Oblique biases: An instance of domain- and modality-general spatial representation

Sami R. Yousif[1] & Samuel D. McDougle[2]

[1]University of Pennsylvania, Department of Psychology

[2]Yale University, Department of Psychology

Running Head : Generalized Oblique Biases

Addresses for correspondence : Sami Yousif
Department of Psychology
University of Pennsylvania
425 S. University Ave, Stephen A. Levin Bldg.
Philadelphia, PA, 19104-6241

Email : sryousif@sas.upenn.edu

Word Count : 9650
Version : Preprint

Abstract (230 words)

A variety of phenomena related to the oblique regions of space have been observed across modality (e.g., in vision and in action) and across domain (e.g., for properties like orientation and location). For instance, the classic 'oblique effect' describes a deficit in visual acuity for oriented lines in the oblique regions of space, and classic 'prototype effects' describe a bias to mis-localize objects towards the oblique regions of space. While there has been speculation that some 'oblique-related effects' share a common mechanism, many of these effects are explained in very different terms. The oblique effect itself is often understood as arising from coding asymmetries in orientation-selective neurons in the brain, whereas prototype effects have been described as arising from categorical biases in higher-level cognition. But is it mere coincidence that there are so many distinct effects linked to the oblique regions of space? Here, we explore the possibility that most, if not all, known 'oblique-related effects' may stem from a single, underlying spatial representation. In two first experiments, we show that individuals show stable oblique biases across domain and across modality, and we explore how both biases may have a common cause. Then, in a final experiment, we show that this perspective correctly predicts behavior in a novel spatial judgment task. Thus, we argue that a single (distorted) spatial representation may be the root cause of dozens of known phenomena.

Significance Statement (118 words)

Spatial representation spans multiple modalities and domains. The visual and somatosensory systems are involved in *perceiving* space, for instance, while the motor system is responsible for *acting* in it. To what extent are spatial representations across modalities and domains shared? Are there shared representational formats for perception and action? As a case study, we explore 'oblique-related effects'. Many such effects have been observed, yet they have been explained in radically different ways. Using an individual differences approach, we show that these effects — of location and of orientation, in perception and in action — arise from a single, underlying, distorted representation of space. These findings represent a striking example of a highly general representational format common across modalities and domains.

The oblique effect describes the phenomenon whereby observers are worse at discriminating oriented bars presented in the oblique (diagonal) regions of space compared to the cardinal (horizontal/vertical) regions. It is one of the most robust psychophysical effects ever studied. But what is the nature of the oblique effect? Typically conceived as a bias of orientation, it has traditionally been explained by appeal to coding asymmetries in orientation-selective neurons in the visual cortex (e.g., Li et al., 2003). However, a range of related effects have been observed not just in orientation judgment tasks, but also in location judgment tasks (Yousif et al., 2020), location placement tasks (Huttenlocher et al., 1991), various haptic/motor tasks (e.g., Gentaz & Hatwell, 1995; Gordon et al. 1995; Smyrnis et al., 2007), and even various aesthetic judgment tasks (Latto et al., 2000; Latto & Russell-Duff, 2002; Plumhoff & Schirillo, 2009; Youssef et al., 2015). Moreover, 'oblique-related' effects come in several different forms: Some of these effects are about reduced visual *acuity* in the oblique regions of space (e.g., Appelle, 1972; Yousif et al., 2020), whereas others involve memory errors and mis-localizations *towards* the oblique regions. Some involve vision (Huttenlocher et al., 1991; Latto et al., 2000; Yousif et al., 2020), while others are observed in the absence of visual input (e.g., Gentaz & Hatwell, 1995; Gordon et al., 1995; Smyrnis et al., 2007). Finally, some effects are characterized as attraction to certain regions of space, whereas others are characterized as effects of repulsion (see, e.g., Huttenlocher et al., 1991; Rademaker et al., 2017; Wei & Stocker, 2015). Do all of these effects reflect one underlying phenomenon, or many?

Surprisingly, these biases are often explained in radically different ways. While the standard visual oblique effect is explained by variance in neural representations across specific orientations (see, e.g., Furmanski & Engel, 2000; Li et al., 2003; see also Nasr & Tootell, 2012), oblique biases in spatial localization tasks have traditionally been explained by categorical effects of spatial representation (Huttenlocher et al., 1991). Meanwhile, oblique effects in haptic perception have been linked to gravitational cues (Gentaz & Hatwell, 1995) and oblique biases in motor responses (e.g., reaching) have been explained by the physical constraints of the human arm (Gordon et al., 1995).

Here, we consider the possibility that all of these 'oblique-related' effects originate from a singular deficit in angular acuity in the oblique regions of space — a deficit that is not specific to any modality (i.e., it may span visual perception and motor control) nor stimulus domain (i.e., it may manifest as a bias of orientation or as a bias of location). If correct, our unified framework of oblique effects points to a generalized spatial representation that is universally deployed for perception and action, and likely for all sorts of spatial tasks.

The oblique effect(s)

The oblique effect typically refers to the phenomenon whereby observers, human and non-human, are faster and better at discriminating oriented lines near the cardinal axes as opposed to the oblique axes (Appelle, 1972; Bonds, 1982). That is, a line oriented at, say, 3°, would be more readily discriminated from a line at 1° versus lines oriented at 48° and 46° (see Figure 1A). This phenomenon is well-replicated and exceptionally robust (see, e.g., Essock, 1980; Vogels & Orban, 1985; Furmanski & Engel, 2000). More recent work on the oblique effect has focused on the nitty-gritty details of its implementation; for instance, there has been considerable interest in the reference frames over which the oblique effect operates (e.g., Cecala & Garner, 1986; Luyat et al., 2001; Luyat et al., 2005; Luyat & Gentaz, 2002; Rademaker et al., 2017).

Details aside, there is consensus that the oblique effect is well-understood: It is thought that the oblique effect arises directly from the number of orientation-selective neurons devoted to processing certain orientations (Furmanski & Engel, 2000; Li et al., 2003; see also Nasr & Tootell, 2012). In other words, the idea is that there are more neurons specifically tuned for cardinal (and cardinal-adjacent) orientations than there are for oblique (and oblique-adjacent) orientations, likely reflecting the natural image statistics of the environment (Keil & Cristobal, 2000; Girshick, Landy, & Simoncelli, 2011; Henderson & Serences, 2021; Wei & Stocker, 2015).

However, there are a number of effects not specific to visual orientation perception that involve biases near the oblique regions of space. For instance, simple location memory tasks reveal strong biases *towards* the obliques (Huttenlocher et al., 1991). Huttenlocher and colleagues (1991) famously proposed that spatial localizations simultaneously depend on 'coarse' and 'fine-grained' representations, the former of which is dictated by higher-level spatial knowledge. They proposed that biases towards the oblique axes reflected a bias towards the 'prototype' – the center of the quadrant in which the point originated (see subsequent work on the 'Category Adjustment Model'; Holden et al., 2010; Holden et al., 2013). While not mutually exclusive with this category-based explanation, recent work has shown that, coincidentally, angular acuity for the location of visually presented dots is lower near the oblique axes of space (Yousif et al., 2020). This raises the possibility that a reduction in angular acuity for object position at the obliques may be related to the placement biases first observed by Huttenlocher and colleagues. (The relation between acuity and bias is one that is considered throughout this paper; see *Interim Discussion.*)

The story is further complicated by the fact that oblique effects have been observed in other modalities. Indeed, there are biases in both touch and motor control that resemble visual oblique effects. For instance, there is an analogous "motor oblique effect" (i.e., a bias for motor movements to err towards the oblique regions of space; Gordon et al. 1995; Gourtzelidis et al. 2001; Mantas et al., 2008; Petersik & Pantle, 1982; Sainburg et al. 1995; Smyrnis et al. 2000; Smyrnis et al., 2007) as well as a "haptic oblique effect" (i.e., a reduced ability to discriminate angled rods based on haptic information in the absence of vision; e.g., Gentaz & Hatwell, 1995).

Are all of these biases – visual, somatosensory, and motor – connected in some way? However convenient it would be to lump these effects together, it is not obvious that they reflect the same underlying processes. In fact, current explanations for motor and visual oblique effects could not be more different. Motor oblique effects have been explained by the "inertial field of the arm" (Gordon et al., 1995, p. 846), for instance, whereas visuospatial oblique effects have been explained by scene statistics and the superimposition of category knowledge on perceived locations (Huttenlocher et al., 1991).

