Investigating the Limits of Allocentric Memory

Edward Heywood-Everett

Doctor of Philosophy

University of York

Psychology

September 2021

Abstract

There are at least two distinct ways in which the brain encodes spatial information. In egocentric representations locations are encoded relative to the observer, whereas in allocentric representations locations are encoded relative to the environment. Both inform spatial memory, but their contributions to behaviour are not fully understood. Each system has specific advantages and disadvantages for different tasks, and these strengths and weaknesses relate to fundamental characteristics of the underlying representation.

This thesis uses a novel method (developed in Chapter 2), combining approaches from spatial memory research and psychophysics to measure spatial precision in change detection tasks where the observer's viewpoint changes between presentation and testing (viewpoint-independent memory, which relies more on allocentric representation).

Chapter 3 uses these methods to investigate the effect of parametric changes in viewpoint on spatial change detection thresholds. A monotonic but non-linear effect of viewpoint on precision was found, consistent with a preregistered model that shows how the precision of spatial memory changes lawfully as a function of viewpoint shift. The model separately quantifies viewpoint-dependent and -independent parameters reflecting the way ego- and allocentric representations combine to determine performance.

Chapter 4 builds on these results to investigate spatial memory precision with regards to changes in the scale of the stimulus and environment. In both viewpoint-dependent and -independent memory, precision is found to scale with the extent of the stimulus in the observer's field of view rather than its absolute dimensions. This finding suggests that egocentric encoding plays a part in limiting the precision of viewpoint-independent memory.

Chapter 5 investigates the limits of capacity in viewpoint-dependent and -independent memory with two distinct tasks. Here, working memory-like capacity limits determine how many items can be retained in both viewpoint-dependent and -independent spatial memory with some indications that these limits are distinct, perhaps due to additional task-specific demands.

Overall, these studies highlight the way that ego- and allocentric systems interact to determine the limits of spatial memory.

Contents

List of Figures

List of Tables

List of Equations

Acknowledgements

Tom. Thank you so much giving me this opportunity. I could not have asked for a better supervisor. I have enjoyed the project from start to finish, have developed so many skills and it is in such large part due to your wonderful supervision and guidance. I am deeply grateful for the time you have dedicated to me. I feel honoured to have apprenticed under you and lucky to have been able to spend the last 4 years working with you. A massive thank you for the patience, support and inspiration you provided.

Thanks to my TAP and progression panel members, Aidan Horner, Beth Jefferies and Karla Evans. I have enjoyed our meetings since the first. Sincere thanks for the impactful conversation, the thoughtful and constructive questions and all the support you have given me through some personally challenging times. A special thank you to Daniel Baker for the invaluable expertise and insight you provided.

Thanks also to the wonderful people I met in York during this fantastic time. Especially to Bryony, Nicole, Mandy, Joanna, Charlotte, Leah, Yujin, and Bardur for your friendship.

To my good friends, David, Joe, Elliot and Toby. Thank you for letting me practice my presentations on you, thank you for completing many hours of boring pilot studies, thanks for the tasty meals, the board games, and the weirdness. You guys keep me happy.

In part, I owe this to my family. To my Mum, the best Mum, who I can always rely on, thanks for always being there for me. To my Dad, my favourite conversation partner who inspires me to be interested in everything. To my brother Samuel, the quickest and funniest person I know. To my Nanny, who taught me to be kind and generous and to my Grandad who taught me to love maths and science. To Uncle Andrew- I wish I could have shared this with you and I am thinking of you, thank you for teaching me to be passionate about my hobbies. Thanks to Karen for the love and snack provisions and thanks to Io and Tekla, the cats, for helping me get through the final weeks. You are all the best and I would not be here without you all.

Most of all, I am thankful to Emily. Thank you for helping me realise my potential and having my best interest at heart. Thanks for the patience, encouragement and love you have given. You are *the best one*, and you are why I am here today.

Author’s declaration

I, Edward Heywood-Everett, declare that this thesis is a presentation of original work and that I am the sole author. This work has not previously been presented for at award at this or any other university. All sources are acknowledged as references.

The material presented in Chapter 3 has been published in the following paper:

Heywood-Everett, E., Baker, D. H., & Hartley, T. (2020). Testing the precision of spatial memory representations using a change-detection task: Effects of viewpoint change. *Journal of Cognitive Psychology*, 1-15.

All authors contributed to the design of this study. E. Heywood-Everett collected and analysed the data under the supervision of T. Hartley. E. Heywood-Everett drafted the manuscript and both T. Hartley and D. Baker contributed to a number of revisions. D. Baker provided valuable advice regarding psychophysical methods and analysis.

Chapter 1. Investigating the limits of allocentric memory: A review of the literature

1.1 Introduction

How do we remember where things are? The question is more complicated than it first appears as there is more than one way to store spatial information. Over brief periods of time (seconds and minutes), we can store information about stimuli in our immediate vicinity (Wang & Brockmole, 2003). Over longer periods of time (hours and days, weeks and years), we may need to store information about places well beyond our current surroundings. The act of touch-typing, packing a bag or reaching out to turn off the alarm clock in a dark room requires a type of spatial memory but so does finding a diversion when a regular route is unavailable, recalling the context of an event in the distant past or reading a map. While the common element to all of these scenarios is a dependence on remembered spatial information, different scenarios demand different types of representation. Spatial memory is not a unitary phenomenon, it is used in different ways and in conjunction with numerous other processes.

There are multiple ways to encode space in the brain and these representations can be distinguished in terms of their brain bases and distinct contributions to behaviour (Burgess et al., 1999; Zaehle et al., 2007). One key distinction is between 'egocentric' and 'allocentric' representations. In an egocentric representation, the locations of objects or stimuli are represented with respect to their relationship to the body, or part of the body (see Stein, 1989 for review). For example, when we reach out for a coffee cup, visual information from the retinal image, combined with proprioceptive information about the relative location of our eyes, head, torso and limbs allow us to represent the relative spatial relationships between cup, eyes, head, arm and hand. This type of representation is useful in guiding action to targets in the immediate environment (for example, moving across a room, picking up an object) and as such, egocentric spatial representation is often associated with short time frames and relatively small spaces (Loomis et al., 1992; Philbeck & Loomis, 1997; Rieser et al., 1990). Over short intervals egocentric information remains valid or can be updated to take account of our movements (see 'spatial updating' literature reviewed below). For example, in everyday life, we can briefly store the location of nearby objects so that our current environment guides our actions. In the laboratory, this form of memory is measured in working memory tasks, such as Corsi Blocks (Milner, 1971), which require participants to remember sequences of spatial locations.

By contrast, an 'allocentric' or geocentric representation is abstracted from any particular position or point of view. Rather than encoding the relationship between the body and the objects and stimuli around it, it encodes the location and heading of the observer within the environment, and the locations of landmarks and objects, independent of an observer, analogous to the way a map represents spatial information (Burgess et al., 1999). As such, it has been thought that allocentric spatial representations are well suited to behaviours that unfold over longer time frames and in larger spaces. Therefore, allocentric representations are thought to be most useful when one's destination is beyond the range of perception, and in situations when a familiar environment may be perceived from a novel perspective. For example, re-orientating when navigating over long distances, finding a shortcut or new route (Hartley et al., 2003; Hölscher et al., 2011; Wolbers & Büchel, 2005; Wolbers & Hegarty, 2010).

Critically for the current project, there is robust evidence from neuroscience (reviewed below) that both ego- and allocentric spatial representations exist in the brain (Galati et al., 2000). As far as allocentric representations are concerned, the types of information encoded by spatial cells in the hippocampal formation are reasonably well understood (for example see Hartley et al., 2014 for review). However, the limits of the system, the limits of egocentric and allocentric representations separately and/or together in terms of its capacity and precision, have not been systematically investigated and psychological research is only beginning to address the possible structure and content of these representations (Aagten-Murphy & Bays, 2019; Bicanski & Burgess, 2018; Sargent et al., 2010; Segen et al., 2021).

This contrasts with egocentric spatial memory whose limits have been well investigated. In particular, visuo-spatial short-term memory (VS-STM) is a key component of working memory (Baddeley, 1992), and, like other working memory systems, is proposed to have limited capacity (Baddeley & Hitch, 1974; Cowan, 2001; Miller, 1956). This thesis will adopt some of the methods that have been applied to egocentric working memory tasks to investigate the limits of allocentric spatial memory.

In this review I will first investigate the neuroscientific evidence underpinning the allocentric representations and their contributions to spatial memory. Turning to psychological investigations of spatial memory representations in humans, I will then review the behavioural methods used to test spatial memory and highlight research that has informed us about its allocentric component. I will also focus on what this research reveals about the representation's capacity and precision. Finally, I will review how change detection has been applied to investigate the limits of (egocentric) visual working memory.

1.2 Neurobiology of Spatial Memory

The concrete basis of evidence underlying the distinction between allocentric and egocentric spatial encoding originates in insights from neuroscience. Many neurons in the mammalian brain have egocentric firing fields - that is, their firing rates are determined by the location of stimuli or responses relative to the body or parts of the body. This is true of neurons in the sensory and perceptual systems, for example retinotopic responses of cells in the primary visual cortex relate the location of external stimuli to their position in the image formed on the retina (i.e., determining an object's location by reference to the angle it makes with the eye) (Gardner et al., 2008; Guo & Li, 1997; Sereno et al., 2001). It is also true of motor neurons whose firing represents the direction of a planned action relative to the current location of the limb to be moved (Graziano, 1999). In the parietal lobe there are many other, intermediate, neurons whose firing rates reflect the location of somatosensory stimuli, or are sensitive to multimodal spatial inputs (Avillac et al., 2005). Some neurons respond to stimulus position relative to multiple egocentric reference frames (for example, responding maximally to a stimulus at a certain combination of eye- and head-centred coordinates: 'gain fields'; Andersen, 1995). These neurons are thought to be involved in sensorimotor computations underpinning the control of action (Andersen, 1995; Pouget & Sejnowski, 1997; Salinas & Abbott, 2001). In some cases egocentric sensorimotor neurons are able to sustain their activity over brief delays when stimuli are absent, suggesting a role in the maintenance of visuo-spatial short-term memory (see D'Esposito & Postle, 2015 for review). Consistent with this, fMRI studies in humans reveal that the dorsal occipital and parietal regions are active during spatial- visual working memory tasks (Chafee & Goldman-Rakic, 1998; Cohen et al., 1997; Compte et al., 2000; Pessoa et al., 2002; Qi et al., 2010).

Neurons with allocentric firing fields are found in the hippocampal formation. For example, place cells in the hippocampus fire whenever an animal is at a particular circumscribed part of its environment, irrespective of its current heading (Muller et al., 1994; O'Keefe & Dostrovsky, 1971) or the presence or absence of individual local cues (O'Keefe & Conway, 1978). Just upstream of the hippocampus itself, grid cells in the medial entorhinal cortex have similarly localized firing fields but these form a regular triangular grid that tessellates the environment (Moser et al., 2008). The firing rates of head direction (HD) cells, found in the presubiculum and elsewhere, are determined by the animal's heading (e.g., Northeast), independent of its location (Taube et al., 1990). The critical feature of this type of neural representation is that in each case, the spatial cells' firing fields are anchored with respect to the environment and are insensitive to the presence of individual stimuli or their relationship to the body.

The discovery of place cells led O'Keefe and Nadel to propose that the hippocampus acts as a "cognitive map", a latent store of allocentric spatial information that is accumulated incidentally as the animal explores an environment (O'Keefe & Nadel, 1978). Consistent with this idea, in animals' the integrity of the hippocampal formation is important for environmental spatial learning and memory (R. G. M. Morris et al., 1990). Similar cells are found in humans (Ekstrom et al., 2003; Miller et al., 2013), and fMRI evidence suggests that the hippocampal formation is important in human navigation, spatial learning and memory especially when learning is incidental, long-term and linked to stable features of the environment (Doeller et al., 2008; Rosenbaum et al., 2004).

Object vector cells (Hoydal et al., 2018) provide a mechanism through which allocentric representation can be applied to objects. Høydal and colleagues found that a large proportion of cells in the medial entorhinal cortex fire specifically when mice are at a given distance and direction from a certain object. As with other spatial cells, these firing fields are anchored with respect to the environment and in was not affected by the animals heading. This evidence shows that object spatial information, is also encoded allocentrically despite it's more transient nature, for a range of object types. This information presumably aiding tasks such as reorientation and navigation.

The distinction between ego- and allocentric spatial representations in the brain has important implications for behaviour. Broadly the roles of the different reference frames are compatible with the visual streams framework put forward by Goodale and Milner (1992). They describe two neural processing pathways that visual information follows after initial visual processing in the occipital lobe. The dorsal stream includes the parietal regions involved in egocentric processing required for the online control of action, while the ventral stream (terminating in the hippocampal formation) is concerned with encoding object identity, memory and learning (Goodale & Milner, 1992; James et al., 2002). The ventral stream is thus oriented towards lasting and static properties such as object shape and colour, and (critical for spatial memory and identifying places) stable, allocentric, configural relationships (Dijkerman et al., 1998; Gallistel, 1990; Klatzky, 1998; Neggers et al., 2006).

Within the ventral stream, spatial stimuli (for example, images of places; Epstein & Kanwisher, 1998) are preferentially processed in the more medial regions including parahippocampal gyrus which includes the parahippocampal cortex (posterior) and entorhinal cortex (anterior). Neuropsychological evidence indicates that damage to the parahippocampal cortex leads to topographical agnosia and disorientation (Aguirre & D'Esposito, 1999). Neuroimaging studies show a parahippocampal sensitivity to spatial layouts and an insensitivity to both configuration of objects and scene 'type' (i.e. environmental properties Epstein, 2008; Epstein & Kanwisher, 1998; Kravitz et al., 2011).

