $t(9) = 0.38, p = .71$. This allowed us to combine the data from both tilt directions in order to analyze the results with more power. We standardized responses so that an object's true position was designated 0,0; in leftward-tilt trials, negative X coordinates designated rightward positions, and in rightward-tilt trials, negative X coordinates designated leftward positions. This was done so that the predicted biases would appear in the same hemisphere in depictions and analyses. One way to think about these analyses is as follows: For trials with rightward head tilts, we analyzed the results exactly as we had in Experiment 1—left on the observer's screen was also left in the analyses and depictions. In contrast, for leftward-tilt trials, we first standardized responses in the same way, and then rotated the output by 180 deg—causing left on the screen to be right in the analyses and depictions.

In Experiment 3, we found the predicted directional bias. Response distributions were shifted significantly laterally (by 0.20° , on average) in the direction opposite each trial's head tilt, $t(9) = 5.175, p < .01, d = 1.16$ (Fig. 5).

We also found a significant 0.20° downward response bias, $t(9) = 4.277, p < .01, d = 1.26$. This was not surprising, since the head tilt was not perfect: It resulted in a diagonal head orientation, as opposed to a perfectly horizontal one. Moreover, two participants were considerably shorter than the others. Accommodating them while making the monitor visible required an unplanned partial vertical rotation of the monitor. Excluding these two participants had no effect on the lateral biases observed, but it reduced the size of the downward bias. We also note that with all participants included, the downward bias observed in this experiment was significantly smaller than the one observed in Experiment 2, $t(16) = 3.48, p < .01$.

In retrospect, it might have been better to measure each participant's head tilt angle in order to make predictions about



Fig. 4 Experimental setup of Experiment 3: Participants tilted and rested their heads on a stack of books either to the right (left panel) or to the left (right panel)

the degree of downward bias to expect.² Since we could not do this after the fact, we instead ran one additional participant to use as a proxy for the group (the participant was naïve and drawn from the same population as the participants in all the reported experiments). We thus sought to provide a reference for examining the degree to which an observer could comfortably tilt his or her head and the exact impact that this should produce on responses.

The participant completed the same experiment as the previous ten participants. Her response directional bias was similar to the group's laterally ($M = 0.3^\circ$ opposite the direction of head tilt), but she did not display any real vertical bias ($M = 0.05^\circ$ upward). The critical difference in the testing procedure was that during the course of the study, we measured the angle of her head tilt. Before starting the experiment, the participant was asked to place her head comfortably on the stack of books, and then to tilt her head as far as she could. She was unable to tilt farther than 74° , and she reported feeling most comfortable at 60° . Measurements were taken in both tilt directions. This pair of values characterize her widest range. Next, on four occasions (two in each direction) during the course of the experiment, the experimenter spontaneously entered the testing room, and when the current trial was completed he asked the participant to remain in the same position, and he then measured her head tilt. Averaging together the measures taken during these four instances, we found an average head tilt of 64° . Taken as a proxy for the other ten participants, this value conformed well with the partially downward and partially lateral biases we observed in the group. The lateral bias in Experiment 3 was about 70 % as big as the downward bias in Experiment 1, and a head tilt of 64° is about 70 % of a 90° tilt.

² We thank an anonymous reviewer for pointing this out.

Discussion

Experiment 3 further supports the hypothesis that a spatial response bias can arise because of non-uniformities in attentional resolution relative to the axis adjoining the two eyes. More broadly, Experiment 3 demonstrates that spatial response biases can arise egocentrically—here, in relation to one's head orientation. When spatial response biases are observed, they are typically attributed to representational contents that reference objects in the world to one another—for example, representations of one object's position with reference to the position of another relevant object (e.g., Liverence




Fig. 5 Standardized plots of all mouse responses in Experiment 3. 0, 0 designates the true position of each memory item in each trial. All trials were treated as if the head was tilted to the right. Ten participants produced a total of 4,233 responses. The accompanying histograms show the distribution of standardized X and Y response coordinates

& Scholl, 2011; Sheth & Shimojo, 2001). Future research should also consider potential egocentric causes of response biases among the possible mechanisms that lead observers to respond to space in distorted ways.