If these various oblique biases *do* stem from the same underlying spatial representation, this suggests that the "oblique effect" is less well-understood than previously thought. To wit: Current explanations of the oblique effect emphasize *orientation-selective visual* neurons,

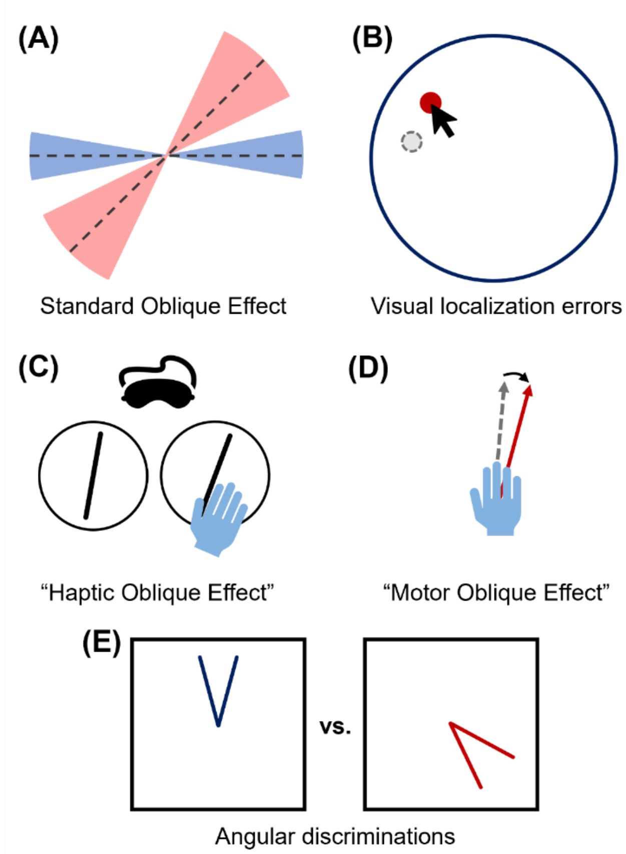

Figure 1. Various oblique biases. (A) A depiction of the classic oblique effect, whereby orientated lines/bars near the cardinal axes are perceived with greater acuity (e.g., Appelle, 1972). (B) A depiction of common localization errors whereby people are biased to remember items as having been closer to the oblique axes than they really were (e.g., Huttenlocher et al., 1991; Yousif et al., 2020). (C) A depiction of the "haptic oblique effect" observed in blindfolded, sighted individuals as well as blind individuals (e.g., Gentaz & Hatwell, 1995). (D) A depiction of the "motor oblique effect", wherein points and reaches err towards the oblique axes (e.g., Smyrnis et al., 2007). (E) A depiction of other plausible oblique biases, such as differences in angular perception, which are studied in this paper.

whereas there are known oblique biases that are not about orientation (i.e., they are about location; e.g., Huttenlocher et al., 1991; Yousif et al., 2020) and do not depend on vision (i.e., they depend on haptic/motor processing instead; e.g., Gordon et al., 1995; Smyrnis et al., 2007). What would it mean for all of these 'oblique-related' biases to reflect a single underlying representation, and, if they do, what is the nature of that representation?

Current Study

Here we explore the possibility that various known oblique biases result from a singular deficit in acuity at the obliques, one that is stable across tasks. First, we evaluate oblique biases across several tasks and explore relationships between them. In Experiment 1, we show that oblique biases related to orientation and location may share a common basis. In Experiment 2, we show that oblique biases are stable not only across tasks, but across modalities — we reveal stable individual differences in oblique biases across both visual and motor tasks. Finally, in Experiments 3a and 3b we explore the consequences of these oblique biases on the perception of technically empty spaces, demonstrating that angles spanning the oblique regions of space are perceived as more acute than those spanning the cardinal regions.

Experiment 1 — Oblique biases are task-general

The canonical oblique effect reflects a deficit in acuity for oriented lines in the oblique regions of space. However, there are similar deficits in acuity for angular position that are not about orientation *per se* (see Yousif et al., 2020). Are these phenomena related, or is this a coincidence? In a first experiment, we have participants complete both a location discrimination and an orientation discrimination task. We ask (1) whether we do in fact find an 'oblique effect' for both orientation and localization tasks, and (2) if those effects are related in some way (i.e., whether there are individual differences in visual angular acuity in the oblique regions that are stable across both tasks). If these biases share an underlying source, then we would expect oblique acuity biases to be correlated across tasks. If these biases are distinct, then we would expect no such correlation across tasks.

Method

This experiment, and all subsequent experiments, were preregistered. Those pre-registrations, as well as raw data and analyses, can be accessed here: https://osf.io/7tcbh/

Participants. 100 participants were recruited via Prolific. Here, and for all subsequent experiments in this paper, the sample sizes, primary dependent variables, and key statistical tests were chosen in advance and were pre-registered. The exclusion criteria we pre-registered were highly conservative, and thus no participants were excluded. This study was approved by the Yale Institutional Review Board.

Stimuli. There were two stimulus types: Dots for the location discrimination task and lines for the orientation discrimination task. Both sets of stimuli had similar properties. For the dot stimuli, a small black dot (10 pixels in diameter at default browser zoom distance) was presented relative to a central grey dot (10 pixels in diameter at default browser zoom distance). The black dots could initially appear at one of eight 'axes' around the grey dot (0, 45, 90, 135, 180, 225, 270, or 315 degrees), always 200 pixels away from the central grey dot. During the second presentation (see Procedure & Design), either the angle or distance of the black dot relative to the grey dot would change. Angle could change by +/- 4, 8, or 12 degrees, or 0 degrees (for a total of seven possibilities); distance could also change in seven increments, and the magnitude of the changes was set to match the difference of the angle changes in Euclidean distance (and so would change by either 0, 14, 28, or 42 pixels).

For the orientation discrimination task, the stimuli were similarly administered. Lines would initially appear along one of the same eight axes, and at second presentation would be altered by the same amounts in angular/distance space as described above. For distance changes, the line would simply shift along the set axis by that number of pixels.

Both the orientation task and the location task involved angle changes as well as distance changes. Distance changes were included to serve as a control. Including these changes makes it possible to ask whether any relation observed between angular judgments across tasks is specific to orientation or is more general (i.e., extending to all spatial judgments).

Each task had a total of 8 initial axes along which items could appear x 2 change types (angle, distance) x 7 increments (e.g., -12, -8, -4, 0, 4, 8, and 12 degrees) for a total of 112 trials. Thus, across both tasks, there were a total of 224 trials. A visual depiction of the task design and trial types can be seen in Figure 2.

Procedure & Design. The trials were blocked such that half of the participants completed the orientation discrimination task first and the other half completed the location discrimination task first. In both tasks, the initial image was presented for 1000ms before disappearing. After another 1000ms, the second image would appear, at which point participants were prompted to press 's' if the second image was the same as the first and 'd' if the second image was different form the first. Throughout the trials, there was a thin black border (4px) around the stimuli (800 pixels wide; 680 pixels tall). During the response window, that border briefly turned green, in order to signal to participants that they were able to respond. Between the blocks, there was a brief break, during which participants were reminded that the stimuli were going to change but that the task would remain the same. Prior to the first block, participants completed two representative practice trials (these data were not analyzed).

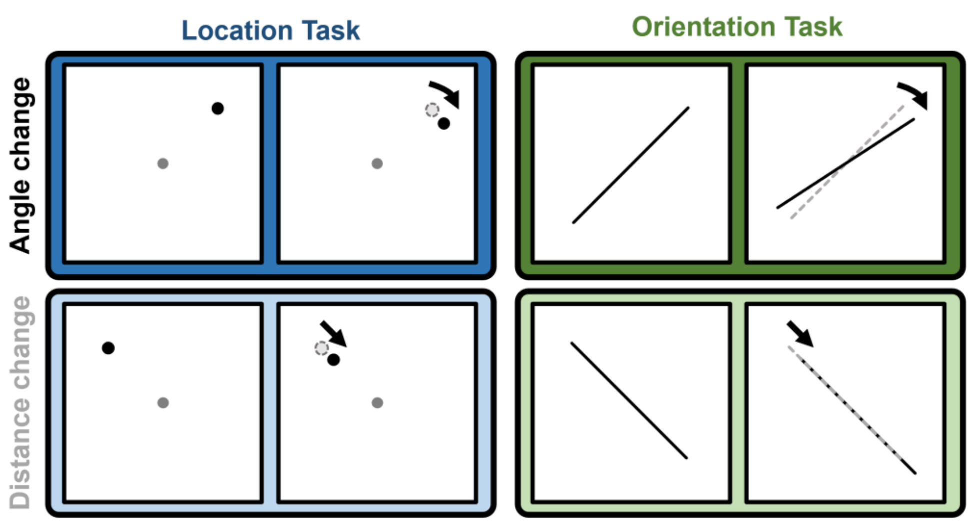

Figure 2. Schematic of task and trial types for Experiment 1 (relative stimulus locations not to scale).

Results and Discussion

First, we quantified the oblique biases themselves. We separately calculated trial accuracy for each trial type for each region of space. Then, we calculated a difference score between those two accuracy values, leaving us with a single value for each trial type that indicates whether participants were more accurate for trials in the cardinal regions vs. trials in the oblique regions (see Figure 3A). We observed significant oblique biases across both angle change detection conditions (location-task-angle: $t(99)=13.95$, $p<.001$, $d=1.40$; orientation-task-angle: $t(99)=16.20$, $p<.001$, $d=1.620$). Interestingly, we also observed significant oblique biases across both distance change detection conditions (location-task-distance: $t(99)=2.83$, $p=.006$, $d=.28$; orientation-task-distance: $t(99)=2.05$, $p=.04$, $d=.20$). Note however that these biases for the distance changes were in the opposite direction (i.e., better change detection in the oblique regions), were significantly smaller than their angular equivalents (ts>10.00, ps$<.001$, ds>1.00), and were inconsistent with prior work (e.g., Yousif et al., 2020). Moreover, the difference in the orientation task did not survive Bonferroni correction. Thus, while there may be small differences in acuity for distance changes between the oblique and cardinal regions, our data do not provide strong evidence in support of that conclusion.