The hippocampus itself is implicated in a range of spatial and related tasks. Neuroimaging studies of virtual navigation show hippocampal activation is correlated with spatial task accuracy (Ekstrom et al., 2005; Hartley et al., 2003; Parslow et al., 2004). The hippocampus is also activated in situations in which participants construct spatial scenes in imagination (Hassabis & Maguire, 2007; Mullally et al., 2012). An intact hippocampus is also necessary in imagining scenes in past, future and fictitious atemporal contexts (Maguire et al., 2010) and these findings have prompted theories that the hippocampal structure is associated with 'mental time travel' (Bartsch et al., 2011), self-projection (Buckner & Carroll, 2007), relational binding (Schacter et al., 2012, p. 201) or scene construction (Hassabis et al., 2007; Hassabis & Maguire, 2009; Mullally & Maguire, 2014).

Damage to the hippocampal formation leads to profound episodic amnesia (discussed below) which can sometimes obscure specifically spatial problems. Nonetheless, patients with hippocampal pathology show deficits in route learning as well as navigation and scene learning (right hippocampal lesion - (Astur et al., 2002; Kessels et al., 2002; Spiers et al., 2001). When a patient with bilateral hippocampal pathology is asked to recall the location of group of objects in 3D space, they show considerable deficit when performing the task from a novel, rotated viewpoint demonstrating a deficit in line with a view-independent perception of space (King et al., 2002).

Overall, there is clear evidence for an allocentric encoding of space in the brain and that this representation supports parts of viewpoint-independent spatial behaviour. While an egocentric representational system could in principle work alone for some simple tasks, in more complex, large scale or long-term spatial tasks an allocentric representation is needed. In this case it seems that the two spatial representations interact to inform behaviour and action (Burgess, 2006; Burgess et al., 1998; Lenck-Santini et al., 2002; O'Keefe & Speakman, 1987; Robertson et al., 1998).

The dissociation and interaction of ego- and allocentric systems is perhaps most clearly demonstrated in cases of representational neglect in which damage has disconnected the posterior parietal areas (Bisiach & Luzzatti, 1978; also see Guariglia et al., 2013). For example, Bisiach and Luzzatti asked participants with reported hemi-spatial neglect to describe the Piazza del Duomo in Milan from either end of the square. In retrieving the memory and describing the scene, it was found that participants neglected one side of the environment relative to their position (i.e., the buildings to the participant's left were neglected regardless of direction faced). The patients' long term-spatial representation of the scene was intact however, their ability to egocentrically attend to the representation was not. This phenomenon suggests that recall is based on an intact allocentric representation which must be reconstructed in an egocentric reference frame in order for participants to be able to describe it from a particular viewpoint. In this case it is this latter capacity, to reconstruct the scene in egocentric terms at recall which seems to be impaired.

Bicanski and Burgess (2018) have put forward a detailed and implemented model of spatial memory built on the earlier Byrne, Becker and Burgess (Byrne et al., 2007) model. They describe a convincing mechanistic account of the way ego- and allocentric representations interact. In this model, spatial memory and imagery can be driven by a 'bottom-up' egocentric representation in which an initial encoding of space (as well as the locations of objects using object vector cells) is represented in an egocentric reference frame in the parietal lobes (in the model termed the 'parietal window'). It is then automatically transformed into an allocentric reference frame. This is done in the retrosplenial cortex, a medial region of parietal/posterior cingulate neocortex lying anterior and superior to visual cortex, with input from head direction cells which give information about the current heading within the environment. The resulting allocentric, viewpoint- invariant encoding is stored within the spatial cells of the medial temporal lobe described above. As well as a 'bottom-up' operation, hypothesised during perception, to account for visuo-spatial imagining, the system can also work in a 'top-down' capacity. Here the spatial context of the environment can be reformed from spatial cells and converted into an egocentric representation in the 'parietal window' which can be brought into conscious imagination.

Neurobiological evidence shows that allocentric spatial memory is not just used in esoteric tasks but that it is an important component of everyday memory for personally experienced events; episodic memory, in which the hippocampus plays a crucial role (Tulving & Markowitsch, 1998; Vargha-Khadem, 1997). Spatial context is often regarded as a foundational aspect of episodic memory (Hassabis et al., 2007; Hassabis & Maguire, 2007) and allocentric spatial memory is often characterised by its long term and stable nature (see Burgess, Becker, et al., 2002, for discussion). There are many models that tie these two forms of memory together, in some cases suggesting that episodic memory may well be organised using its spatial context (Moscovitch et al., 2005; O'Keefe & Nadel, 1978a). There is increasing evidence to suggest a shared functionality, for example, evidence that place cell activation can be triggered when participants recall contextual elements associated with particular environments (Hoscheidt, Nadel, Payne, & Ryan, 2010; Miller et al., 2013). Links between episodic memory and allocentric memory are substantial and provide motivation for a better understanding of allocentric spatial encoding. Deconstructing our complex remembered experience requires an understanding of its very basic structure, encoding and organisation.

Neuroscience has provided a concrete evidence base showing two distinct ways in which space can be represented in the brain. The dorsal stream supports an action based egocentric representation while the ventral stream and hippocampus support a view independent allocentric representation. The two systems interact to support spatial memory with egocentric sensorimotor information being converted to (and from) allocentric form for long-term memory and complex, large-scale spatial behaviours such as navigation. In the next section we will look at different behavioural tests of spatial memory and the way they draw on ego- and allocentric representation and processing.

1.3 Behavioural methods for testing spatial memory representations

As we have seen from the neuroscientific evidence, it is apparent that space can be encoded in the brain in both an egocentric, and allocentric way. However, from the literature explored above it is unclear *when* each these representations are used to inform behaviour and *how* these representations might interact when an observer makes use of remembered spatial information in different contexts.

As we will see below, there are a variety of different experimental methods which are designed to explore the separate contributions of ego-and allocentric information in behavioural tasks. By understanding the contexts in which these representations seem to be used, behavioural methods have been developed that attempt to isolate the ego-and allocentric spatial memory contributions to behaviour.

After exploring the methods used in spatial memory research, I will then return to the focus of the review and ask how the behavioural literature has contributed to our knowledge of limits in allocentric spatial memory in terms of capacity and precision.

1.3.1 Egocentric and Allocentric Contribution to Spatial Memory

At any one moment a myriad of different spatial cues, both egocentric and allocentric, are being used to inform behaviour. Active spatial updating is the automatic process that incorporates body and object movement into an egocentric representation of space allowing object location to be maintained over short time spans (Klatzky, 1998). Vestibular, proprioceptive and motor efference information and optic flow (together known as idiothetic cues) information, track the movement of the body and can be used to inform and update our current location in the form of path integration (Etienne & Jeffery, 2004). It is also argued that we are also able to locate ourselves through matching of visual snapshots from previous views (Spiers et al., 2001). This is to name but a few of the different sources of information available to us naturally at any given point. To separate and control information available in spatial memory tasks and to attempt to isolate and test discrete facets of spatial memory, careful methodological strategies must be implemented.

One such strategy uses cue-conflict and systematically invalidating different cues, allowing the contribution of individual cues to be investigated. This has been accomplished in the past with an elegant rotating table-top design which allows experimenters to invalidate visual snapshot information, self-motion and the use of distal landmarks (Burgess et al., 2004; Motes et al., 2006; Simons & Wang, 1998; Wang & Simons, 1999). In these studies, participants view an array of objects on a table-top. The participant is then blindfolded while either an item on the table is moved or the array remains the same. The participant is shown the items again and asked, (depending on the study) if an item had moved *or* which item had moved. Between presentation and testing, one of two experimental manipulations can occur depending on the trial; (i) the participant either walks to an adjacent chair or walks halfway and returns to their starting position, (ii) the table is either rotated by a corresponding angle, or left in place. This creates situations of cue conflict. For example, if the table rotated but the participant viewed the array from the original position, their proprioceptive information from walking is in conflict with their viewpoint, which has changed. Likewise, if the participant walks around the array and the table also rotates - they are viewing the array from the same perspective, despite changing position.

Wang and Simons (1998 and 1999) used this paradigm to show that in conditions where test arrays are consistent with self-motion (i.e., in situations where the table does not rotate), participants were more accurate than in conditions with viewpoint dependence (a matching viewpoint despite cue conflict). They also found that change detection performance is worse when the array is rotated compared to when the participant moves around the array. These results indicate that observer movement is contributing to updating spatial representations and suggests that an element of egocentrism is present in updating the memory of the environment.

Burgess, Spiers and Paleologue (2004) adapted this paradigm by using luminous objects in a completely dark room to exert control over environmental visual cues outside of the stimulus array. In this study, Burgess et al added an extrinsic, luminous, landmark to the environment that could also be rotated with or separately from the array making it consistent or in conflict with the other spatial information. The results replicated Wang and Simons' earlier findings despite the object array being in darkness (emphasizing the role of physical movement rather than optic flow in spatial updating) and showed increased accuracy in conditions where the array was consistent with the extrinsic landmark. This indicates that participants were using more than purely automatic spatial updating of egocentric information and suggests that they were also encoding and making use of a static (allocentric) representation of the wider environment to orient themselves and the array. This provides evidence that a viewpoint-independent representation of space may also contribute to spatial memory in short term tasks of this type.

While such results suggest that viewpoint-independent memory contributes to spatial memory in laboratory tasks, the relative contribution of both ego- and allocentric reference frames is harder to pin down and other results suggest that it can vary depending on subtle differences in task requirements. For example, Motes, Finlay and Kozevnikov (2006) used a similar rotating table-top paradigm, manipulating the degree to which the table was rotated. They found that contrary to the previous studies, there was no improvement in performance due to self-motion, and that angular distance from the initial viewpoint was the only predictor of performance, suggesting that the decline in performance was an example of an “alignment effect” indicating a preferred allocentric environmental orientation (described in detail below). In an effort to replicate the locomotion effects of the previous experiments, Motes et al. conducted a follow up in which the delay between encoding and retrieval was shortened and the number of objects in the array was reduced from 10 to as low as 4, making the task more suitable for an egocentric representation, however, no movement updating effects were found.

The authors suggest that the way in which a participant has encoded the environment can change the balance of egocentric/allocentric information in spatial memory tasks. Tasks that require participants to learn locations from within the array (Farrell & Robertson, 1998; Wang & Simons, 1999) place more emphasis on egocentric representation useful for action, whereas for tasks that entail learning locations from outside the array (Holmes et al., 2018; Motes et al., 2006; Mou et al., 2006a) an allocentric representation is preferred (although Burgess et al., 2006 does not show this).

Whether an array is encoded as a whole or in piece-meal also seems to influence the contribution of allocentric reference frames. In Holmes et al. (2018), movement (active and passive) around an array increased participants’ ability to accurately locate array items in an open table-top environment, whereas only self-motion, walking around the array, improved performance when an environment was experienced in piecemeal. If idiothetic cues, as well as the perspective from which we encode the array matter for the ego/allocentric representational balance, this raises methodological concerns regarding the use of screen-based virtual-environments and the effects of immersion in spatial memory research.

A recent study has shown that the nature of the environment itself might influence the weighting of representational contribution to remembering. In Aagten-Murphy and Bays (2019) participants were required to remember object locations in a 2D array and replace a cued object when a landmark was present or not. The results showed an increased accuracy when a landmark was present. This further increased the closer the landmark was to the cued object. The authors interpreted this effect to be evidence of an allocentric contribution even though the task could have been completed purely with egocentric cues such as view matching. This suggests that both egocentric and allocentric sources of information are available in tasks but that representational sources may be weighted based on their usefulness.

Taken together, these studies show that both egocentric and allocentric reference frames do contribute to spatial memory but that the individual contributions of ego- and allocentric representations can vary based on task requirements. In the next section I will look how the ways in which the use of allocentric and egocentric reference frames can be detected behaviourally in spatial memory.

1.3.2 Detecting Reference Frames in Memory

1.3.2.1 Retrieval Orientation

One characteristic of the use of an allocentric representation is ''alignment effects'' a term used to describe the effect that retrieval orientation has on performance. In a study by Sholl, (1987) participants in a university building were pointed in the direction of familiar locations on campus and to more geographically distant locations. It was found that although orientation on recall had no effect on accuracy, participants performed significantly faster, locating landmarks in front of them when their notion of a wider spatial context was aligned with their current heading.

Alignment effects have been shown to suggest that the preferred orientation is often the same as the orientation of the environment at first encounter (Diwadkar & McNamara, 1997; Kelly & McNamara, 2010; Shelton & McNamara, 2001). This finding has led to some confusion in the literature, being used as evidence to support a solely view dependent (single representation) model of spatial memory in which we are relying on egocentric viewpoint memories of space in order to orientate ourselves (Wang & Spelke, 2002).

However, when there is an intrinsic axis to the array (for example, all array items are in a line) the alignment effects are no longer seen for the initial viewpoint and instead participants respond fastest from an optimum position relative to the array (Mou & McNamara, 2002). Alignment effects have also been demonstrated to coincide with intrinsic environmental axes (Adamou et al., 2014; McNamara et al., 2003). These novel-environment alignment effects have also been shown with large scale environments, virtual environments and cartographic maps (Evans & Pezdek, 1980; Frankenstein et al., 2012; Richardson et al., 1999; Roskos-Ewoldsen et al., 1998; Werner & Schmidt, 1999). Together these results suggest that rather than reconciling with a prior egocentric viewpoint, alignment effects are a sign of an intrinsic directionality to our allocentric representations. This relationship may be defined by factors such as how the environment was originally viewed, its intrinsic axis or even the environments extrinsic axis.

The response latency shown in these studies might reveal something about the interaction between egocentric and allocentric representations. As described in the Bicanski and Burgess model, the process of reorientation using a stored allocentric representation of space requires conversion from initial egocentric encoding. This could either be accomplished by mentally rotating ("mental rotation") a representation of the environment so that it aligns with the viewer or mentally travelling ("perspective taking") the path between known to current viewpoints (Meilinger, Berthoz, & Wiener, 2011). Either way, transforming between a held allocentric representation and an egocentric viewpoint is an effortful process. There is a cognitive cost in time or accuracy of mental rotation or mental perspective-taking and storing a default orientation might make this process more efficient. As environments become more familiar these alignment effects diminish (Easton & Sholl, 1995; Meilinger et al., 2007).