General discussion

We sought to investigate one possible way that responses derived from spatial memory could become biased: In particular, the resolution of attention is better below than above fixation. This led us to consider the possibility that vertical response biases could arise even with unbiased memories. Experiment 1 evidenced such a bias—specifically, a downward bias for reproducing the location of a single dot. Experiment 2 demonstrated that eye fixation during response was unbiased, and that responses were biased to land below individual fixations (and below the correct position). Finally, Experiment 3 supported the theory that the observed downward biases were caused by egocentric features of attention relative to the axis adjoining the eyes, by demonstrating that the bias could be shifted to the right or left simply by requiring participants to tilt their heads.

Theories of spatial representation and distortion

As we noted early on, spatial distortion has played a role in research on spatial representation from the outset. Broadly, there are two kinds of theories for how distortions can arise. The first involves unbiased or undistorted representations, but distorted responses arising from the coordination of multiple mechanisms. Perhaps the best known of such theories is that of Huttenlocher et al. (1991). They investigated memory for spatial position in environments with obvious landmarks and boundaries. Their theory was that observers naturally employ simultaneous representational formats in these situations—representing the absolute location of an object with some uncertainty, and also representing its position in more categorical and descriptive terms with respect to any boundaries or landmarks (e.g., the object was in Quadrant 1). Distortions then arise when producing a discrete response based on noisy estimates in position memory and through reconciliation with category information. For example, sampling from a noisy distribution can become truncated if the distribution crosses an obvious boundary.

Our results expand the scope of mechanisms by which distortions can arise from the challenges of estimating and sampling a response from an unbiased but noisy position memory. Specifically, the challenge of generating a reproduction response from a noisy memory requires the selection of a point in the current display to serve as a target for one's response. At the same time, an observer should assume that the target might not be hit head-on (Trommershäuser et al.,

2006). In other words, the observer contends with two kinds of noise: noise in the memory representation, and noise arising mechanically while producing a response. To minimize the combined effects of such noise, the observer should select easier-to-hit targets—target positions that are more easily discriminable from their neighbors. Returning to the analogy used earlier, all else being equal, one should prefer a target with a precisely specified location as when launching a missile. Since the resolution of attention is more precise below than above fixation, observers should systematically select points below. This bias is independent of the spatial representation in memory, and thus reflects a new kind of interaction between attention, fixation, and memory-derived spatial reproduction.

The second kind of theory concerning spatial distortion when utilizing memory appeals to distortions present in the contents of memory—that is, distorted memories per se. Although in some contexts distortions likely do exist in the contents of spatial memory—and certainly in the contents of perceptual representations (e.g., the Müller-Lyer illusion)—our results suggest some practical considerations for carefully identifying such distortions. A quick glance at Experiment 1 might have led to the conclusion that an observer's representation of a single object in memory is distorted below its true location. However, further exploration revealed that responses were distorted, but not memories.

Prescriptions

To more concretely identify the practical implications of our results for identifying genuinely distorted spatial representations, we suggest some prescriptions for future research. Specifically, the present set of studies serve as a reminder that reaching and clicking actions do not map directly from the representation of a position, but that they are mediated by a variety of mechanisms, including eye movements. Accordingly, eye-tracking should routinely be employed in the study of spatial distortions. In particular, when reproduction responses evidence distortion, they should be related to fixation tendencies. Second, we suggest that attempts should be made to manipulate distortions systematically on the basis of hypotheses concerning their origins. This may not be possible in every case, but it is worth attempting, since manipulating an effect can aid in discriminating between hypotheses. For instance, in Experiment 3 we shifted distortions laterally by requiring head tilts, a prediction that would seem hard to reconcile with any other theory. More specifically, Experiment 3 took advantage of the fact that our theory predicted distortions that arise relative to egocentric reference points. By devising a manipulation that dissociated egocentric and allocentric references, we could test this theory more directly. Since many spatial distortions are accounted for in terms of representational contents, the reference frame of the relevant

contents should be specifiable, as well. As a result, many theories could likely be tested relative to simple egocentric manipulations.