For good measure we also conducted the same between-region analysis but for response time rather than accuracy. Participants were significantly faster to respond to angular differences on trials with items nearer to the cardinal axes in both the location (cardinal axes: $M=849$ms, $SD=226$ms; oblique axes: $M=912$ms, $SD=337$ms; $t(99)=2.37$, $p=.020$, $d=.24$) and orientation tasks (cardinal axes: $M=772$ms, $SD=215$ms; oblique axes: $M=894$ms, $SD=418$ms; $t(99)=3.13$, $p=.002$, $d=.31$), though the former result did not survive Bonferroni correction. Response time differences for judging cardinal versus oblique stimuli were not seen for distance change detection in either the location (cardinal axes: $M=941$ms, $SD=391$ms; oblique axes: $M=897$ms, $SD=244$ms; $t(99)=1.15$, $p=.25$, $d=.11$) nor orientation task (cardinal axes: $M=985$ms, $SD=483$ms; oblique axes: $M=966$ms, $SD=512$ms; $t(99)=.30$, $p=.77$, $d=.03$).

To begin to examine relations between the tasks, we first calculated the accuracy for each of the trial types across tasks (i.e., the overall proportion of trials for which participants pressed ‘same’ when there was no change and ‘different’ when there was a change). We found that accuracy across all trial types was correlated to a moderate or strong degree (location-task-angle vs. orientation-task-angle: Pearson’s $r=.63$, $p<.001$, Spearman’s $r=.61$, $p<.001$; location-task-distance vs. orientation-task-distance: Pearson’s $r=.54$, $p<.001$, Spearman’s $r=.53$, $p<.001$; location-task-angle vs. location-task-distance: Pearson’s $r=.77$, $p<.001$, Spearman’s $r=.81$, $p<.001$; orientation-task-angle vs. orientation-task-distance: Pearson’s $r=.71$, $p<.001$, Spearman’s $r=.74$, $p<.001$). In other words, participants who did well in the task tended to do well across all conditions (see Figure 3B-C).

The key question in this experiment, however, was whether there is a unique relationship in oblique biases across location and orientation judgments. Using the above metrics of oblique biases, we evaluated the same cross-task and cross-trial-type correlations we assessed before. Here, to ensure that these critical correlations were robust, we also ran bootstrapped correlations, resampling trials from each participant with replacement. The results of those bootstrapping analyses are shown as confidence intervals alongside the other correlation values. Unlike the overall accuracy correlations, we found that oblique biases were reliably correlated only for angle discriminations (location-task-angle vs. orientation-task-angle: Pearson’s $r=.43$, $p<.001$, bootstrapped 95% CI = [0.21, 0.44], Spearman’s $r=.38$, $p<.001$, bootstrapped 95% CI = [0.19, 0.43]; see Figure 3B). All other cross-task and cross-trial correlations were nonsignificant (location-task-distance vs. orientation-task-distance: Pearson’s $r=.10$, $p=.31$, bootstrapped 95%

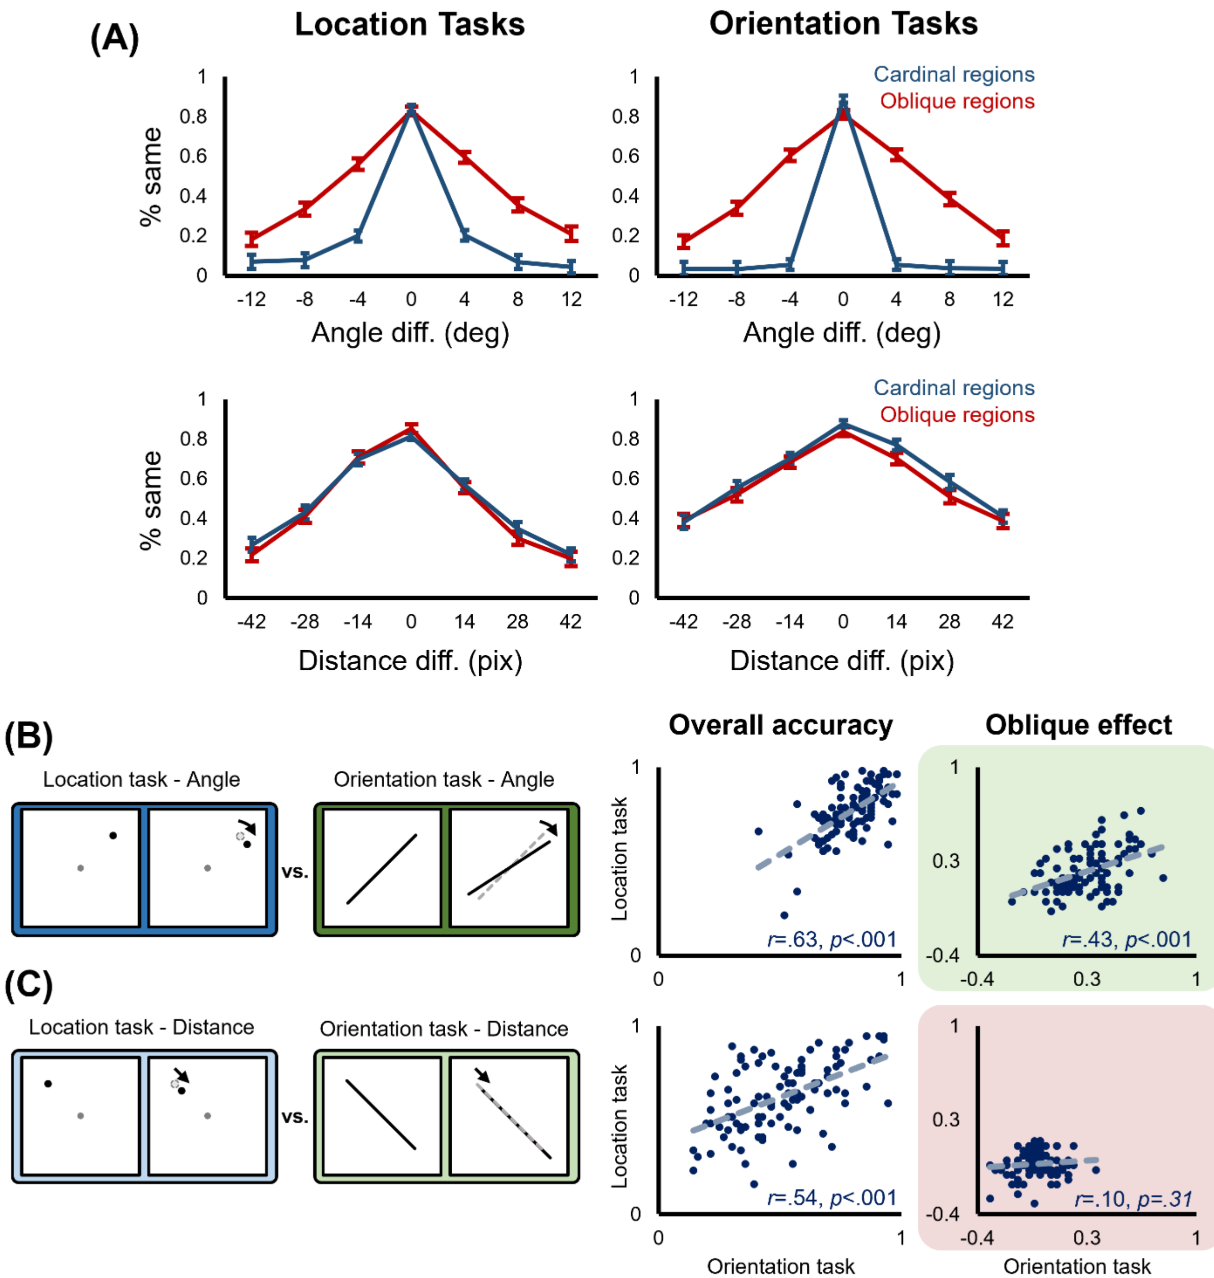

Figure 3. Design and results of Experiment 1. (A) Results for the four different task/trial-type combinations. (B) Correlations for overall accuracy and oblique biases between the location-angle discrimination trials and the orientation-angle discrimination trials. (C) Correlations for overall accuracy and oblique biases between the location-distance discrimination trials and the orientation-distance discrimination trials. The key result here is that oblique biases are correlated for angular discriminations only, indicating there are stable deficits in oblique acuity across tasks. (The depictions of the stimuli and distances between them shown here are not to scale; they are modified to increase readability of the figure.)

CI = [-0.12, 0.22], Spearman's r=.05, p=.60, bootstrapped 95% CI = [-0.14, 0.21]; location-task-angle vs. location-task-distance: Pearson's r=.16, p=.13, bootstrapped 95% CI = [-0.09, 0.25], Spearman's r=.15, p=.15, bootstrapped 95% CI = [-0.10, 0.24]; orientation-task-angle vs. orientation-task-distance: Pearson's r=-.09, p=.40, bootstrapped 95% CI = [-0.20, 0.09], Spearman's r=-.12, p=.23, bootstrapped 95% CI = [-0.21, 0.10]; see Figure 3C).