1.3.2.2 Directional judgement

The literature on directional judgment shows another example of how task differences change the balance of ego-and allocentric contributions to memory. It also provides a useful technique in how to develop tasks which rely more on the use of viewpoint-independent memory.

Tests of *absolute* directional judgment generally involve participants learning the locations of objects within a scene or array before being disoriented and blinded from the stimulus. The participant's accuracy and reaction time in estimating object directions can be used as measures of spatial memory accuracy. Wang and Spelke (2000) showed participants an array of objects (a room full of items with the participant in the centre) before asking them individual object locations in one of three conditions; (i) with the objects removed, (ii) while blindfolded and (iii) after being disoriented (spun on a chair) and while blindfolded. The results showed that disorientation produced increased variation in directional judgement performance. The authors argued that this was evidence for an egocentric model of spatial memory and a lack of an allocentric representation. They reasoned that if a coherent allocentric representation were available, disoriented participants would use this and their directional judgements would show a systematic error towards a constant rotation. The fact that disorientation produced increased variation in judgement was taken as evidence for an egocentric representation coupled with spatial updating; under this model each object's stored location could be disrupted independently.

However, by providing participants with a task that without the blindfold would usually be completed egocentrically (i.e. indicating direction from position and point of view), it has been argued that participants may still try to use egocentric information despite it being quite inaccurate (see Burgess, 2006 and Ekstrom et al., 2014 for discussion). To address this, Waller and Hodgeson (2006) replicated these findings, adding another condition in which a judgement of *relative* direction (JRD) was assessed. Here, after blindfolding and disorientation, the participants were also asked to point to the location of an object from an imagined viewpoint (i.e., imagine standing at the chair, facing the table, and pointing to the TV). It was found that, in the JRD task, findings were reversed, and the disorientated group outperformed participants who had been blindfolded but not disoriented. The authors argue that this is evidence of a *representational switch*: in situations where allocentric representations are available, the conflict produced by disorientation necessitates a move from a precise action based egocentric representation to a comparatively coarse allocentric one. This also suggests that in tasks that do not include a novel viewpoint or require a reorientation process, a more easily available egocentric representation might be preferable.

The inclusion of the JRD in the Waller and Hodgeson study indicates again that the contribution of ego-and allocentric information to spatial memory is variable and task dependent. For many tasks spatially updated egocentric representations may be preferred as faster, more precise, and less effortful to create and maintain. However, tasks such as the JRD, which require participants to adopt a new imagined perspective, make available the use of allocentric information for use in the task.

1.3.2.3 Error Characteristics

In Wang and Spelke's absolute directional judgement task, increased error variation in response was used as evidence for an egocentric model of spatial memory. While further evidence reveals that this may be task dependent (Waller & Hodgson, 2006), the method of characterising error as way to identify global or piecemeal updating of spatial elements is a clever way to detect ego-and allocentric representational use in spatial memory.

In theory, spatial memory reliant on an egocentric representation would allow for random error-localised to individual array elements, while a globally updated allocentric representation would show global and systematic error. With this in mind, Sargent, Dopkins, Philbeck and Modarrez (2008) measured directional pointing error following repeated and cumulative disorientation. In their study, participants learned the location of four objects before undergoing whole body rotations. After each rotation participants were asked to point to the objects. They found that for some trials error was seen to be similar across objects, suggesting global updating, and in some cases random and isolated to individual objects suggesting a piecemeal updating. This implies that for different stimuli, in the same experiment, egocentric and allocentric representations were contributing different amounts. In a second experiment (Sargent et al., 2010) spatial layout was recorded and varied. They found that for spatial arrays that were characterised by physical closeness of objects or by regular configuration (such as in recognisable geometric shapes), elements were encoded globally whereas more random configurations were encoded by piecemeal. This shows that a kind of spatial chunking is taking place for object configurations that show intrinsic geometry (see also McNamara et al., 1989) and that these chunks are spatially updated as one. The authors suggest that, in agreement with the Wang and Spelke (2000) model, the objects around the room were being updated egocentrically but that more stable environmental features (such as the room corners, which did not show increased variation in response error) were updated allocentrically. However, as the authors point out - the notion of object must be relaxed to also include groups of objects. The question of what constitutes a unit of spatial representation and memory is central to the question of spatial memory capacity and one that we will return to later.

There are a number of different ways to detect and identify egocentric and allocentric contributions to spatial memory for example using alignment effects, allocentric-oriented tasks such as JRD, and by using error to show how elements in the environment were updated. These insights can be used to characterise viewpoint-independent tasks. In the next section I will look at studies that aim to isolate allocentric representations explicitly and how we can judge their success.

1.3.3 Isolating Allocentric Contributions to Memory

Studies detecting egocentric and allocentric representations in spatial memory have provided insights into how tasks can be constructed to weight the use of particular reference frames. Using a virtual ‘town square’ environment (King et al., 2002) tested participants’ ability to recognize the locations of objects presented sequentially above randomly selected markers during an encoding phase. At retrieval the participant viewed the environment either from the same viewpoint or (having been relocated to disrupt any spatial updating of an egocentric representation) from a new rotated viewpoint. A patient, Jon, with bilateral damage to the hippocampus, showed a significant deficit in performance for the rotated view condition while performing in line with a healthy control group when tested from the same viewpoint.

The fact that a hippocampal patient was specifically impaired on only the version of the task in which the scene was viewed from a novel viewpoint suggests that an allocentric representation of space was being used by the healthy controls but was not accessible for Jon. This is an important result as it seems in line with neuroscientific evidence and confirms that we have an accurate picture of what role allocentric spatial memory takes and the kinds of tasks that utilise it.

The four mountains task (4MT, Hartley et al., 2007) adopts a similar 'novel view' approach to look at the content of allocentric memory in more naturalistic scenes, taking several steps to encourage allocentric and discourage egocentric strategies. Rather than a sequential, object-based memory task as seen in King et al. (2002), it uses a forced choice recognition task. The participant is shown an initial stimuli- a carefully constructed range of four mountains of varying shape and size textured with noise to create naturalistic variation. The participant is then asked to choose the same spatial environment from a selection of similar images. An unpredictable viewpoint shift is incorporated between encoding and presentation meaning that the target item is seen from a novel view that could not be imagined beforehand. Non-spatial, contextual properties of the environments (such as lighting, lighting direction, cloud cover, vegetation) are changed between encoding and testing so that participants cannot rely on local visual cues. The incorrect 'foil' images are constructed with similar elements to the target landscape in different spatial configurations. The combined effect is that memory for individual features, visual cues or local configurations is not reliably sufficient for correct recognition.

In the initial study, the results showed a significant task deficit for the group of hippocampal patients compared to healthy controls. There was also no deficit in the hippocampal group's performance for the non-spatial equivalent of the task (matching non-spatial characteristics of the landscapes). Using the 4MT as a measure of allocentric spatial ability, subsequent studies have investigated other forms of hippocampal pathology including Alzheimer's Disease (Bird et al., 2010; Pengas et al., 2010) and Mild Cognitive Impairment (Moodley et al., 2015). Perhaps most notably in the current context, variation in the volume of the hippocampus is associated with performance in healthy young adults (Hartley & Harlow, 2012) suggesting that the task is sensitive to physiological constraints that may lead to detectable individual differences in the limits of allocentric spatial memory. This theory is supported by similar 'brief-delay, viewpoint-rotation' tasks carried out with healthy participants in an fMRI scanner and show increased activation in the hippocampal formation associated with successful performance in the rotated view spatial memory task (Schmidt et al., 2007).

Generally, the research showing performance deficit in patients with hippocampal damage or atrophy indicates that carefully designed laboratory tasks that resist egocentric or visual strategies may be able to isolate hippocampal-dependent allocentric representations which can support spatial memory even over brief intervals. However, at the start of this thesis, such tasks had not yet been adapted to investigate the limits of allocentric spatial memory in any detail. In the following section, I turn to these limits, and what existing research tells us about them.

1.3.4 Limits of Allocentric Performance

Despite using a participant's spatial precision as a measure of accuracy in many tasks, limits of allocentric spatial memory have not been systematically explored in previous studies.

In terms of precision, allocentric spatial representations are generally characterised by their ''coarse'' nature compared to a detailed, egocentric representation (Waller & Hodgson, 2006). In terms of capacity, allocentric spatial memory is associated with the hippocampus and with episodic memory. When long-term memory for familiar/unfamiliar scenes is tested, participants are able to distinguish familiar and unfamiliar images from extremely large collections of images at a high level of performance, suggesting an extremely large capacity (Brady et al., 2008; Standing, 1973; Voss, 2009), However, tasks such as the 4MT also suggest a role of allocentric representations over a very short time frame in which, even healthy participants do not perform flawlessly.

In the spatial updating table-top studies, accuracy was measured in terms of 'proportion of changes correctly identified'. The Burgess et al (2004) study had participants remember 5 objects, one of which moved between pre-allocated point from 22 to 73cm apart. In the most egocentric combination of factors (no movement of participant, no rotation of table or distal cue) participants were able to correctly identify which object had moved an extremely high proportion of the time (around 93%). Whereas in the condition where participants approached the array from a novel direction, whilst the array did not rotate (the most ''allocentric'' version of the task without cue conflict) participants performance suffered (correct 82% percent of the time). In this case the viewpoint-independent condition was only slightly less effective than the viewpoint-dependent condition. This could be due to their range being slightly too low and performance hitting a ceiling. In comparison, Motes et al's (2006), first experiment tested participants with arrays of 11 objects which moved a similar distance, and participants were correct in their same/different judgement around 80% of the time consistently regardless of view rotation. In the second experiment, only 5 objects were in the array, here proportion correct was actually much lower, peaking at around 70% for 0 degrees rotation and around 40% correct (lower than chance) for any view rotation over 72 degrees. This stark difference in performance is very surprising and performance seems to hit a floor based on a certain degree of viewpoint rotation. This, combined with the results from above provide rationale for the development of a task with a dynamic difficulty range that varies viewpoint rotation.

We can see from the rotating-table studies that viewpoint-dependent memories seem to be more precise and have more capacity. It is interesting that between the Burgess et al study which used 5 object arrays and the Motes et al study which initially used an 11-object array, there was a similar, very high level of performance (between 80 and 100% correct). This suggests that between the experiments there was no informational capacity limit reached and participants were able to accurately recognise scenes with up to 11 objects. While participants seem to be able to recognise scenes with large set sizes with little difficulty, other literature suggests that the precision at which a participant can *reconstruct* a changed array is strongly affected by set size (Wang et al 2006 showing effects of set size for 1-3 items). These contrasting results as well as the significant drop in performance between the first and second Motes et al study is suggestive of the highly task dependent nature of the results and that it would be unwise to compare results between experiments despite the use of a very similar method and stimulus. This is a subject that will be explored in depth later in the thesis.

Wang et al (2006) specifically tested the effect of set size on performance in a virtual locating task. In this study participants learned the locations of objects around them in a virtual room. The objects were removed, and the participants were required to re-place them in their original locations. In one condition the participants were also required to move to a new location between encoding and retrieval. The study found that an increase in set size from 1 to 3 objects both increased response latency and response error. In the ‘same view’ condition the error increased by approximately 5 cm however, in the condition where participants moved from their original position, the error increased by approximately 15cm. The authors interpreted this result as evidence of piece-meal updating system that was negatively affected by the number of objects in the array. However, unlike the examples of explicit allocentric testing, self-motion cues were still present in this study and so it may be that egocentric cues were present and preferably weighted. The results could also be interpreted as evidence that the process of conversion or translation between egocentric and allocentric is made more difficult with more complex environments. Despite these caveats, the results do show an effect of set size on spatial precision. This is an important result and could be indicating a capacity/precision trade-off similar to that proposed in visual working memory research (discussed below).

The study by King et al required participants to identify the locations of between 2 and 7 objects in the view-change (allocentric) condition. In the control group, the number of objects in the array did have some effect on performance, which fluctuated between 75% and 60% correct with a slight negative trend. Interestingly, during the ‘view consistent’, egocentric task the control group showed very little decline in performance at all up to and including the maximum 15 object array. However, as the task involved participants selecting possible location markers rather than freely estimating location, it is possible that the measure of accuracy was not sensitive enough to pick up on a small decline in performance.

The 4MT task employs a view rotation task using an image of a large-scale environment suitable for allocentric representation. In the topographical memory condition, controls were able to correctly identify a remembered environment approximately 80% of the time. Some of the 'foil' images showed mountains in the same position as the target but differently shaped. For one, this tells us that even with only four items in the array, the task is sometimes too difficult for healthy controls. Secondly, participants would require a highly precise representation to be able to distinguish between a target and foil which has the same spatial layout but where two only two items are switched.

Information on the relative precision and capacity of allocentric spatial memory shows a mixed picture. Overall, it seems that participants find tasks that require an allocentric component harder, and their spatial precision suffers. Set size in an array seems to have a particular effect for allocentric rather than egocentric performance however, there are conflicting results. Further research on the quantifying capacity and precision of spatial memory is necessary and that is what this thesis intends to do.

1.3.5 Summary of Behavioural Methods

In sum, behavioural research into spatial memory has contributed a great deal in terms of characterising the use of allocentric reference frames in spatial memory and in constructing tasks that are reliant on view-independent information. However relatively little is known about the detailed limits of allocentric spatial memory, its capacity and precision. In the next section I will review the contribution of the change detection paradigm in assessing the limits of visuo-spatial working memory.

1.4 Visuo-spatial Working Memory

One way in which the structure and content of cognitive representations can be explored is by looking at a system's limits. The psychophysical methodology of change detection has been used in the field of visual working memory, using a participant's ability to discriminate change as a gauge of how much information is stored and how precisely it is represented representation (for a review see Brady, Konkle, & Alvarez, 2011). When a participant cannot notice a change, the memory representation either does not contain the element that changes or the representation lacks the definition to detect said change. It is also important to note that the vast majority of the following studies involve 2D displays which can be represented and processed without recourse to allocentric memory. However, as we shall see, the methods that have been developed to probe 2D visuospatial working memory may also be applicable to questions about the limits of ego- and allocentric memory in 3D spatial tasks.