These prescriptions can be viewed through the lens of prior research in which distortions were identified and attributed to representational contents. For instance, Dent and Smyth (2006) investigated capacity limitations in spatial working memory through a reproduction task much like ours, but with multiple items. Their mean reproduction error increased with a memory load up to three objects, and then it stabilized. These results were interpreted as evidence of a discrete capacity limit in spatial working memory—in other words, a limit on memory contents. However, only mean errors were provided, without characterizing the distribution of responses relative to each item and the order in which they were responded to. This makes the results difficult to interpret. Multiple serial responses add to the challenge of estimating from noisy representations, by requiring movements between reasonable response positions and also saccades between them. Thus, executing the task involves the coordination of multiple mechanisms in ways that likely interact with memory load. Some effects are likely to reflect restrictions on memory contents, but others may reflect complex dynamics that remain otherwise unexplored. A study requiring these coordinated demands should utilize multiple manipulations and eye-tracking.

Conclusion

Attention is among the most extensively studied and interesting aspects of human visual cognition. We often describe it as facilitating, speeding, or otherwise making responses more accurate. Another way to put this is to say that attention distorts responses: They look different for attended versus unattended stimuli. Typically, we manipulate attention (e.g., with cues) to study the response distortions that it produces. When we do not manipulate attention, however, we should not assume that it is uniform in the visual field. Our results demonstrate that a known nonuniformity can influence reproductions from spatial memory.

Author note We thank Kaan Zaimoglu for assistance with the data collection, and Howard Egeth for suggesting Experiment 3 and for providing additional helpful comments.

References

- Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854. doi:10.1126/science.1158023
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436. doi:10.1163/156856897X00357
- Dent, K., & Smyth, M. M. (2006). Capacity limitations and representational shifts in spatial short-term memory. *Visual Cognition*, *13*, 529–572. doi:10.1080/13506280444000760
- Diedrichsen, J., Werner, S., Schmidt, T., & Trommershäuser, J. (2004). Immediate spatial distortions of pointing movements induced by visual landmarks. *Perception & Psychophysics*, *66*, 89–103. doi:10.3758/BF03194864
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376. doi:10.1037/0033-295X.98.3.352
- Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology*, *43*, 171–216. doi:10.1006/cogp.2001.0755
- Liverence, B. M., & Scholl, B. J. (2011). Selective attention warps spatial representation: Parallel but opposing effects on attended versus inhibited objects. *Psychological Science*, *22*, 1600–1608. doi:10.1177/0956797611422543
- Marr, D. (1982). *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco, CA: W. H. Freeman.
- Nelson, T. O., & Chaiklin, S. (1980). Immediate memory for spatial location. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 529–545. doi:10.1037/0278-7393.6.5.529
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. doi:10.1163/156856897X00366
- Ross, J., Morrone, M. C., & Burr, D. C. (1997). Compression of visual space before saccades. *Nature*, *386*, 598–601. doi:10.1038/386598a0
- Sheth, B. R., & Shimojo, S. (2001). Compression of space in visual memory. *Vision Research*, *41*, 329–341. doi:10.1016/S0042-6989(00)00230-3
- Suzuki, S., & Cavanagh, P. (1997). Focused attention distorts visual space: An attentional repulsion effect. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 443–463. doi:10.1037/0096-1523.23.2.443
- Trommershäuser, J., Landy, M. S., & Maloney, L. T. (2006). Humans rapidly estimate expected gain in movement planning. *Psychological Science*, *17*, 981–988. doi:10.1111/j.1467-9280.2006.01816.x
- Tsal, Y., & Bareket, T. (2005). Localization judgments under various levels of attention. *Psychonomic Bulletin & Review*, *12*, 559–566. doi:10.3758/BF03193805
- Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, *13*, 407–433. doi:10.1016/0010-0285(81)90016-5
- Zimmermann, E., Fink, G., & Cavanagh, P. (2013). Perifoveal spatial compression. *Journal of Vision*, *13*(5), 21. doi:10.1167/13.5.21