As predicted, there was a unique relationship between the oblique effects across tasks. How should we interpret this correlation? One of the critical aspects of our design was the inclusion of distance change detection as a control. These distance change trials allowed us to ask if the relationship we observed for the angular oblique effects was unique. Even though we did observe tiny oblique effects for the distance changes, those effects were unreliable, were not related across tasks, and were in the opposite direction of typical oblique effects. Moreover, oblique effects for distance *within* each task were not related to the angular oblique effects in those same tasks. Thus, the only key factor that was reliably related across tasks was the angular oblique effect. We note that if this crucial correlation was due to some general factor like effort or attention, we would expect the control distance metrics to also be correlated, but they were not. Thus, we think that this pattern reveals a unique, meaningful relationship between oblique biases across these rather different tasks.

In short, the observed correlations point to two primary conclusions: (1) There are stable individual differences in visual acuity in the oblique regions of space that are *not* specific to orientation tasks, and (2) These biases are specific to angular acuity (as opposed to a more general deficit in spatial acuity; see also Yousif et al., 2020). These results suggest that the 'oblique effect' may not be unique to orientation after all, and that, instead, the oblique effect may be a product of a more general deficit in spatial acuity near the obliques.

Experiment 2 — Oblique biases are modality-general

Other than the oblique effect itself, perhaps the second most-well-known 'oblique-related' effects are biases of spatial localization towards the oblique regions of space, away from the cardinal axes (Huttenlocher et al., 1991; Yousif et al., 2020; Wei & Stocker, 2015). Rather than being explained by differences in angular acuity, though, these biases are traditionally described as arising from a categorical bias — a tendency to place points towards the center of the 'category' (often, a quadrant of Cartesian space) in which they originated. Here, we ask whether these localization biases — like the oblique effects in the previous experiment — are more general in nature. Specifically, we ask whether they are stable across modality. Participants will complete two separate tasks: A visual localization task (in which they remember and recreate locations based on visual input) and a motor localization task (in which they remember and revisit locations based on kinesthetic input). As with the previous experiment, we are asking whether we observe oblique biases in both tasks, and, if so, whether those biases are related.

Method

This experiment consisted of two separate tasks. One was a visual location matching task in which participants saw dots briefly presented on a computer screen and then, after a delay, had to recreate the location of that dot relative to a landmark. The other was a motor (proprioceptive) location matching task in which participants were passively guided by a motorized robot to a location in space (sans any visual input) and then, after a delay, had to move the robotic arm back to that location.

Participants. 40 undergraduate students participated in exchange for course credit. Half of the participants completed the visual location matching task first, and the other half completed the motor location matching task first. Four additional participants were excluded prior to further data analysis based on pre-registered exclusion criteria (three because of their responses during the debriefing survey; one because their overall accuracy was low).

Procedure & Design. The visual location matching task was modeled after the tasks used by Yousif & Keil (2021). Participants saw a blue target dot (10 pixels in diameter) presented in a random location relative to a central grey dot (25 pixels in diameter). The dots could *not* appear further than 120 pixels away from the central grey dot, nor could they appear within 30 pixels of the central grey dot. The dots would appear on the screen for 1500ms before disappearing. After another 500ms, the grey dot would reappear in a different location and the blue dot would be absent. The participants were asked to place a new blue dot to match the location of the previous dot, relative to the current grey dot. The central grey dot would initially appear in one of the four quadrants (always 250 pixels away from the center of the screen horizontally, and 150 pixels away from the screen vertically); the grey dot would always reappear in the opposite quadrant from where it had been initially. The initial position was counterbalanced so that the grey dot appeared in each quadrant an equal number of times. Once participants had clicked a single time, a blue dot would appear. However, participants could drag and drop or click additional times to replace the blue dot as they wished. They had an unlimited amount of time to respond, although they were encouraged to respond as quickly and as accurately as possible. To submit their responses, they pressed the spacebar. There were 120 trials in total. Participants completed two representative practice trials before beginning the task.

The motor location matching task was designed to be as similar as possible to the location matching task. Participants sat at a desk in front of a robotic manipulandum (henceforth referred to as the 'robot arm'; KINArm Endpoint, Ontario Canada). The robot arm could be dragged by the participant, but it could also move autonomously (thus dragging the participants hand with it). Participants wore a black 'bib' that obfuscated their vision of the robot arm and the desk itself. However, they were able to see visuals which displayed helpful information throughout the task (e.g., signals for when they could respond, start the next trial, etc.); these minimal stimuli/prompts were reflected from a horizontally mounted LCD screen onto a semi-silvered mirror positioned below it (the mirror provided further visual occlusion, thus making the full arm and hand invisible to participants).

Each trial began with a grey dot presented centrally on the screen. During this portion of the task only, there was a small cursor (a white dot) that corresponded to the location of the participants hand on the desk below. Participants were told to move the cursor onto the central dot to begin the trial. As soon as they did this, both the central grey dot and the cursor would disappear. At this time, the robot arm would move the participant's hand to a random location in the 2D workspace. The random location could not be more than 7cm away from the center in each dimension (so that the maximum distance any point could be from the center was ~10cm), nor could it be within 3cm of the center in each dimension. The robot arm would guide the participant's hand directly to the location on each trial (this passive movement was designed to always take 1000ms), pause for 1000ms, then return to the center. After another 500ms, a green dot would appear on the screen, which signaled to participants that they could respond. Participants were instructed to move immediately and directly to the point that had been indicated by the robot. After the robot detected no significant movement (velocity <.5cm/s) for 500ms, it would register the participant's current hand position as the response on that trial. At this point, the cursor and central grey dot would reappear, and the participant could control the cursor to return to the home location and begin the next trial.

Participants were explicitly told prior to the task that they should not rely on any special strategies or heuristics to localize the points in space. Instead, they were told to rely only on their sense of space, even if it meant they were slightly less accurate. This was done to prevent participants from surreptitiously using strategies like placing their arm against the table or pressing it against their body and trying to remember how their arm had been positioned, rather than the locations themselves. As with the visual location matching task, participants completed 120 trials. They completed 8 representative practice trials before beginning the task, during which they were given extensive verbal feedback (about the task itself, not their accuracy) to ensure that they understood the task.

Results and Discussion

The full data set for each task is displayed in Figure 4. As is evident from the figure, there were robust oblique biases that resemble those observed in prior work (e.g., Huttenlocher et al., 1991; Yousif et al., 2020). There are many ways to quantify these biases. One simple metric is to simply count all the trials in which participants erred towards the oblique axis vs. towards the cardinal axis. For the visual localization task, an average of 72% of trials (SD=.07) moved towards the oblique axes, $t(39)=20.07$, $p<.001$, $d=3.25$. For the motor localization task, an average of 59% of trials (SD=.07) moved towards the oblique axes, $t(39)=8.14$, $p<.001$, $d=1.29$. We can also quantify the magnitude of these biases: Are errors that move towards the oblique axes *larger* than errors that move towards the cardinal axes? For the visual localization task, the errors towards the oblique axes were an additional 3.91 degrees larger on average (points moving toward oblique: M=8.81deg, SD=2.29deg; points moving toward cardinal: M=4.91deg, SD=1.74deg; $t(39)=14.71$ $p<.001$, $d=2.33$). For the motor localization task, the errors towards the oblique axes were an additional 1.37 degrees larger on average (points toward oblique: M=6.39deg, SD=1.23deg; points toward cardinal: M=5.02deg, SD=1.33deg; $t(39)=6.66$, $p<.001$, $d=1.05$. These analyses confirm what is evident from Figure 4: Participants exhibited a robust tendency to err towards the oblique axes. For the remainder of this section, we'll refer to this analysis as 'differences-by-error-direction'.