When varying the number of items within an array as well as the complexity of the items to be encoded, certain patterns in performance emerge, suggesting informational constraints on the visual working memory system. This evidence is contentious, with some interpretations proposing a slot based system in which 3-4 items can be maintained within the visual representation before performance declines (Cowan, 2001; Luck & Vogel, 1997; Zhang & Luck, 2008), others suggest a more fluid, informationally constrained model in which the resolution at which an item is remembered is proportional to the 'cost' of its representation (Alvarez & Cavanagh, 2004; Ma & Wilken, 2004). The change detection methods allow for the investigation of the capacity and precision of the representations by quantifying the smallest change the system is able to detect.

Capacity testing follows a tradition in working memory research (Baddeley, 2003 for review). Its success and suitability to this application is perhaps due to the conceptualisation of working memory as a unitary system with well-defined limits (for review see Logie, 2011). Within the working memory model separate cognitive systems manage different information modalities. The visuospatial sketchpad and phonological loop manage short term 'RAM-like' memory of visual/spatial information and sound information respectively. This can be further dissociated into more fundamental aspects of information such as speech, music, sound, vision and space (Baddeley & Hitch, 1974; Schulze & Koelsch, 2012 and for visual/spatial dissociation see for example Tresch et al., 1993, for evidence of dissociable spatial and visual memory disruption and neuropsychological studies of Williams syndrome patients see Vicari et al., 2003)**.**

1.4.1 Capacity and precision in VWM

With an overall aim in this thesis to study spatial memory representations through their limits of capacity and precision – it is worth reviewing the gains that visual working memory research has made by looking at the system demands and organisational structure of the storage system. The capacity limit of a system can be defined as a maximum limit of information stored within a representation. The precision or fidelity can be defined as the detail present in a representation or its resolution. These features are often investigated using a 'two alternative forced choice' (2AFC) approach in which participants indicate whether or not they notice a change from one stimulus to the next. By changing the difficulty of the task, performance can be modelled using a psychometric function- similarly to tests of psychophysics in, for example, vision or hearing. Modelling data in this way is an attempt to quantify a standard measure of experience. This is done by finding the maximum task difficulty where participants reliably perform better than chance (i.e., 75% correct responses for a two alternative choice task – this concept is explored in greater detail in Chapter 2). Using these techniques in VWM, research has found a trade-off relationship between memory capacity and the fidelity of memory representation.

Luck & Vogel, (1997) conducted a change detection task in which participants completed a 2AFC task with a stimulus of coloured squares. The ‘change’ condition was a change in either colour or location of an array item. The results showed that increasing the array size to over 3-4 items incurred a steady incremental decline in performance that increased with the number of array items. This finding led to a limited, slot-based theory of VWM. This theory describes discrete slots within VWM that each house items that can be remembered at a high level of fidelity. Beyond this limit, the quality of information retained steeply declines due to a processing ‘bottleneck’ somewhere in the system. Using ROC curves to predict performance in this task rather than linear or psychometric functions, Rouder et al., (2008) found ROC to be a good model of performance. This model describes near perfect performance within the slot limit and a steep fall-off when this limit is exceeded. This would imply an ‘all or nothing’ situation well described by the slot-based model. These results have been replicated with a variety of stimuli (Donkin et al., 2013, 2014).

Subsequent experiments (Alvarez & Cavanagh, 2004; Awh et al., 2007; Ma & Wilken, 2004) include fidelity of representation, as well as capacity, as a measure. This is achieved by controlling the amount of change in the 2AFC task. For example, noticing the difference between red and blue would require a lower fidelity representation than a change from red to pink. Similarly, a change in array item location of 10 cm requires a lower fidelity representation than a change of 0.5 cm. In their initial experiment, Alvarez and Cavanagh demonstrated an interaction between capacity and fidelity by increasing the complexity of the array items. By using 3D style cubes and varying either rotation of the cubes (high fidelity change) or by switching the cubes for another item (low fidelity change), they found that change detection ability decreased with more complex objects. This led to an informational limit theory in which rather than discrete item slots, an informational resource is distributed fluidly between items. A more demanding representation would use more of this resource and therefore would be limited in terms of the number of items represented. By analogy, one could think of a camera memory card which could store a certain number of high quality photographs or a larger number of photographs at a lower quality. Awh et al., (2007) found similar results, testing participants with small and large changes to small and large arrays, they found that participants were able to detect large changes (from a cube to Chinese 'Hanzi' character) at 4-items, but were only able to detect a small change (cube rotation) with a one item array. This model of memory limitation has no bottleneck limit, instead it is proposed that error and misremembering, is due to internal noise which increases with set size (Ma & Wilken, 2004).

The debate as to whether this limitation to memory is best characterized in unitary slots or as a fluid resource is very much ongoing (Luck & Vogel, 2013). Donkin et al., (2016), in light of both convincing and contrasting information on both sides proposes a third option, describing a system in which a resource-based model might behave as though it were based upon slots. They argue that, whether or not performance looks 'slot-like' or 'resource-like' depends on the task itself. When a resource demand is predictable, for example when the number of items in an array remains constant throughout the experiment, the participant is better able to control resource allocation. When the task demands are unpredictable, the participant opts for a strategy in which a small number of items are encoded with a high degree of precision, producing data that looks as though it is limited by a slot-based system.

Other alternatives to the 'slot-based' or 'resource based' debate have been suggested. Schurgin et al., (2018) propose an alternative account of the evidence, that the reason behind confusion and contrasting evidence relates to how we apply physical differences in the physical world to 'psychological space'. The change detection framework, and the previously mentioned studies, presuppose that the difference between two stimuli is comparable to the difference between those stimuli in perception. By instead taking into account a nonlinear psychological space and accounting for the psychological distance between two objects, the two seemingly separate elements of capacity and precision instead can be thought of as a single dimension of distance in psychological space. For example, the difference in hue between two pairs of colours (a light green and dark green, a light red and dark red) might be equal in terms of hue contrast or difference in wavelength, however, the distance between their psychological representations might not be equivalent. Shurgin et al first tested a stimulus pair's distance in psychological space using a triad task (Maloney & Yang, 2003) and used this information to inform performance modelling (i.e. psychological distance as well as physical difference informed change detection trial difficulty). By taking into account distance in psychological space, task difficulty and performance (both in terms of number of objects in the array and the precision of the representation) could be explained with a single factor of distance in psychological space.

Finally, in one recent study, similar methods to those described above have been applied to spatial representations. In Aagten-Murphy and Bays (2019, mentioned above) it was found that in a 2D spatial array, the presence of stable, unmoving landmarks increased the participants' accuracy when asked to replace cued objects. Any cues drawn from the spatial relationship between the landmark and target were interpreted as allocentric. They found that the closer the landmark was to the cued object, the larger the effect the landmark had. This advantage was interpreted as additional allocentric information being used to inform the task. It was also found that when a landmark was present but not used (i.e., when the distance between landmark and cued object was so great that no advantage was seen), performance was the same as when no landmark was present at all. The authors took this to mean that the presence of a landmark did not incur a resource cost on the items stored egocentrically, suggesting separate resource pools for egocentric and allocentric information. These results are very compelling and can form the basis of research investigating these claims within a more ecologically valid and spatially complex environment. As discussed previously, when egocentric cues are available, in some cases they are preferentially used. In this experiment, with a 2D stimulus and no cue conflict, the task could be completed with purely egocentric strategies. Interestingly, this study also presumes that when no landmarks are present, allocentric information is not being generated (or if generated is not consuming any information resource). However, it is possible that regardless of whether stable landmarks are available, allocentric representations are being generated automatically. It is interesting and highly relevant to consider whether the resources linked to memory capacity are consumed when the information is available but not utilised.

1.5 Conclusion

Evidence from neuroscience provides a concrete basis for two different ways of coding space in the brain; an egocentric representation that encodes locations relative to the viewpoint and position of the body and an allocentric representation that encodes locations relative to other parts of the environment. These can be dissociated by their discrete neural areas and dissociable neuropsychological case studies. By understanding the function of a number of spatial cells and the functional connectedness of relevant areas we can understand computationally how the spatial memory system might work overall.

Techniques have been developed to help characterise behaviour associated with egocentric and allocentric representational use respectively. Studies have focussed on representational flexibility, situational demands and which discrete spatial cues contribute to our knowledge of the environment. Despite often using error and spatial accuracy as a comparative dependant variable in many of these studies, the size and detail of our cognitive maps is largely unclear.

The VWM literature has used signal detection methodology to test the limits of egocentric visual representations. What has emerged are specific predictions on informational limits of visual working memory and what systemic properties exist to produce such limits.

The application of such methods on allocentric spatial representations would provide a nuanced approach to investigating our spatial codes and perhaps help us better understand the quality and structure of these representations.

Chapter 2. Methods

The aim of this thesis is to better understand the limits of spatial memory by quantifying its precision and capacity, using these measures to investigate the part played by allocentric and egocentric representations in determining these constraints. The methods used combine techniques from two distinct fields; spatial memory testing and psychophysics. From spatial memory research we can use viewpoint changes to selectively disrupt the spatial cues available to observers to investigate precision of viewpoint-dependent and viewpoint-independent memory separately. From the field of psychophysics, we can use the concept of perceptual thresholds, adaptive staircases and change-detection techniques to find the limits of spatial memory precision and capacity.

Regarding working definitions, precision refers to the degree to which the location of objects are exactly (or approximately) represented in spatial memory. Except where indicated no more technical interpretation is implied. We can similarly investigate the number of items that can be accurately stored in spatial memory, this we refer to as capacity, and again no more technical interpretation is implied except where indicated.

As discussed in Chapter 1, the capacity and precision of a memory system may interact, with the nature of such interactions being of theoretical interest. In the thesis I aim to quantify and describe capacity and precision under varying conditions including under viewpoint manipulations that may influence the involvement of allo- and egocentric representations.

In this chapter, methods used throughout this thesis will be described in detail. This includes a description of the general theory behind the view rotation paradigm, psychophysical thresholds, and their combination. Additionally, methods used to dynamically create spatial stimuli will be described as well as details on the process of estimating psychophysical thresholds from datasets. Pilot data collected during the development of these methods will also be provided to evidence and contextualise the process of development.

2.1 Isolating allocentric spatial memory

As established in Chapter 1, tasks involving a fixed viewpoint or 2D spatial display are often preferentially solved through a combination of egocentric representation and spatial updating whereas tasks employing an abrupt and discontinuous change in viewpoint prevent the automatic visual and vestibular, egocentric updating that would normally occur and so require or emphasise allocentric representations (Ekstrom et al., 2014 for review).

Because we cannot directly infer the involvement of ego- and allocentric representations from behaviour alone, and because it is possible that behavioural responses combine information from ego- and allocentric representations, in this thesis I will use the term "viewpoint-dependent" to refer to behaviour and performance measures that linked to a particular viewpoint and "viewpoint-independent" to refer to behaviour and performance measures in tasks where the viewpoint changes abruptly and unpredictably between presentation and testing. The latter type of behaviour is expected to rely more heavily on allocentric representation since egocentric representations cannot be reliably updated to take account of such viewpoint changes, whereas allocentric representations are by definition anchored to the environment and independent of the observer's viewpoint.

In the methods developed in this chapter we use a novel-viewpoint, change detection task. By instantaneously moving the participant around an array between presentation and testing, we provide a novel viewpoint of a familiar environment in which participants must use viewpoint-independent information or processes to inform their memory. By comparing performance on a viewpoint shifted task with performance on a static viewpoint task, varying the angle of rotation away from initial viewpoint (i.e., how similar/dissimilar the novel view is) we can manipulate the level of viewpoint independence in the task and potentially isolate or bias egocentric and allocentric contributions to spatial memory.

2.2 Psychophysics and memory limits

As for a dependent variable, our aim is to gauge the limits of capacity and precision in spatial memory representations. This brings us to psychophysics, a discipline dedicated to quantifying experience.

2.2.1 A brief history of the psychometric function

The origins of psychophysics can be linked with the emergence of early experimental psychology and empirical philosophy. It aims to establish a measurable relationship between conscious experience and the physical world, mind and body, a ‘psychophysics’. In order to quantify the relationship between the physical world and sensory experience, a quanta or unit is needed. While the physical world is full of appropriate unitary measurements such as weight, length, brightness, loudness etc, experience of these stimuli are subjective and different for each person. The difficulty is to measure and compare the subjective experience of a stimulus, and psychophysics aims to do this by describing subjective experience intensity as a function of the physical stimulus intensity.

The concept of sensory thresholds, introduced by Fechner (1860), is based on determining the magnitude of a stimulus dimension necessary for perception; the 'absolute threshold' being the minimum stimulus intensity required for a stimulus to be perceived at all. Similarly, a perceived stimulus must undergo a certain physical magnitude of change for a corresponding qualitative change to be noticed (for example a light getting brighter, a sound changing pitch or an object moving). This magnitude of change is the 'difference threshold' and represents the amount of change required to produce a 'JND' or just noticeable difference. These threshold measurements create the basis of a quantifiable and comparable metric which allows for sensory experience to be represented as a function of the physical stimulus responsible for its generation. As the experience can be expressed in quantifiable and physical terms, threshold levels can be tested empirically and experience can be compared between people and situations, something that many dualist philosophers had, at the time, thought impossible.

Psychophysics, and the related signal detection theory, have shown to provide very versatile and robust methods that continue to be used in modern psychology and neuroscience. They have been successfully applied to a wide range of phenomena, from classical applications such as vision and hearing sensitivity to more nuanced applications such as memory, language recognition and even rock climbing (Bosch & Sebastián-Gallés, 1997; Gallotta et al., 2015; Luck & Vogel, 2013).

2.2.2 Understanding precision in terms of the psychometric function

Finding a participant's threshold first requires data showing the participant's performance in trials across a range of stimulus intensities. The test range should include the participant's threshold/JND and so should include both trials in which the stimulus change is too difficult for the participant to detect and also trials in which the stimulus change is very easy for the participant to detect.

Figure represents the idealized performance of a participant in a change detection task. The objective is for the participant to identify which of two objects in an array have moved. The y-axis shows performance; the proportion of trials in which the participant was able to detect the stimulus change. The x-axis shows stimulus change intensity across a test range, (in decibels) or to put it simply, how much the target object moves. We can see that before '0' on the x-axis participants are scoring 50% each time - they are guessing. After point '0' the participants are performing at 100%. Their absolute threshold in the task is '0dB', before this point the participant was unable to detect a difference in the stimuli, afterwards the participant was always able to differentiate the stimuli.

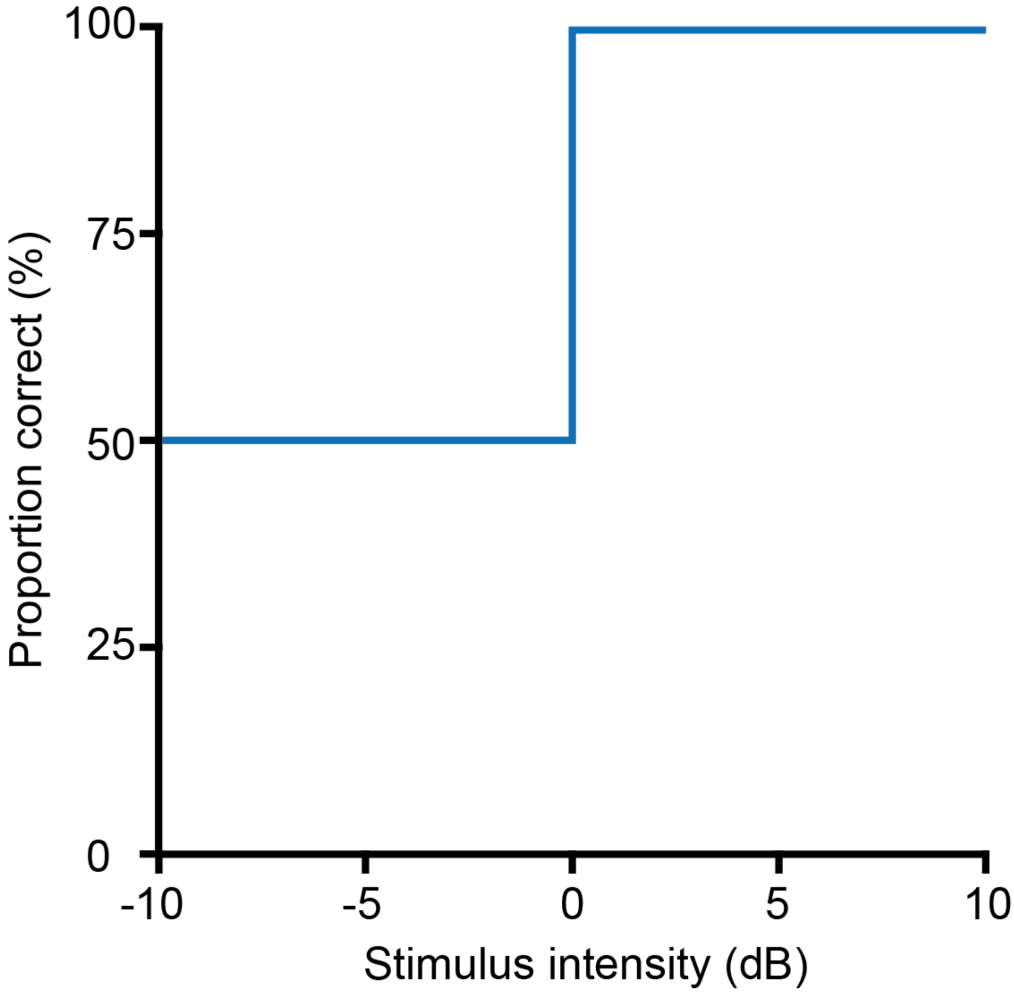

Figure 2.1 A noiseless psychophysical curve. This shows a participants ability to detect a change in a stimulus. The x axis shows change in intensity in dB and the y axis shows proportion correct. When the stimulus change is below 0dB, the participant is not able to detect any change and they are guessing– scoring a 50% proportion correct. Above 0dB the change is always big enough for the participant to detect. The threshold level of performance is at 0dB.

However, tasks are rarely so binary and judgements can be uncertain due to a wide variety of effects including internal and external sources of noise and interference (apparatus, imperfections, inattention, neural/biological fluctuations etc - Swets, 1961). This noise can be represented by a normal distribution centred around the absolute threshold. The nearer the trial intensity is to the threshold, the more influential the noise is in determining a decision. As such, performance is generally better described by a sigmoidal curve, as shown in the example below (Figure 2.2) rather than the step function above. The sigmoid or psychometric curve is described by two parameters:

the threshold, which describes where the graph lies on the x-axis and the slope, which describes its steepness.

The threshold parameter indicates the stimulus intensity when the participant's performance is at 75% correct and represents a measure of representational precision. The slope parameter is governed by the noise-gaussian standard deviation, changes the steepness of the sigmoid and so indicates the amount of noise and uncertainty in the task. This can vary dramatically depending on individual differences and stimulus type.

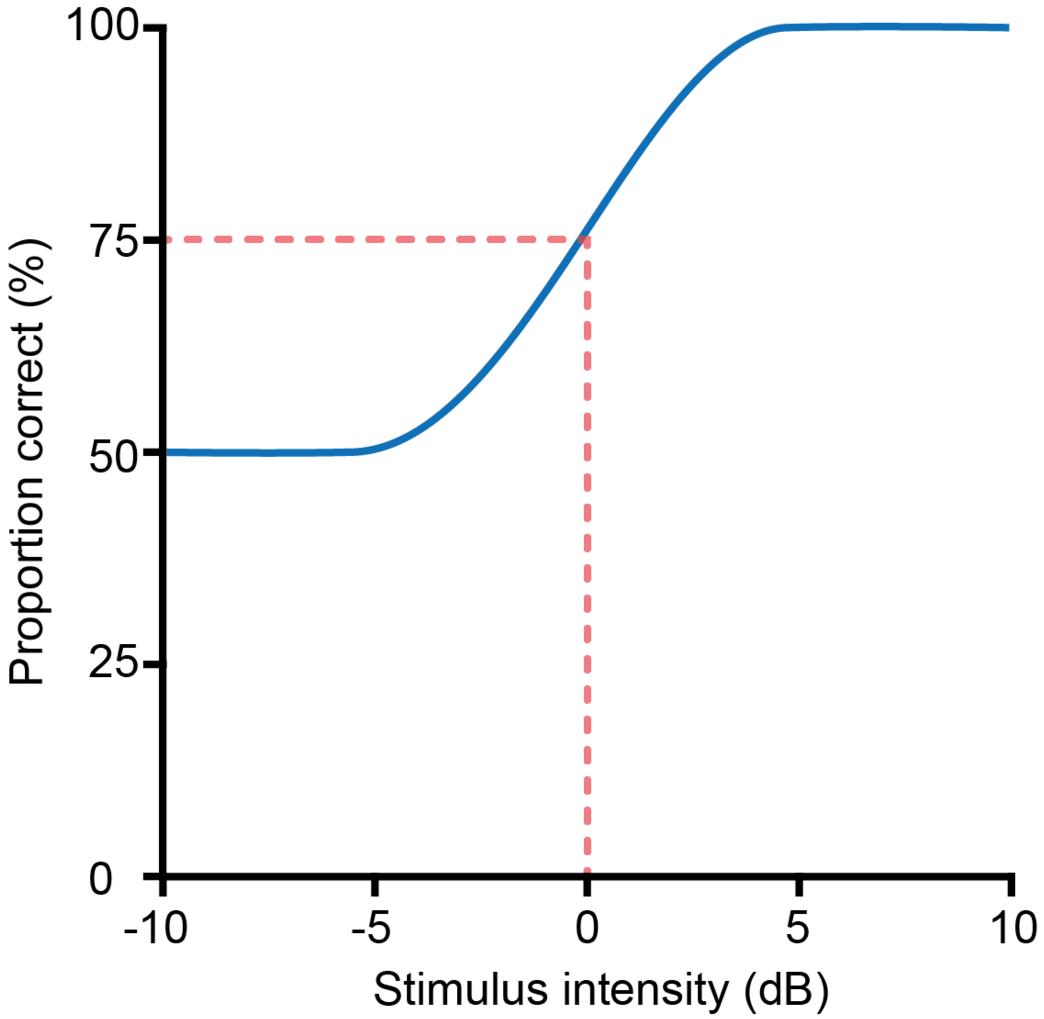

Figure 2.2 A psychometric curve the example gaussian error distribution around the threshold. The blue sigmoid is a summation of the gaussian error distribution and the sawtooth wave in figure 2.1. The red dotted line denotes the change detection threshold corresponding to the stimulus intensity where participants are correct on 75% of trials.

The psychometric function is the sigmoid which is fit to the distribution of participant proportion correct scores across a test range (Figure 2.4). The 'JND', or absolute threshold is the magnitude at which a change is reliably detected. As we are using two alternative answers, the lowest performance we would expect to see is 50% (chance or guess rate), the highest is 100% and so the threshold would be between these, at 75% proportion correct (denoted by the dotted line in Figure 2.2). Overall, the psychometric function describes 'proportion correct' as a function of our independent variable, stimulus intensity allowing us to precisely track and compare ability levels across tasks.

2.2.3 Applying psychophysics to spatial change detection

In the current thesis we will be applying methods used in psychophysics to an allocentric spatial memory paradigm. During a single trial of the spatial change detection task, at presentation a participant will view an array of virtual objects in a random 3D spatial configuration. At test the participant will view the same stimulus either the same, or a rotated viewpoint and will be asked to judge whether an object has moved (or, later, which of a pair of objects has moved). We will manipulate the degree of spatial change between presentation and testing dynamically based on performance in order to acquire a measure of participants performance across a range of stimulus intensities. We will then model the data with a psychometric function with the aim to find the participant's spatial change detection threshold.

In the next sections, repeated components of the materials and methods described above will be explored in depth. These methods document repeated elements of experiments used throughout the following chapters (except where indicated). This includes the processes involved in creating stimuli, the trial-by-trial procedure and following this, the process of estimating thresholds from collected data.

2.3 Constructing the spatial stimuli

The program used to generate and present the scenes was developed and compiled in C# using Unity 3D, a graphics engine and coding environment used to create a variety of media including computer games and CG cinema.

The stimuli consist of an array of objects generated at the start of each trial. A similar procedure can be applied to any environment and items (for our initial experiments, described below, we used objects on a table-top). The stimulus generation process (see Figure 2.3) is as follows (for full reference to the original code please see https://github.com/edheyev/SCD_SpatialArrayGenerationTools).

First, a random seed is generated which defines all randomisation in this trial and allows for exact reproducibility if required.

A specified number of objects are given initial locations equidistant around the circumference of a circle, centred in the middle of the testing area.

Before being displayed, each object is subject to pseudo-random displacement and rotation in order to make the spatial organization of the array unique to each trial. Constraints are placed on how far and to what angle from the centre the objects can move (see Figure 2.3B). These constraints provide random configurations while avoiding overlap between objects, ensuring visibility of all objects and maintaining the approximate overall footprint and centre-point of the array. This constitutes stimulus 1 (i.e., before change).

To create stimulus 2, one object is chosen randomly and moved a distance specified by an adaptive staircase (or none are moved in the case of a 'same' trial in the 'same/different' experiment). The object is moved in a random direction. If the new location collides with any other object or is placed outside of the test area/ view of the scene camera, a new location will be chosen.

Stimulus 1 and 2 constitute a matched pair that are then shuffled at random so that one is used as the presentation stimulus and the other is used as the test stimulus.

In order to generate an image of the scene it is rendered from a specific viewpoint (a virtual "camera"). The camera can be rotated around the environment between presentation and testing to encourage the use of allocentric spatial information. The camera always points toward the centre of the array after movement.

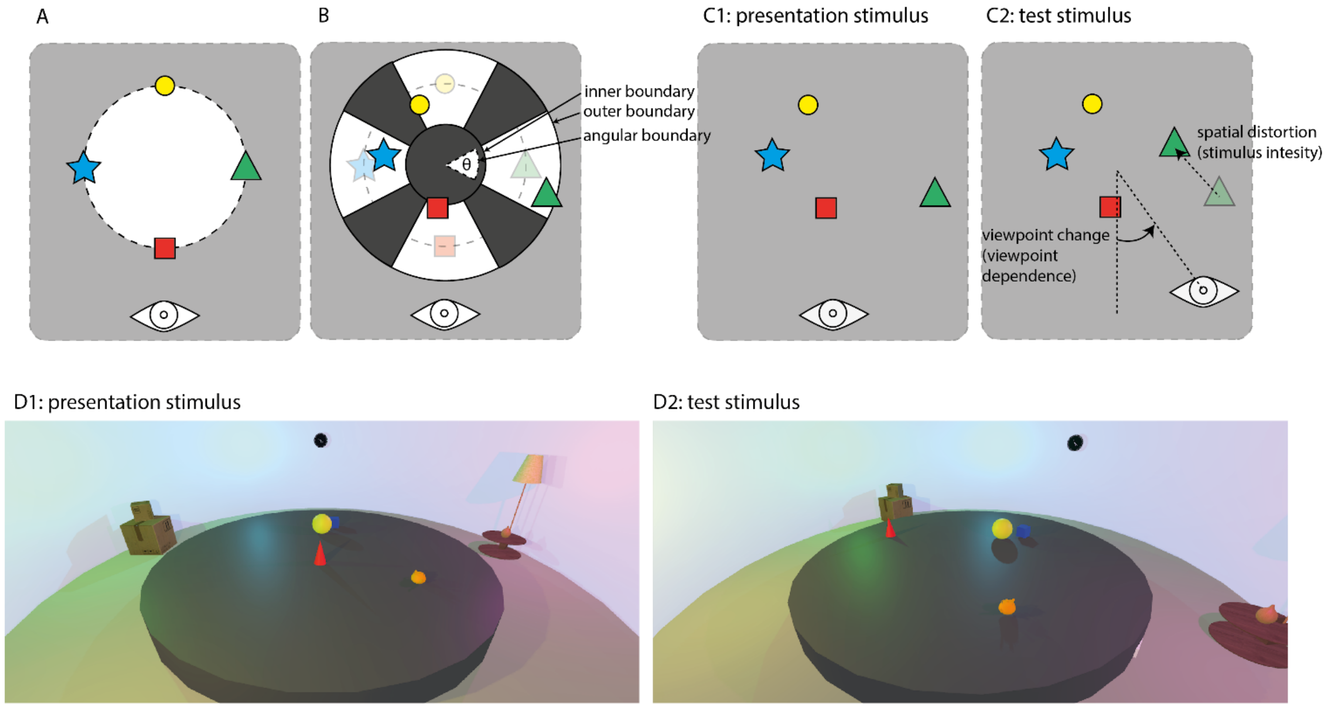

Figure 2.3 A diagram showing stimulus generation process and example screenshots of stimulus. A shows the initial placement of objects in equidistant placement. B demonstrates initial distortion to create a unique configuration, each object can move in the area defined by the inner and out boundary and the angular boundary (valid placement area shown in white). Faded shapes represent the old locations of the objects. C1 shows the presentation stimuli (the same as B, without example boundaries). C2 shows 'different' test stimulus, the green triangle is moved an amount defined by an adaptive staircase (spatial distortion). The viewpoint angle is also changed with the intent to bias the use of allocentric information. D1 and D2 show example screenshots from the actual task used later in the pilot study. D1 and D2 are the virtual equivalent of the diagrams in C1 and C2.