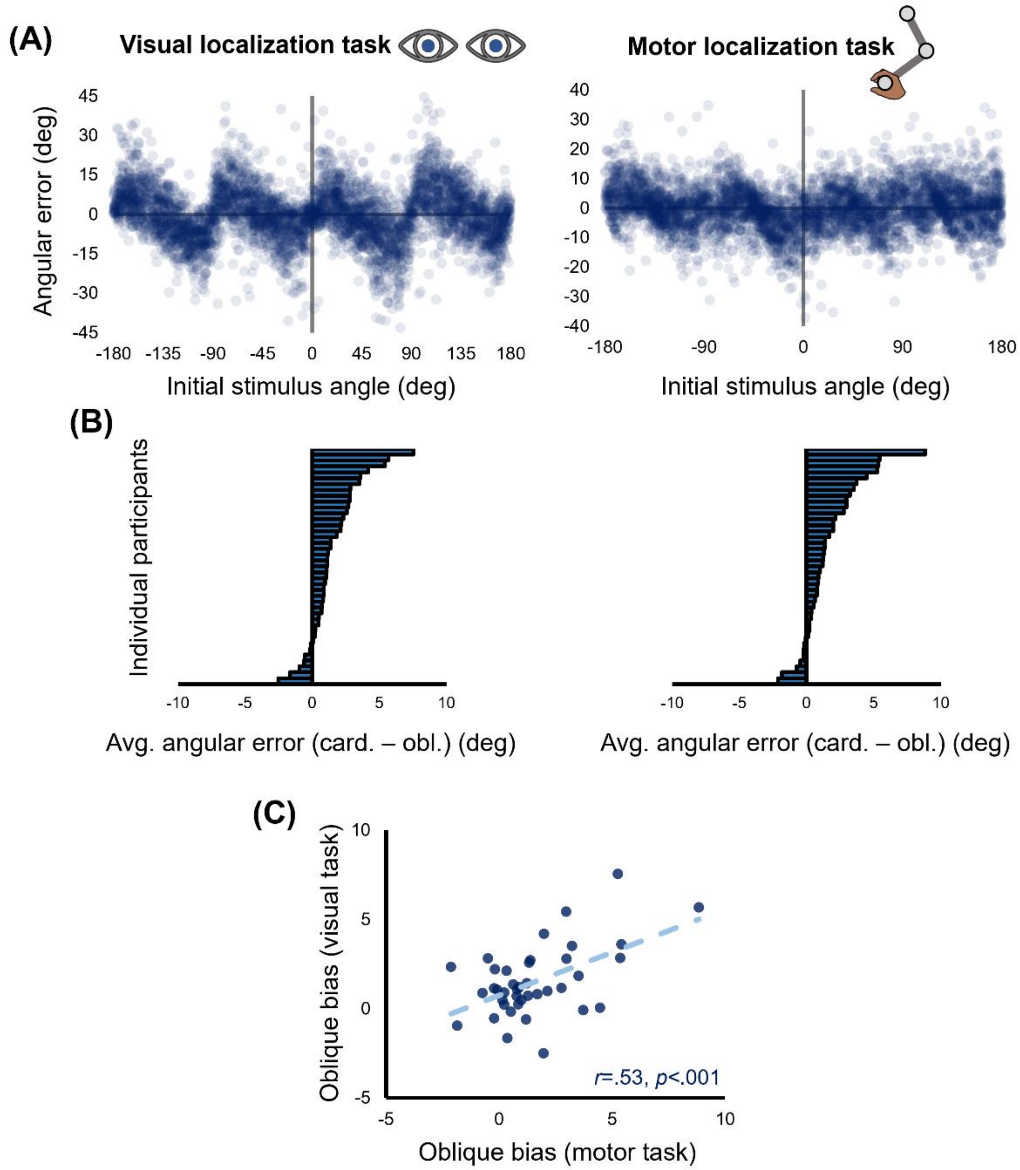

Figure 4. Results of Experiment 2. (A) Angular error as a function of initial angular position. (B) Oblique biases, quantified as the difference in angular error for points that originated near the cardinal axes vs. the oblique axes. In other words, we took all the trials with points that originated closer to the cardinal axes and calculated the average absolute angular error for those points; then we did the same for all the trials with points that originated closer to the cardinal axes. This x-axis here reflects the difference between those two values, for each participant. (C) The correlation between the oblique biases in (B).

Separately, we quantified the magnitude of angular errors for points that originated near the cardinal axes vs. those that originated near the oblique axes (unlike the previous analysis, which was based on where points erred towards, not where they originated). For the remainder of this section, we'll refer to this analysis as 'differences-by-origin-point'. For the visual localization task, errors were on average 1.26 degrees larger for points that originated near the cardinal axes, $t(39)=4.40$, $p<.001$, $d=.70$; for the motor localization task, errors were on average

1.19 degrees larger for points that originated near the cardinal axes, $t(39)=6.19$ $p<.001$, $d=.98$. Combined with the previous analysis, these results suggest that points originating near the cardinal axes (1) tend to move towards the oblique axes and (2) tend to move farther than points which had originated near the oblique axes.

While analyzing the data, we realized that our preregistered exclusion criteria may result in bias. Insofar as there are larger errors for certain trial types (e.g., those that move towards the oblique axis), more trials of that type would be excluded. This is true in practice: In the visual task, for instance, over 80% of the trials that would be excluded per the accuracy exclusion are points that erred towards the oblique axes. While there are some genuine outliers (i.e., trials on which observers appeared to forget the location entirely), many of the presupposed 'outliers' seemed to reflect large oblique biases. To account for this unusual source of bias, all of the analyses reported so far were conducted with and without these exclusion criteria. The results were identical. That said, correlation analyses (reported below) are more sensitive. Because we did not want to disproportionately exclude trials of one type or the other, we ran the correlation analyses first without exclusion. To ensure that these results were not primarily driven by a subset of trials, and consistent with the approach taken in Experiment 1, we ran bootstrapped correlations on the relevant values, resampling trials from each participant with replacement. The confidence intervals for these bootstrapped correlations are displayed alongside the other values.

Are the oblique biases we observed in each task related to one another? Surprisingly, there was no correlation between the magnitude of errors that moved towards the oblique axes (Pearson's $r=.07$, $p=.68$, 95% CI = [-0.33, 0.38]; Spearman's $r=.09$, $p=.55$, 95% CI = [-0.25, 0.32]). Crucially, however, there was a significant correlation between the magnitudes of errors that originated near one axis vs. the other (i.e., the difference in angular error between points that originated near the cardinal axes vs. near the oblique axes; Pearson's $r=.53$, $p<.001$, 95% CI = [0.06, 0.57]; Spearman's $r=.39$, $p=.014$, 95% CI = [0.03, 0.49]). We note that the differences-by-origin-point are more comparable to the sort of effect that we observed in Experiment 1, again suggesting a shared source for the oblique effect.

Consider what it means to observe any correlation between these tasks: The values being correlated here are *differences* in angular accuracy between two different regions of space, in two different modalities and in two different spatial planes (vertical in the visual task, horizontal in the motor task). This means that participants that happen to make larger errors near the cardinal axes in a visual localization task also happen to make larger errors near the cardinal axes in a completely nonvisual proprioceptive/motor localization task. This relation cannot be parsimoniously explained by purely visual or purely motor biases alone. It also cannot be explained by general inattention or inaccuracy, as there is no reason that errors due to attention or low effort should be localized to specific regions of space (except if a genuine oblique bias exists, as we propose). Thus, these results appear to reflect an oblique bias that arises from a modality-general system of spatial representation.

Unlike those in Experiment 1, the effects observed here are not about acuity *per se*. Although both effects (i.e., the acuity differences observed in Experiment 1, and the localization biases observed in this experiment) involve some difference between the cardinal and oblique regions of space, the effects themselves may not be related to one another. In the following section, we consider how these two distinct effects may be related to one another.

Interim Discussion

So far, we have described two distinct oblique-related effects. In Experiment 1, we showed that there are deficits in angular acuity in the oblique regions of space that are not specific to orientation. In Experiment 2, we showed that people mis-localize points *towards* the oblique regions of space in both visual and motor localization tasks. And in both cases, we showed that biases across tasks (Experiment 1) or across modalities (Experiment 2) were related to one another. Combined, these results imply a relation amongst all these effects.

How could it be possible, in computational terms, for all these biases to share a common representational format? Our suggestion is simple: One way or another, at some level(s) of representation, there is a gradient in angular acuity such that acuity is highest near the cardinal axes of space and lower near the oblique axes of space. That gradient is the direct cause of effects like the classic oblique effect, or the location effects observed in Experiment 1. Straightforwardly, observers are better at detecting changes near the cardinal axes because acuity is higher near the cardinal axes. (Note here that we are not making a claim as to why acuity is higher near the cardinal axes, only observing that this is the case.)

The trickier question is: How would biases in angular visual acuity might give rise to localization biases towards the oblique axes (e.g., Huttenlocher et al., 1991; Yousif et al., 2020), which have traditionally been explained by appeal to higher-level categorization? To answer this question, we draw on the model of Wei and Stocker (2015) which shows how acuity deficits and efficient coding principles can indeed lead to these exact sorts of localization biases. Under this model, the effects observed in Experiments 2 would directly fall out of acuity differences, such as those observed in Experiment 1. In other words, acuity differences are the likely root cause of the localization biases observed in Experiment 2. If true, this would mean that we would no longer need a higher-level explanation (i.e., categorical biases, or 'prototypes') to explain these effects.

Such biases need not be 'built in' to the visual system. They could arise naturally from experience. For instance, if the visual system more frequently takes horizontal and vertical orientations as input, it could become more sensitive to those orientations. Thus, the visual system could learn to 'efficiently encode' information in a way that would result in the exact sorts of oblique biases we have described here (see, e.g., Benjamin et al., 2022). If this is true — that these effects arise from the input of the natural environment — it remains unclear why these biases would manifest in modalities separate from that input. In other words: It is not clear why biased input in the visual modality would result in kinesthetic oblique biases. Yet this raises the intriguing possibility that visual input shapes spatial representations even beyond the visual modality.

Can we use this perspective to make novel predictions? Here we have argued that several 'oblique-related' effects may share a common basis — that region-specific deficits in angular acuity may be the root cause of many known effects. If this is true, we should expect to observe other sorts of biases in the oblique regions of space. For instance: Might perception of the space itself be distorted in some way? In a final pair of experiments, we address this question.