2.3.1 Adaptive Staircase Design

To collect data for a psychometric function, performance must be recorded over a range of stimulus intensity levels, ideally covering a saturated performance at both the most and least difficult condition respectively. There are a number of different methods used to systematically vary stimulus intensity so that performance can be plotted over range of stimulus intensity. Our approach uses an adaptive staircase design (Cornsweet, 1962) to vary the stimulus intensity – in this case the distance an object moves between presentation and test. This involves starting with a very high or very low stimulus intensity (i.e., very easy or very hard task) and modulating the intensity level dynamically based on the participants performance. In the most simple adaptive staircase, if a participant answers correctly, the task gets more difficult. If a participant answers incorrectly, the task gets easier. This allows for the greatest number of trials to be centred around the threshold (75% proportion correct) which is critical for our purposes. The staircase is completed when a set number of reversals is achieved, i.e., the staircase changes from ‘getting harder’ to ‘getting easier’ or vice versa due to consecutive correct and incorrect responses.

At the beginning of the procedure the size of the steps can be varied to more quickly adapt the task difficulty to an appropriate level based on the participant’s ability. For example, if the staircase begins at maximum or minimum intensity and we assume that the participants ability lies somewhere in the middle of the testing range, the initial steps, up or down, can be larger - skipping the extreme ends of the scale. This is very useful as the dependent variable (proportion correct) requires many data points in order to get a good estimate of ability and skipping unimportant trials allows for more efficient testing. The more data surrounding the participant’s approximate threshold, the more accurate the model will be at this critical point. This mechanism, designed to focus the majority of trials around the most critical stimulus intensity can thus improve the efficiency of the staircase procedure. Efficiency is extremely important in many psychophysical designs as they generally derive statistical power by using a large number of trials (Baker et al., 2020). Techniques such as these, which help to reduce the length of experiments, can counteract the negative consequences of long and mentally arduous tasks.

The criterion deciding traversal up or down the staircase also offers an opportunity to focus trials around the threshold level. By changing the number of correct/incorrect responses required to change intensity the trials can home in more rapidly and concentrate around the threshold. In tasks involving a binary decision, research has shown that employing a ‘3 down 1 up’ approach in which participants are required to get 3 correct answers before the task gets harder, this has been shown to better target the critical range of difficulty levels that is centred on the participant’s threshold of 75% correct (Brown, 1996; Levvit, 1970).

Together, these techniques allow us to design a 2AFC or Same/Different experiment that effectively and efficiently measures the threshold of a participant’s ability. Using an adaptive staircase, the task adjusts its difficulty level based on participant performance, dynamically changing to focus testing on the critical range required to provide the threshold level of 75% correct performance. We are left with data showing participant performance across the testing range, with the majority of trials focussed around the threshold.

2.4 Fitting the psychometric function

The raw trial data is processed by combining and averaging performance at each recorded level of the staircase. This results in a proportion correct at each level tested. In order to estimate a participant's spatial change detection threshold, we must fit a sigmoid shaped curve to model to the given data. We can then estimate the threshold of performance by finding the point on the curve where stimulus intensity corresponds to 75% performance (for example see Figure 2.4).

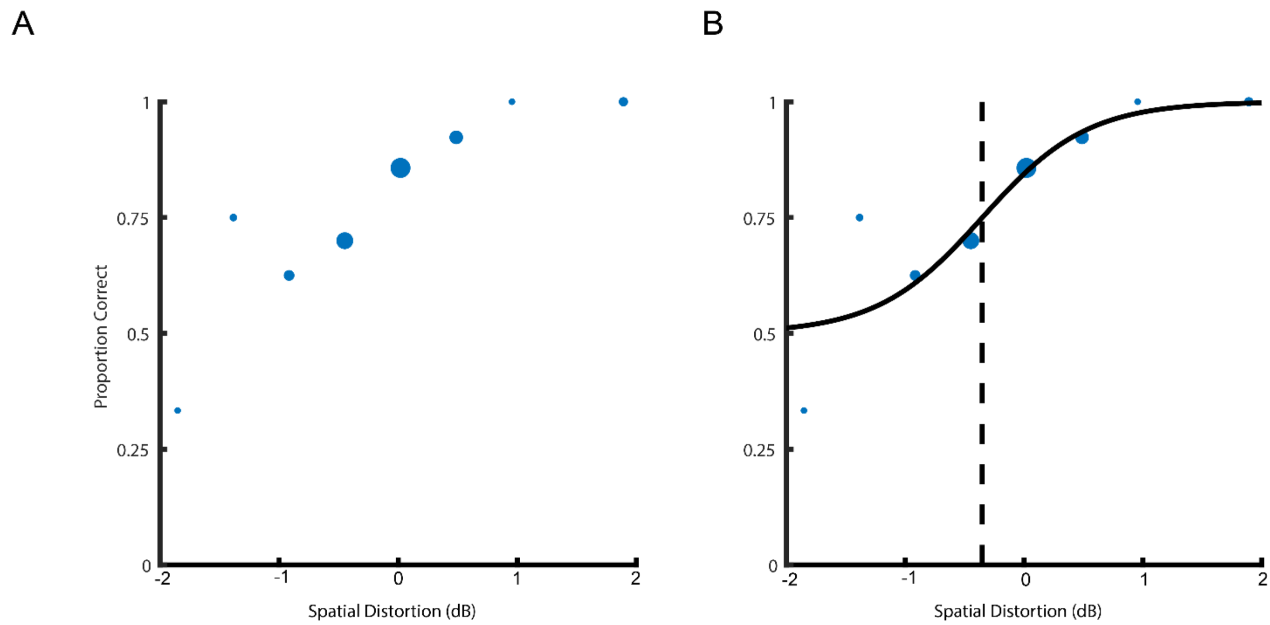

Figure 2.4 Demonstrating psychophysical curve fitting. A shows example proportion correct data across a test range. Note that this is not real data and for explanatory purposes only. B shows a logistic psychometric curve fit to the same data. The dotted line shows the spatial distortion at the threshold, corresponding to 75% performance. In this type of plot, the number of trials at each staircase level is indicated by the diameter of the circles; here the majority of trials are focussed on staircase levels near to the threshold.

The process of fitting the sigmoid curve to the proportion correct data that is used in this thesis involves finding the parameters (threshold and slope; x_0 and k) that maximise the likelihood of the data (x) under the model (f). To do this we use MATLAB's fminsearch procedure for estimating the parameters of a Logistic curve.

Equation 2.1

$$f(x) = 0.5 + 0.5(\frac{1}{1 + e^{-k(x-x_o)}})$$

The fminsearch procedure implements the Nelder-Mead simplex algorithm, iterating over parameter value combinations and minimizing a cost function. In this case the cost function is 1 – the log-likelihood of the data under the logistic model (Equation 2.1) given that iteration's parameters. The fitted curve describes the probability of a correct/incorrect response at any staircase level (x). Using this we can calculate the probability of the observed data, given the model on that iteration using

Matlab's binopdf function. The probability of each response is multiplied to get the overall likelihood and logs are used as the numbers are too small to calculate directly. The fminsearch procedure stops when the likelihood is maximised, resulting in the parameters that describe the best fitting model.

After the adaptive staircase has dynamically chosen the distance to move an object and then curve fitting and parameter estimation has generated a sigmoidal model that describes the data well, we are left with a model of participant performance across the testing range. Using this model we can estimate the point at which the intensity of a spatial change in a stimulus goes from not noticeable to noticeable. We can use the threshold value (x_0) to assess the observer's spatial memory precision.

If the data is too noisy or disorganized (for example, with proportion of correct responses varying non-monotonically with spatial distortion), the fminsearch procedure may not converge, or else the resulting fitted curve may be so flat that it yields a threshold estimate far beyond the testing range. For each change-detection experiment in the upcoming chapters, this case comprises an exclusion criterion.

2.5 A note on decibels

Throughout this thesis, the measurement of spatial change detection thresholds and the spatial distortion applied to individual stimuli, distances will be converted into decibels- a practice consistent with the traditions of psychophysics. The reason why decibels are preferable for many stimuli relating to perception is due to the Weber (or Weber-Fechner) law. The Weber Law describe two interrelating hypotheses which describe nature of the relationship between a physical stimulus and the accompanying subjective experience of perceiving that stimulus. Weber's law states that a change detection threshold is proportional to the pre-existing stimulus. For example, it would be very easy to detect a 100 lumen change in the light that a light bulb emits, it would be very difficult to detect a 100 lumen change in the light of the sun. Weber's law states that this is due to the intensity of the initial stimulus. Fechner's law adds to this, to say that that subjective perception is proportional to the *logarithm* of the physical stimulus intensity. Weber's law has been shown to be applicable across a very wide range of perceptual/sensory measures (Fechner, 1860; Levi & Klein, 1992; Link, 2020; Pardo-Vazquez et al., 2019; Weber, 1834).

Taking into account the Weber law, decibels work well to quantify stimulus intensity or change in intensity as they are both (1) a unit which is proportional to a certain reference value which is appropriate to the stimulus and (2) they also have a logarithmic progression:

Equation 2.2

$$dB = 20\ x\ Log10(\frac{spatial\ change(m)}{dB\ reference\ value})$$

The decibel calculation used here, and throughout this thesis (unless specified), uses a decibel reference value of 0.5m. This value was chosen as an approximate average change detection threshold from pilot and pre-pilot data. Thus, if a participant has a spatial change detection threshold of 0dB it means that, on 75% of trials, they can identify a change to the stimulus when an object has been moved by 0.5m.

2.6 Pilot

After coding the task and analysis, and extensive prepilot work, the final stage of development entailed piloting the paradigm with a sample of naïve participants in a formal experimental setting. Below is documented the pilot studies for both precision and capacity investigations. This involved fine-tuning parameters that govern elements of the psychophysical procedure (e.g., step size, high and low intensity values determining the testing range) and visual aspects of the experimental environment (view distance, object size, angle of inclination, all of which affect the available range of stimulus configurations and spatial distortions). Additional pilot testing, specific to forthcoming experiments described later in the thesis, was also carried out and is documented in the corresponding chapters where appropriate.

2.6.1 Aims

In order to finalise the methods for use in further experiments it was necessary to achieve a number of fundamental features:

- Ensure data is described well by the psychometric function.
- Set testing range for ‘saturated performance’ at both ends of the psychophysical curve.
- Ensure that the staircase procedure converges on a threshold level of performance.

These three elements would indicate that the methods were working as intended.

2.7 Pilot 1: Precision

2.7.1 Methods

2.7.1.1 Participants

A sample of n=6 was used for Pilot 1. Participants were students at York University that applied to the experiment through the University of York student participation pool.

2.7.1.2 Procedure

In Pilot 1 we investigate the precision of memory using spatial change detection thresholds with an array of fixed set size (4 items). A Same/Different spatial memory task was used. Participants were presented a virtual array of 4 objects on a virtual table-top environment for 6 seconds (stimulus 1). This was presented using a standard desktop computer. After a 1 second black screen, participants then saw stimulus 2 for 6 seconds- the same array from a rotated viewpoint, 60 degrees around the table (clockwise or anticlockwise). In 'Different' trials, one object had moved from its original position. In 'Same' trials the spatial configuration of objects was identical in stimulus 1 and stimulus 2. The participant was then asked if the stimuli are 'Same' or 'Different' and given 8 seconds to respond. Response is via a key press 's' for Same and 'd' for Different.

In the Precision Experiment (Pilot 1), the independent variable controlled by the adaptive staircase was spatial distortion, i.e., the amount that an object moves in the different trial and the dependent variable is threshold of precision. A higher intensity stimulus, with an easier to detect change, has a large spatial distortion.

2.7.1.3 Precision: Testing parameters

In Pilot 1, an adaptive staircase governed spatial distortion on each trial (stimulus intensity). The testing range was a log scale of 9 values between 0.27 and 1.5m.

Table 2.1 Showing the distance an object moves (stimulus intensity) at each staircase level. This is given in meters and converted to decibels.

Spatial Distortion									
Meters	0.27	0.33	0.41	0.51	0.63	0.79	0.97	1.21	1.50
Decibels	-5.35	-3.61	-1.72	0.17	2.01	3.97	5.76	7.68	9.54

The experiment finished after 10 reversals and staircase step sizes changed after each reversal. The step sizes were 4,4,2,2,1,1,1 respectively (for example, the first step was 4 stairs and after it had changed direction twice, it would jump 2 stairs and so on). For each condition, one staircases was used, beginning at the lowest difficulty (highest intensity change).