Experiment 3a — Oblique biases explain other perceptual phenomenon

What does it mean to have variation in angular acuity in different regions of space? Most existing work focuses on how differences in angular acuity manifest as differences in precision (i.e., a decreased ability to discriminate oriented lines or points in specific regions of space). Yet if these differences are truly general differences in angular acuity, one may expect that these differences would result in other perceptual distortions. Here, we ask: Is it possible that the empty space of the oblique regions itself could be distorted? Specifically, we ask if observers perceive the empty space between two oriented lines in the oblique regions as larger or smaller in magnitude than equivalent empty spaces in the cardinal regions (for a summary of the design, see Figure 5).

Method

This experiment was identical to Experiment 1, except as stated below.

Participants. 100 participants were recruited via Prolific.

Stimuli. The stimuli consisted of angles subtended by two oriented lines, each of which originated in the center of the display, was 2 pixels wide, and was 200 pixels in length. One of the two angles was always presented centered on a cardinal axis (0, 90, 180, or 270 degrees, randomly selected); the other was always presented centered on an oblique axis (45, 135, 225, or 315 degrees, randomly selected). The size of the first angle was always either 20, 30, or 40 degrees. The size of the second angle was always the size of the first angle +/- 3, 5, or 7 degrees.

Procedure & Design. There were 2 starting orientations (cardinal, oblique) x 3 base angle sizes (20, 30, 40 degrees) x 3 possible size increments (3, 5, 7 degrees) x 2 directions (size increases, size decreases), resulting in 36 unique trial types. Participants completed each trial type 4 separate times, resulting in a total of 144 trials. However, note that these 4 repeated trial types were not necessarily identical, as the specific axis that was chosen within the set of possible cardinal or oblique axes was random. All trials were completed within a single block. There were two representative practice trials before the beginning of the task. All other aspects of the design were identical to Experiment 1.

Results and Discussion

The primary question of this experiment is whether empty spaces in oblique regions are perceived as smaller or larger than angles in cardinal regions. To that end, we conducted a repeated measures ANOVA with two factors (3 levels of angular change and 2 trial types [cardinal angle larger, oblique angle larger]). There was a main effect of angular change such that larger angles were better discriminated ($F[2,98] = 125.26$, $p<.001$, $\eta_p 2=0.72$), a main effect of trial type such that trials in which the cardinal angle was greater were better discriminated ($F[1,49] = 76.86$, $p<.001$, $\eta_p 2=0.61$), as well as an interaction between the two ($F[2,98] = 17.00$,

$p<.001$, $\eta_p 2=0.26$; see Figure 5). Moreover, this main effect of trial type — the critical result in this experiment — was consistent across all three base angle sizes (i.e., whether the initial angle was 20, 30, or 40 degrees; all ps<.001). This main effect indicates that people generally perceive the empty space at the cardinals as larger than at the obliques. Thus, even the perception of empty space may be influenced by the sort of oblique biases we are studying here.

Experiment 3b — Angle size differences are not a function of line orientation biases

The results of Experiment 3a may imply a difference in perceived space in the oblique regions, as we originally hypothesized. However, there may be a simpler explanation: Because of the nature of the design in Experiment 3a, it could be that the results are driven not by the perception of the space itself but by the percept of the oriented lines that form the angles in the first place. If lines near the oblique regions are perceived as closer to the oblique axes, for instance, this could cause those angles to be perceived as subtending a smaller area. Here, we address this possibility by using larger angles. That is, in our new experiment the angles centered on a given axis were made up of oriented lines that are closer to the orthogonal axes, eliminating the possibility of an oblique effect driven by the lines themselves. If the empty spaces centered on the oblique regions are still perceived as smaller (i.e., replicating Experiment 3a), the phenomenon cannot be explained by the orientations of the individual lines (indeed, a line orientation bias here would generate an effect opposite to what was observed in the previous experiment).

Method

This experiment was almost identical to Experiment 3a, with one notable difference: The base angle sizes were changed from 20, 30, and 40 degrees to 80, 90, and 100 degrees. The purpose of this change was to de-confound angle size and the axes with which the constituent lines were colinear (or near-colinear). 100 new participants were recruited via Prolific.

Results and Discussion

As with the previous experiment, we conducted a repeated measures ANOVA with two factors (3 levels of angular change and 2 trial types [cardinal angle larger, oblique angle larger]). There was a main effect of angular change such that larger angles were better discriminated ($F[2,98] = 45.79$, $p<.001$, $\eta_p 2=0.48$), a main effect of trial type such that trials in which the cardinal angle was greater were better discriminated ($F[1,49] = 12.76$, $p<.001$, $\eta_p 2=0.21$), but not interaction between the two ($F[2,98] = .16$, $p=.85$, $\eta_p 2=0.00$; see Figure 5). However, the main effect of trial type — unlike the previous experiment — was not consistent across all three base angle sizes (i.e., whether the initial angle was 80, 90, or 100 degrees). Observers were more accurate when the cardinal angle was larger for the 80 degree trials ($p<.001$) and the 90 degree trials ($p=.009$), but more accurate when the oblique angle was larger for the 100 degree trials ($p=.019$). Note that although this p-value is below .05, this value is *not* statistically significant after accounting for Bonferroni correction. Nevertheless, it may be interesting that this effect did not go in the opposite direction, as was the case with the other 5 angle sizes that were tested across both angle judgment experiments. We are not sure why this is the case; but, we think it remains clear that empty spaces spanning the cardinal regions are almost always perceived as larger than equivalent empty spaces spanning the oblique regions.

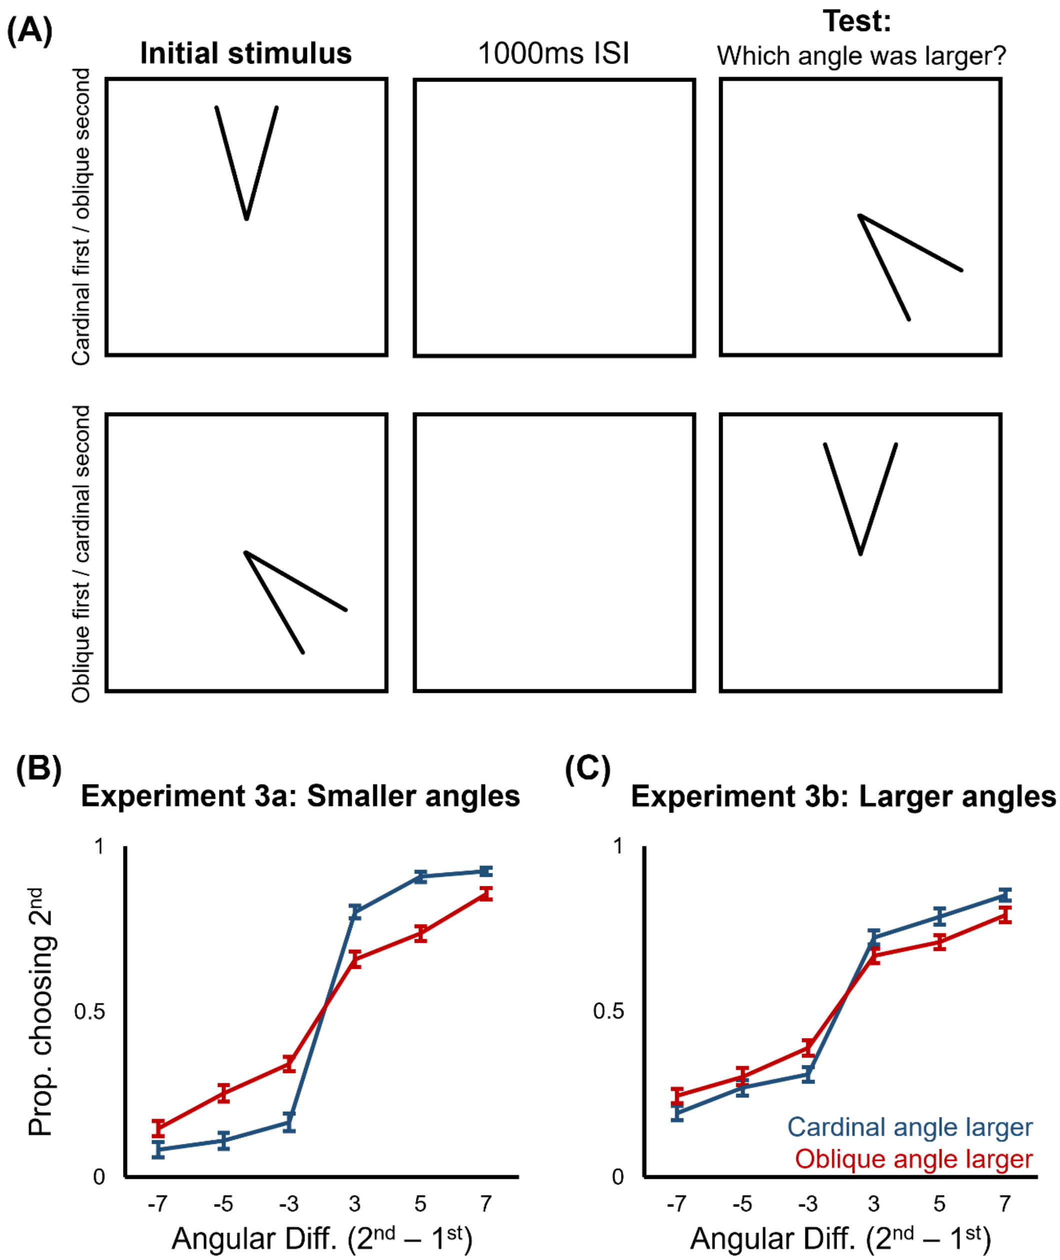

Figure 5. Design and results of Experiment 3a and 3b. (A) A schematic of the design. (B) Proportion of responses selecting the second angle for each of the two trial types, broken down by angular difference, for Experiment 3a. (C) Proportion of responses selecting the second angle for each of the two trial types, broken down by angular difference, for Experiment 3b. The depictions of the stimuli shown here are not to scale; they are modified to increase readability of the figure.

General Discussion

Here, we have proposed that many varieties of oblique biases can all be explained by a shared, underlying spatial representation. Whereas prior work has proposed explanation like cognitive categorical biases (Huttenlocher et al., 1991), kinesthetic limitations (Gordon et al., 1995; but see Baud-Bovy & Viviani, 2004), or 'gravitational cues' (Gentaz & Hatwell, 1995), here we argue that nearly all these effects may share a common cause — a deficit in angular acuity in the oblique regions of space that is not specific to any modality or domain. Moreover, we have shown that this account can explain new biases that have not been observed before — namely, that even the perception of empty space varies depending on whether that space spans a cardinal or oblique region (Experiments 3a and 3b).

What does it mean for these effects to share a common format?

Some of the effects we have discussed here seem obviously related; indeed, some of them — the 'oblique effect', the 'haptic oblique effect', and the 'motor oblique effect' — share a common name. Thus, it seems relatively uncontroversial to say that these biases share a common basis. Some of these effects have long been understood in radically different ways, however. Perhaps the best example of this are the well-known spatial localization biases, whereby people misplace objects closer to the oblique axes than they really were (Huttenlocher et al., 1991; Yousif et al., 2020). These effects, while obviously reminiscent of the oblique effect in some way, have been explained by appeal to a *cognitive* bias, not a perceptual one. Huttenlocher and colleagues famously argued that localization errors result from biases of categorization, a 'coarse' representation of location that is biased towards the 'prototype' of the initial category. The view presented here does not present evidence against that explanation — it continues to be plausible that categorical biases of spatial localization exist — but does offer an alternative way of understanding these localization biases. Specifically, we argue that it is possible that these biases arise not only from discrete, categorical biases but instead from continuous variation in angular acuity. In practice, this means that the same system for spatial representation that biases your visual impression of an oriented line (e.g., the oblique effect) may also be responsible for biasing where you remember something being positioned in space (e.g., prototype effects).

Perhaps even more striking is the fact that these effects span multiple modalities, including vision (Huttenlocher et al., 1991; Yousif et al., 2020), proprioception (Gentaz & Hatwell, 1995), and action (e.g., Baud-Bovy & Viviani, 2004; Gordon et al., 1995; Smyrnis et al., 2007). The connection between these tasks is more than conjecture: In Experiment 1, we observed robust correlations between oblique biases in two distinct domains (i.e., location, orientation). And in Experiment 2, we observe robust correlations between oblique biases in two distinct modalities (i.e., vision, action). The consistency in these biases across disparate contexts opens the door to a provocative conclusion: that beneath these wide range of situations is a single shared representation for representing spatial information. As obvious as this conclusion may seem when stated this way, it is important to remember how differently many of these phenomena have been explained historically. And while others have speculated about a connection between visual and motor effects before (see, e.g., Baud-Bovy & Viviani, 2004), this is the first work to our knowledge to actually demonstrate direct relationships among these disparate modalities (and across domains as well, see Experiment 1).

Putting this all together: We propose that the well-known, thought-to-be-well-understood oblique effect is *neither* an effect only of vision *nor* an effect only of orientation (despite classic explanations that appeal to both vision and orientation, e.g., Furmanski & Engel, 2000; Li et al., 2003). Moreover, well-known localization biases (i.e., 'prototype effects' or 'Category Adjustment Model' effects) are also neither about vision nor about localization. Likewise for the haptic and motor oblique effects. All of these biases may instead reflect a deeply *spatial* phenomenon — one that transcends domain and modality.

Does this mean that canonical explanations for the oblique effect appealing to orientation-selective neurons (e.g., Furmanski & Engel, 2000; Li et al., 2003) are incorrect? Does this mean that 'prototype effects' (Huttenlocher et al., 1991) and the 'Category Adjustment Model' (Holden et al., 2010; Holden et al., 2013) are misunderstood? We think not. While we think the present results suggest that these biases *can* arise from a single spatial representation, we must point out that all of these other explanations could apply in addition to the proposed model. For instance, it is possible that points are biased towards the 'prototype' of a category, and that, in addition to this categorical effect, there is an additional effect of angular acuity, as we have proposed here. (A detailed analysis of these errors may provide further evidence against the categorical model, but we nevertheless wouldn't rule it out entirely; see Yousif et al., 2020; Yousif & Keil, 2021; Yousif, 2022.) Further, it is likely that there are truly more orientation-selective neurons that are tuned to cardinal vs. oblique orientations (Li et al., 2003) and that this difference results in acuity differences. But then a mystery remains: Why would all these related effects — many of which are stable within individual — arise from distinct mechanisms?

Whether the *source* of angular distortions in this general-purpose spatial representation is originally visual (or haptic, or even genetic), and/or directly driven by environmental statistics (e.g., Keil & Cristobal, 2000; Girshick, Landy, & Simoncelli, 2011; Henderson & Serences, 2021; Wei & Stocker, 2015), are intriguing open questions. However, even if the source of these distortions is domain- or modality-specific (e.g., such as coding anisotropies in the primary visual cortex; Li et al., 2003), we contend that these distortions might ultimately shape a higher-level, more abstract representation of space. We think it would be revealing to find that, for instance, some lower-level bias of orientation perception somehow has the downstream effect of distorting other spatial representations, including those that are non-visual and not specific to orientation. For now, however, it remains unclear how so many biases arrange from a shared spatial representation; there is no account of visual processing, to our knowledge, that would predict an early visual process like orientation perception to have such cascading effects.

Other accounts of related phenomena

In addition to the work discussed so far, there is one other recent paper that offers a general account of spatial biases: Based on localization errors in a serial reproduction task (in which one participant's output is presented to another participant as input, much like the game of telephone), Langlois and colleagues (2021) argue that spatial errors are biased towards the regions of an image which are represented with the *highest* acuity. This is of course at odds with perhaps the most famous spatial bias of all (i.e., prototype effects; Huttenlocher et al., 1991), which involves mis-localizations towards the regions of *lowest* acuity (i.e., the obliques; see Yousif et al., 2020; Wei & Stocker, 2015). Could both things be true at the same time? How would the current data be explained by Langlois and colleagues (2021)?

The short answer is that we think that Langlois and colleagues are talking about a different kind of bias than what we have studied here. The phenomena studied and discussed throughout this paper are ones that occur in the absence of any sort of landmark. In the classic work of Huttenlocher and colleagues (1991), participants were simply tasked with remembering the location of a dot with respect to a larger circle. But these exact same sorts of biases emerge even when participants localize a dot relative to a single other dot (Yousif et al., 2020) in the absence of any other visual information that could be used to guide the judgment. This is in stark contrast to the stimuli used by Langlois and colleagues, which consist entirely of naturalistic images (e.g., of a plane, a lighthouse, or a face). This is tantamount to the difference between navigating in an open field vs. in a dense city. When navigating in a city — with copious landmarks and clearly labeled streets — people will call on all of the available information to localize things in space. When giving directions in cities, for instance, people will frequently say things like, "Go over to 24th then up to Spruce past the grocery store, then turn right." But there are not landmarks or street names in a corn field. The sort of spatial representation we use to navigate in complex environments (i.e., a form of representation that depends on propositional knowledge of the environment) is very different from the sort of spatial representation that we use to navigate in more sparse environments (i.e., a form of representation that is influenced by perceptual input, independent of propositional knowledge as much as possible). So it is with the sorts of localizations considered here and by Langlois and colleagues (2021). We are here interested in the latter kind of representations — ones that arise from sparse input.

In other words, we think that it is entirely possible that there are localization biases towards regions of higher acuity in some cases. These are not the biases at-issue in our study. We think those biases may reflect explicit strategies, or even propositional representations of space. The biases we have studied here, in contrast, reflect a foundational system of representation — one that generalizes across tasks and modalities and is invoked even in cases of extremely minimal sensory input.