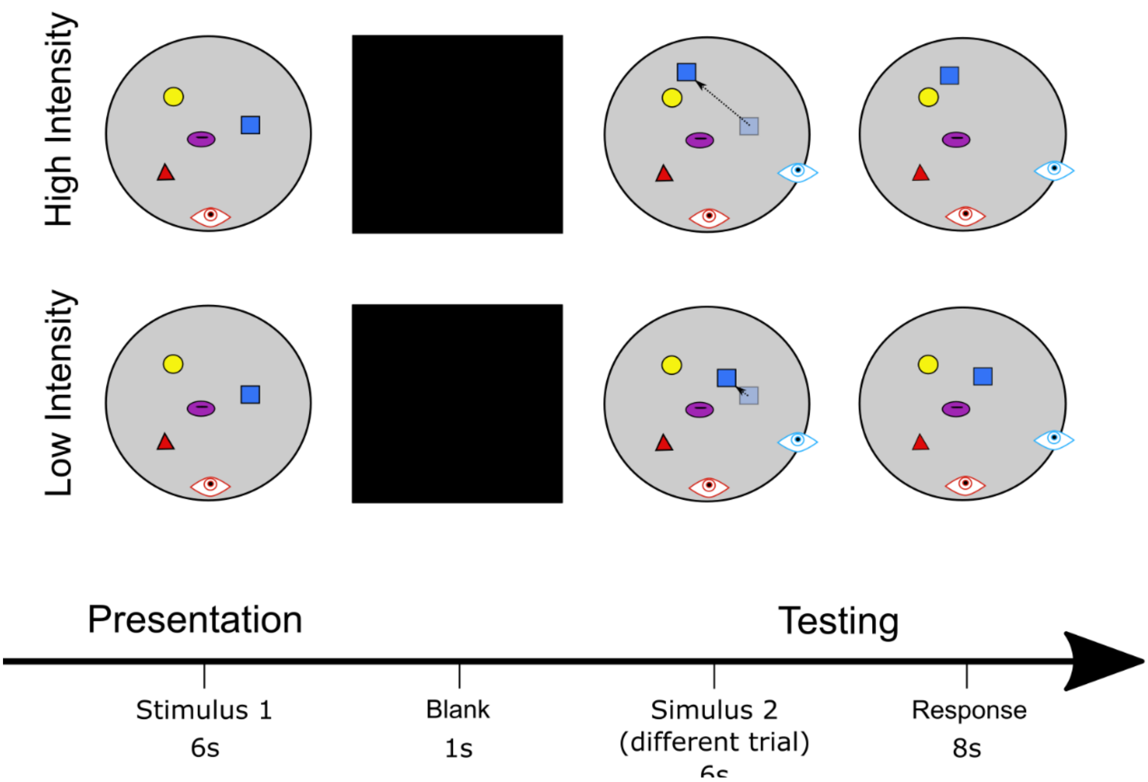

Figure 2.5 Diagram showing the format of a trial in the 'different' trial of the precision based change detection task. Stimulus 1 shows a random configuration of objects to participants for 6 seconds. Then the screen dims for 1 second. Stimulus 2 is then shown to participants for 6 seconds, in the different condition, one object has been moved. In a high intensity trial the object moves a larger distance, in a low intensity trial the object moves very little. Finally, the participant is asked to respond 'same' or 'different' by pressing the 's' or the 'd' key on the keyboard, they are given 8 seconds to do this.

2.7.2 Results

Responses were scored as hit, miss, false alarm and correct rejection. For each staircase level tested, a 'proportion correct' was calculated (Hits + correct rejections / all responses).

As described above, a psychometric curve was fit to the proportion correct responses across the test range and a threshold was estimated for each participant at the 75% performance level. Corresponding threshold levels in decibels were recorded.

In all participants, a psychophysical curve was successfully fit to the data. All thresholds were estimated to be within the testing range, however, there is a diverse range of thresholds which covers much of the testing range.

Table 2.2 Summary statistics for each participant in Pilot 1. Showing the number of trials and the estimated best parameters for the fit psychophysical curve (threshold in dB and slope).

ID	# Trials	Threshold(dB)	Slope
1	91	6.05	1.34
2	88	3.66	3.85
3	85	3.30	0.88
4	86	3.42	1.28
5	93	1.86	4.03
6	95	2.95	2.27

As shown in Figure 2.6, while we are able to fit sigmoid functions to the data, we are not achieving saturation at either end of the testing range with all participants.

The average number of trials was 84.7 trials per participant which resulted in reasonable curve fits. The adaptive staircases seemed to work well and looking at the right hand graphs in Figure 2.6, we can see the trial staircase level covering a fairly wide testing range but ultimately focussing on a portion of the available difficulty levels. This indicates that the curve fitting procedure would have most information around this point.

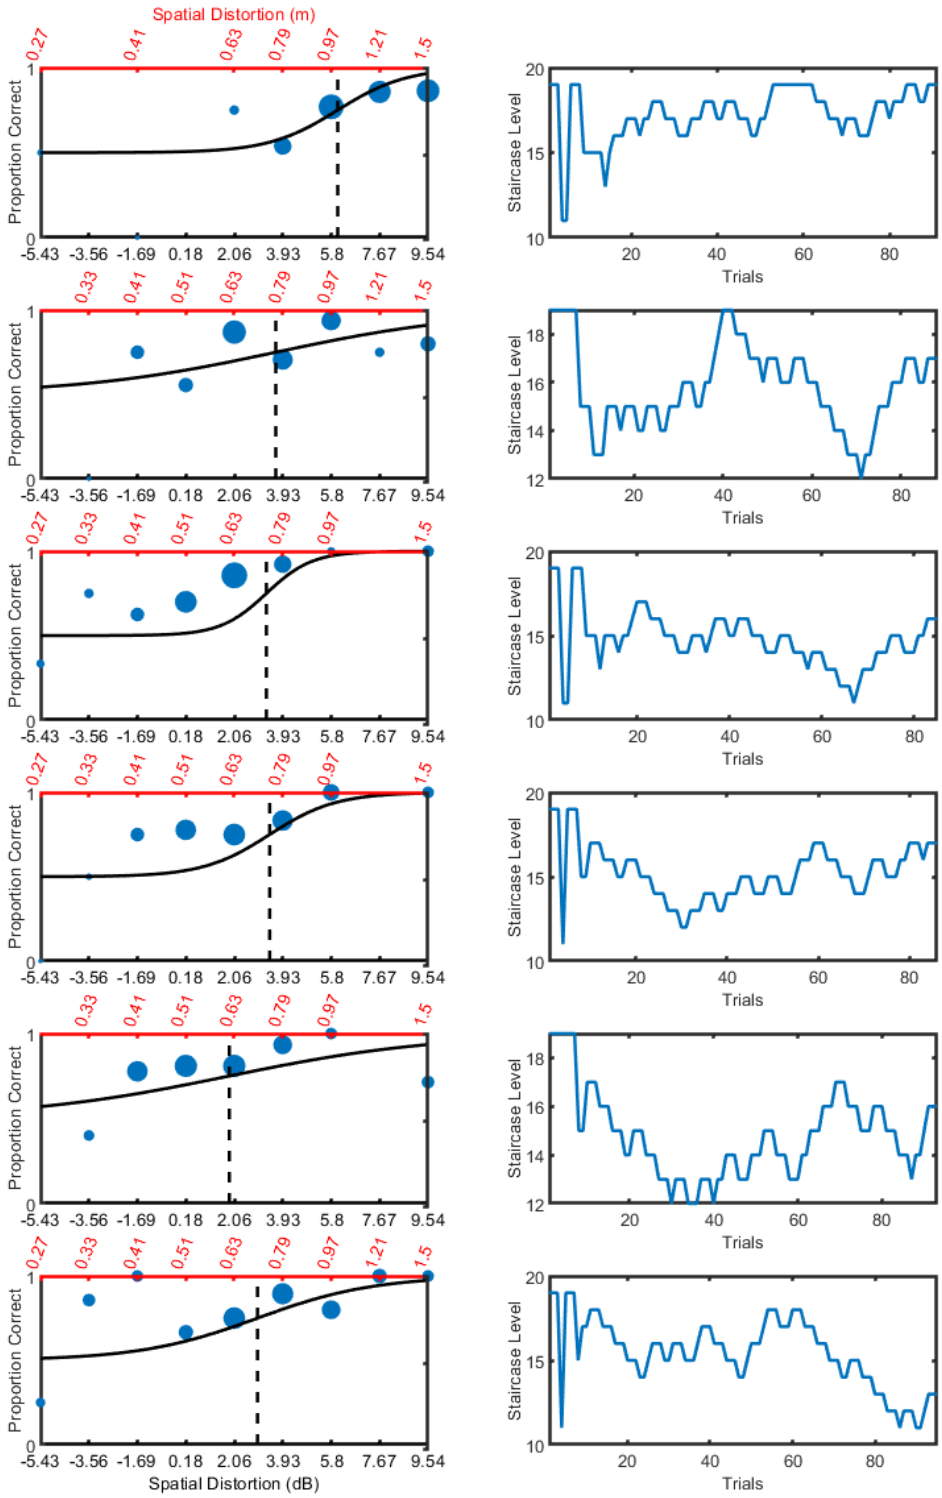

Figure 2.6 Depicting individual participant performance in Pilot 1 . Each row is a separate participant. The left graph is the psychophysical curve fit to proportion correct data. The x axis' shows spatial distortion, the independent variable manipulated by the adaptive staircase. The y axis shows proportion of trials correct at each level. Point size refers to the number of trials probed at that staircase level. The right graph is the progression of the staircase through the experiment, adapting to participant performance. The staircase begins at the highest level at the start of the experiment. The staircase ends after 10 reversals.

2.7.3 Discussion

Regarding the aims of this pilot, the lack of saturation at either end of the testing range suggests that it is necessary to expand the range further. By adding more staircase levels, this will have the added benefit of providing additional data points, leading to a more accurate model.

Secondly, while all thresholds were within the testing range, there was considerable variation between individual participants. This is an indication of the between subject variance in performance and is another indicator that a wider testing range would help to ensure that participants with an unusually good or poor performance can still be reliably tested.

Apart from this change, most elements of the method worked well and will be carried forward. The set-size seemed appropriate and object types were recognisable and easily distinguishable. Based on participant feedback, we know that the environmental layout was clear and meant that participants recognised when the viewpoint rotation took place. Instructions to participants were clear and they understood the task.

There are several points related specifically to the same/different task that have become apparent during this pilot. Firstly, a disadvantage of the same/different method is that the 'Same' trials do not meaningfully contribute to proportion correct statistics at any given staircase level and do not inform us of participant ability resulting in a prolonged experiment. Against this, one possible advantage is that, as we explicitly ask participants to attend to the entire array rather than targeted objects, this was intended to favour the use of allocentric spatial memory by encouraging participants to encode the configuration of the entire array rather than focussing attention on specific objects. An alternative approach, perhaps less well suited to allocentric memory, would be to change the array on every trial and ask participants to identify which object had moved between presentation and testing (for example, offering a two alternative forced choice by highlighting two objects and asking them to select the one that was moved). These issues will be explored in depth during the next chapter where we formally compare different psychophysical methods in a viewpoint rotation task.

2.8 Pilot 2: Capacity

2.8.1 Methods

In the Capacity Experiment (Pilot 2), the independent variable controlled by the adaptive staircase was the number of objects in the array and spatial change was held constant. For example, it was expected that easier to detect change had fewer items and a harder to detect change contained more items. The task would get more or less difficult dependent on participant performance.

As in the Pilot 1, in all trials, viewpoints were rotated 60° around the centre of the array between presentation and test.

2.8.1.1 Participants

A sample of n=6 was used for Pilot 2. Participants will be students at York University that applied to the experiment through the University of York student participation pool.

2.8.1.2 Procedure

As in Pilot 1, a same/different spatial memory task was used. Participants were shown a virtual array of objects on a virtual table-top environment for 6 seconds (stimulus 1). This was presented using a standard desktop computer. After a 1 second black screen, they then saw stimulus 2 for 6 seconds, on 'Same' trials this was the same array from a rotated viewpoint, 60 degrees around the table (clockwise or anticlockwise at random). On 'Different' trials, one object had moved from its original position and the viewpoint was again rotated 60 degrees around the table. In each case, the participant was then asked if the stimuli are 'Same' or 'Different' and given 8 seconds to respond by pressing the 's' key to indicate 'Same' and 'd' for 'Different'.

2.8.1.3 Capacity: Testing parameters

The stimulus arrays had a set size of between 3 and 10 items. There are 7 staircase levels in total each corresponding to a set size (e.g., 3, 4, 5, 6, … 10 items).

Two staircases were used for each condition. One starting at the highest, and another at the lowest difficulty level. The staircases should ideally meet at the threshold level. The experiment completed when each staircase has reached 7 reversals. As the staircase controls the number of objects, the amount of distortion was fixed at 0.5m (an approximation of threshold performance across trials found in pre pilot experiments. As with Pilot 1 the step sizes are 4,4,2,2,1,1,1 respectively.