Essock (1980) proposed a distinction between "class 1 effects" and "class 2 effects" for oblique biases. Class 1 effects are said to reflect basic processing of the visual system; Class 2 effects, in contrast, are said to reflect higher-level encoding. On this view, the classic orientation oblique effect (e.g., Appelle, 1972) would likely be classified as Class 1, whereas localization biases (e.g., Huttenlocher et al., 1991) would likely be classified as Class 2. Some studies have even argued that there are separate, simultaneous repulsion effects from the cardinal axes and attraction effects toward the oblique axes (which can be teased apart by manipulating frames of references; see Rademaker et al., 2017). Here, however, we have argued that most if not all of these biases may arise from a single representational distortion. Classic prototype effects could in principle arise from a gradient in visual acuity (see Wei & Stocker, 2015), as could the distortions of perceived space in Experiments 3a and 3b. The biases described by Langlois and colleagues (2021) seem unlike all of these effects in that they involve a fundamentally opposing bias — one towards the regions of high acuity rather than away from it. It is yet unclear whether, or how, these biases relate to one another. For now, we find it helpful to think of the biases towards regions of high acuity as a more specialized bias, one that arises only in the presence of complex input.

The format of spatial representation(s)

Although the evidence is indirect, it seems noteworthy that the acuity differences and biases we observe are not just about one region of space versus another, but also about one *dimension* of space versus another. The biases we observe are specific to *angular* acuity. This fact alone has some surprising implications. For instance, it means that the classic 'prototype effects' (Huttenlocher et al., 1991) may be conceived not just as biases towards a point in space, but as biases towards an axis of space along a single dimension. It also forces the conclusion that angular information is being represented independently from other dimensions on some level. Because of this, it may follow that the mind is likely to be representing spatial information in some sort of polar coordinate system. Indeed, analyses of errors like those studied here have revealed that polar error are independent from distance errors (while errors in cartesian dimensions are not independent from one another), lending further support for this conclusion (see Yousif & Keil, 2021; Yousif, 2022).

We emphasize this point insofar as it may help to tell a broader story. Perhaps, as prior work suggests, the underlying form of spatial representation(s) is some form of polar coordinate system. And perhaps within that polar coordinate system, the angular dimension is distorted. In other words, it is possible that the effects observed here reflect an underlying representational format that extends beyond merely the consideration of oblique biases. Thus, the findings here can contribute not only to our understanding of oblique biases, but also to the foundations of spatial representation in general. While this connection is speculative, the promise is clear — namely, that an assortment of spatial phenomena can be understood not as distinct effects, but as related through some common underlying representational format.

Conclusion

What do oblique effects in orientation judgments, pointing errors, visual memory errors, and angle-size judgments all have in common? While prior work has offered many different domain- and modality-specific explanations for these phenomena, ranging from cognitive biases to physical limitations, we suggest that they may all boil down to a single representational distortion: deficits in angular acuity in the oblique regions of space. This account is supported by numerous replications of 'oblique related' effects across tasks and paradigms, as well as robust cross-task and cross-modality correlations. Collectively, these findings hint that beneath a wide range of observed phenomena exists a general, flexible, shared system of spatial representation.

References

Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: the" oblique effect" in man and animals. *Psychological Bulletin*, *78*, 266-278.

Baud-Bovy, G., & Viviani, P. (2004). Amplitude and direction errors in kinesthetic pointing. *Experimental Brain Research*, *157*, 197-214.

Benjamin, A. S., Zhang, L. Q., Qiu, C., Stocker, A. A., & Kording, K. P. (2022). Efficient neural codes naturally emerge through gradient descent learning. *Nature Communications*, *13*, 1-12.

Bonds, A. B. (1982). An "oblique effect" in the visual evoked potential of the cat. *Experimental Brain Research*, *46*, 151-154.

Cecala, A. J., & Garner, W. R. (1986). Internal frame of reference as a determinant of the oblique effect. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 314-323.

Essock, E. A. (1980). The oblique effect of stimulus identification considered with respect to two classes of oblique effects. *Perception*, *9*, 37-46.

Furmanski, C. S., & Engel, S. A. (2000). An oblique effect in human primary visual cortex. *Nature Neuroscience*, *3*, 535-536.

Gentaz, E., & Hatwell, Y. (1995). The haptic 'oblique effect' in children's and adults' perception of orientation. *Perception*, *24*, 631-646.

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*, 926-932.

Gordon, J. A. M. E. S., Ghilardi, M. F., & Ghez, C. (1995). Impairments of reaching movements in patients without proprioception. I. Spatial errors. *Journal of Neurophysiology*, *73*, 347-360.

Gourtzelidis, P., Smyrnis, N., Evdokimidis, I., & Balogh, A. (2001). Systematic errors of planar arm movements provide evidence for space categorization effects and interaction of multiple frames of reference. *Experimental Brain Research*, *139*, 59-69.

Henderson, M., & Serences, J. T. (2021). Biased orientation representations can be explained by experience with nonuniform training set statistics. *Journal of Vision*, *21*, 1-22.

Holden, M. P., Curby, K. M., Newcombe, N. S., & Shipley, T. F. (2010). A category adjustment approach to memory for spatial location in natural scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 590-604.

Holden, M. P., Newcombe, N. S., & Shipley, T. F. (2013). Location memory in the real world: Category adjustment effects in 3-dimensional space. *Cognition*, *128*, 45-55.

Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: prototype effects in estimating spatial location. *Psychological Review*, *98*, 352-376.

Keil, M. S., & Cristóbal, G. (2000). Separating the chaff from the wheat: possible origins of the oblique effect. *JOSA A*, *17*, 697-710.

Langlois, T. A., Jacoby, N., Suchow, J. W., & Griffiths, T. L. (2021). Serial reproduction reveals the geometry of visuospatial representations. *Proceedings of the National Academy of Sciences*, *118*, e2012938118.

Latto, R., & Russell-Duff, K. (2002). An oblique effect in the selection of line orientation by twentieth century painters. *Empirical Studies of the Arts*, *20*, 49-60.

Latto, R., Brain, D., & Kelly, B. (2000). An oblique effect in aesthetics: Homage to Mondrian (1872–1944). *Perception*, *29*, 981-987.

Li, B., Peterson, M. R., & Freeman, R. D. (2003). Oblique effect: a neural basis in the visual cortex. *Journal of neurophysiology*, *90*(1), 204-217.

Luyat, M., & Gentaz, E. (2002). Body tilt effect on the reproduction of orientations: studies on the visual oblique effect and subjective orientations. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 1002-1011.

Luyat, M., Gentaz, E., Corte, T. R., & Guerraz, M. (2001). Reference frames and haptic perception of orientation: Body and head tilt effects on the oblique effect. *Perception & Psychophysics*, *63*, 541-554.

Luyat, M., Mobarek, S., Leconte, C., & Gentaz, E. (2005). The plasticity of gravitational reference frame and the subjective vertical: peripheral visual information affects the oblique effect. *Neuroscience Letters*, *385*, 215-219.

Mantas, A., Evdokimidis, I., & Smyrnis, N. (2008). Perception action interaction: the oblique effect in the evolving trajectory of arm pointing movements. *Experimental Brain Research*, *184*, 605-616.

Nasr, S., & Tootell, R. B. (2012). A cardinal orientation bias in scene-selective visual cortex. *Journal of Neuroscience*, *32*, 14921-14926.

Petersik, J. T., & Pantle, A. J. (1982). The oblique effect in a mirror-tracing task. *Bulletin of the Psychonomic Society*, *20*, 69-71.

Plumhoff, J. E., & Schirillo, J. A. (2009). Mondrian, eye movements, and the oblique effect. *Perception*, *38*, 719-731.

Rademaker, R., Chunharas, C., Mamassian, P., & Serences, J. (2017). Dissociable biases in orientation recall: The oblique effect follows retinal coordinates, while repulsion from cardinal follows real-world coordinates. *Journal of Vision*, *17*, 107-107.

Sainburg, R. L., Ghilardi, M. F., Poizner, H. & Ghez, C. (1995). Control of limb dynamics in normal subjects and patients without proprioception. *Journal of Neurophysiology*, *73*, 820-835.

Smyrnis, N., Gourtzelidis, P., & Evdokimidis, I. (2000). A systematic directional error in 2-D arm movements increases with increasing delay between visual target presentation and movement execution. *Experimental Brain Research*, *131*, 111-120.

Smyrnis, N., Mantas, A., & Evdokimidis, I. (2007). "Motor oblique effect": perceptual direction discrimination and pointing to memorized visual targets share the same preference for cardinal orientations. *Journal of Neurophysiology*, *97*, 1068-1077.

Vogels, R., & Orban, G. A. (1985). The effect of practice on the oblique effect in line orientation judgments. *Vision Research*, *25*, 1679-1687.

Wei, X. X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, *18*, 1509-1517.

Yousif, S. R. (2022). Redundancy and reducibility in the formats of spatial representations. *Perspectives on Psychological Science*, 17456916221077115.

Yousif, S. R., & Keil, F. C. (2021). The shape of space: Evidence for spontaneous but flexible use of polar coordinates in visuospatial representations. *Psychological Science*, *32*, 573-586.

Yousif, S. R., Chen, Y. C., & Scholl, B. J. (2020). Systematic angular biases in the representation of visual space. *Attention, Perception, & Psychophysics*, *82*, 3124-3143.

Youssef, J., Juravle, G., Youssef, L., Woods, A., & Spence, C. (2015). Aesthetic plating: a preference for oblique lines ascending to the right. *Flavour*, *4*, 1-10.