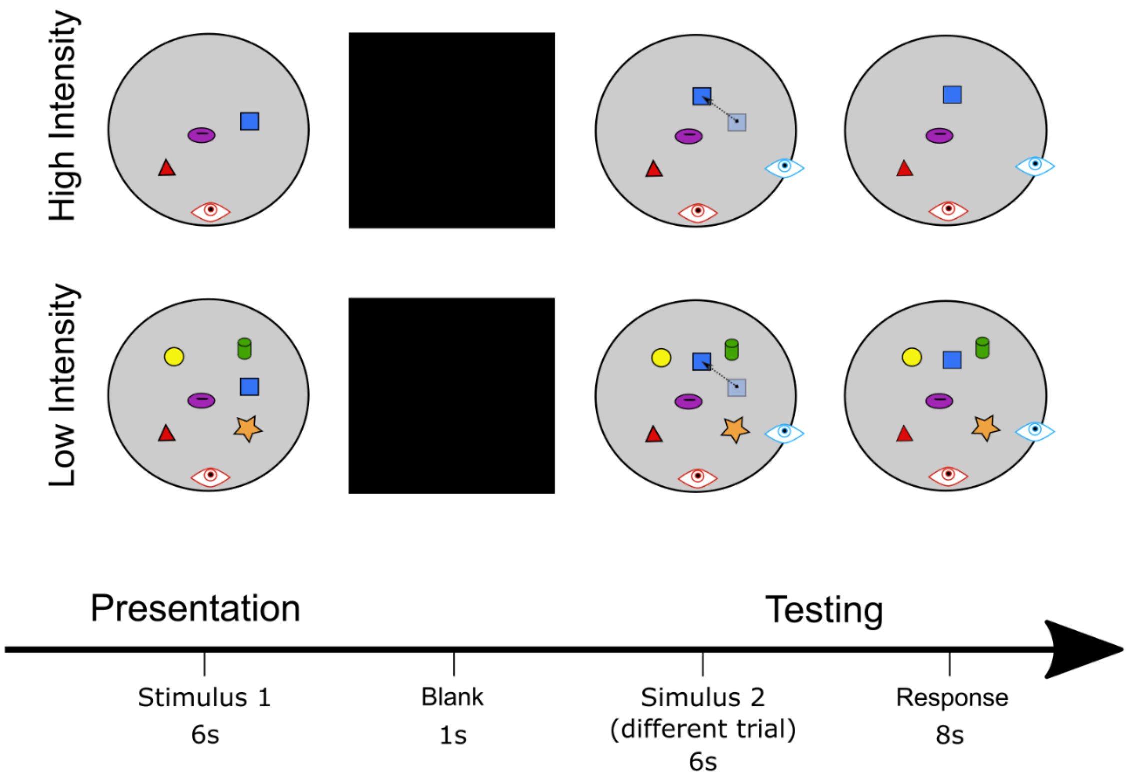

Figure 2.7 Diagram showing the format of a 'different' trial in the capacity based change detection task . Stimulus 1 shows a random configuration of objects to participants for 6 seconds. Then the screen dims for 1 second. Stimulus 2 is then shown to participants for 6 seconds, in the different condition, one object has been moved. In a high intensity trial there are very few items in the array, in a low intensity trial there are very many items in the array. At 'Response' the participant is asked to respond 'same' or 'different' by pressing the 's' or the 'd' key on the keyboard, they are given 8 seconds to do this.

2.8.2 Results

As in Pilot 1, responses were scored as hit, miss, false alarm and correct rejection. For each staircase level tested, a 'proportion correct' was calculated (hits + correct rejections / all responses).

Again, a psychometric curve was fit to the distribution of proportion correct statistics across the test range and a threshold level of performance was estimated for each participant at the 75% performance level. Corresponding threshold levels in decibels relative to 0.5m (the approximate average performance during pre-pilot testing) were recorded.

Table 2.3 Summary statistics for each participant in Pilot 2. Showing the number of trials and the estimated best parameters for the fit psychophysical curve (threshold in dB and slope).

ID	# Trials	Threshold	Slope
7	103	-4.74	2.20
8	97	-2.71	1.34
9	89	0.86	0.74
10	98	-3.01	4.09
11	96	-3.44	2.21
12	88	-4.36	2.95

In terms of the three criteria defined earlier, the data seems fairly well described by the psychometric function although there seems to be more noise in the distribution than in Pilot 1 and with much flatter slopes. We are not able to get a saturated performance in all participants. This looks to be due to participants not reaching a difficulty that corresponds to chance performance (i.e., the most difficult trial; 3 items moving 0.5m is too easy). Finally, due to the flatter curves, in some cases, thresholds were estimated outside of the testing range. Threshold levels of performance were highly variable between participants.

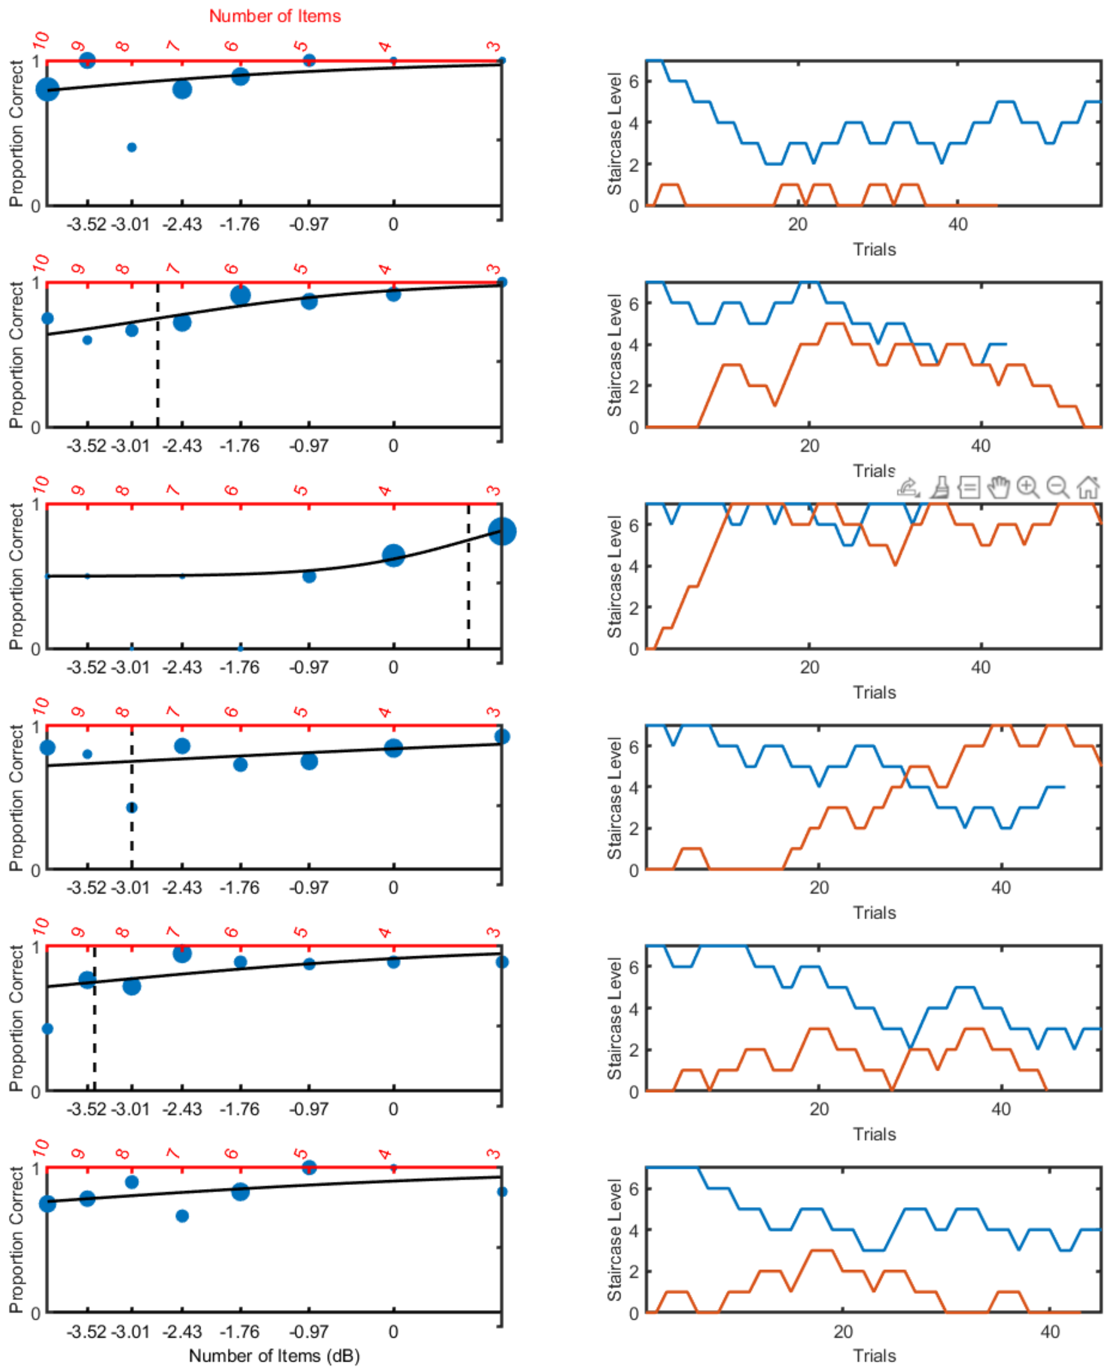

Figure 2.8 Depicting participant performance in Pilot 2. Each row is a separate participant. The left graph is the psychophysical curve fit to proportion correct data. The x axis' shows set size, the independent variable manipulated by the adaptive staircase. The y axis shows proportion of trials correct at each level. Point size refers to the number of trials probed at that staircase level. The right graph is the progression of the two staircases through the experiment, adapting to participant performance. One staircase begins at the highest and lowest level at the start of the experiment. The staircase ends after 7 reversals.

2.8.3 Discussion

The results from the capacity experiment show that the task, in its current form, is much less well suited to the capacity version of the experiment. This is evidenced by more noise in the data resulting in poorer fitting psychophysical curves and also insights from participant comments.

It seems that the issue may be that the task difficulty, as a function of the number of the objects in the array, is not monotonic. That in some cases more items in the array makes change detection easier.

This makes sense in terms of the literature. More items in the array means smaller distances between objects. Research has shown that object-object vector coding cues are more effective at shorter distances (Aagten-Murphy & Bays, 2019) and that shorter distance between array elements also seem to increase the use of allocentric spatial memory- something useful in this viewpoint rotation task. Furthermore, more items, closer together also creates more opportunities for familiar geometries to emerge from the constellations of items. Items may line up in a straight line or form near-regular shapes. These are unintentional cues which would facilitate memory of the spatial environment through crowding or chunking (Sargent et al., 2010; Tamber-Rosenau et al., 2015). For example, the movement target items which are part of a familiar shape and then are distorted and break the configuration would be much easier to notice than an item which is not associated with an easily recognisable geometry.

Whether spatial memory capacity should be expressed in terms of a threshold number of items is a complex issue. Object location is presumed to be encoded relative to landmarks in the environment but in the current experiment, the virtual environment is very sparse and so (within allocentric representations) it is very likely that object location is encoded relative to other objects. The notion and definition of capacity, in this context, quickly breaks down. Is the spatial memory capacity referring to the number of individual items stored? The number of spatial relationships? The amount of information? These complications show that perhaps it is naïve and overly simplistic to presume that more items within a representation equates more capacity.

It is also questionable whether it is useful to characterise a spatial memory capacity threshold as a specific set size. If spatial memory behaves as working memory, it may be the case that participants will still be able to perceive and recall some spatial locations past their capacity, possibly at a lower precision (Bays, et al., 2009). Using this method it is difficult to investigate the possible interaction between capacity and precision.

2.9 General discussion

The results from the pilot studies suggested that a change detection task has the potential to provide robust measures of allocentric spatial memory limits. The adaptive staircase worked well to focus the task difficulty on a portion of the testing range, meaning that the task was efficient and that the resulting psychometric curve fits should be most accurate on around each participants change detection threshold.

However, it was also made apparent that the quality of the data can be improved. Subsequently this pilot led to a number of changes and adaptations to our spatial change detection task.

Regarding precision: the psychometric modelling of the precision trials generally showed better fits and encouraged us to continue using this method. However, there were still improvements that can be made. The lack of saturation on each side of the curve in some cases suggest that we need to widen the testing range to include both more difficult and easier trials. The noise present in the results suggested that we needed to significantly increase the length of the experiment. To improve the definition of the psychophysical curves we will also include more step increments. In future experiments we will also include a formal analysis to assess the goodness of fit of each psychometric curve.

Regarding capacity: our pilot study revealed complications with our initial approach to assessing capacity in spatial memory representations. These complications make the current approach of varying set size using an adaptive staircase and finding a threshold ‘number of items’ undesirable. Taking this important finding into account, the initial focus of the thesis will be on further developing and validating precision measurements under varying conditions (Chapters 3 and 4). Capacity will be returned to later in the thesis and explored in its own right using adapted methods (Chapter 5).

An observation relevant to both pilot experiments is that there were cases where saturation at 100% performance was not achieved. While increasing the testing range and providing a more intense stimulus change for participants will be implemented and may help this, there is a possibility that the task itself can never be made easy enough to reliably achieve 100% performance.

2.10 Conclusions

The methods used in this thesis are a combination of techniques designed to target allocentric spatial memory representations and at the same time measure performance with a high degree of accuracy.

A viewpoint rotation paradigm that prevents automatic spatial updating through unpredictable movement potentially allows us to isolate contributions from a viewpoint-independent memory representation. Techniques from psychophysics such as adaptive staircases and threshold estimation allow us to quantify each participant’s ability level at each tested viewpoint rotation. Using a real time game engine to generatively create stimuli for each trial means that these methods can be abstracted to many different situations. By measuring memory change detection thresholds in this precise and controlled way we can characterise the quality of ego- and allocentric memory. We can look at the effect of viewpoint changes on memory precision and try to better understand how ego- and allocentric memory systems interact. We can also apply the same methods in different environments and examine context specific qualities of these distinct types of memory and their respective suitability to different tasks and settings. Finally, we can also take inspiration from visuospatial working memory studies and look at ego- and allocentric memory in terms of their capacity, investigating whether, for instance, when the number of item in an array increases, the precision at which we remember them decreases.

With some minor parameter adaptations, in the next chapter I will use this method to systematically investigate the relationship between spatial memory precision and viewpoint change and interactions between ego- and allocentric systems, using both the current Same/Different task and a two-alternative forced choice (2AFC) variant.

Chapter 3. Testing the precision of spatial memory representations using a change-detection task: Effects of viewpoint change

3.1 Abstract

There are at least two distinct ways in which the brain encodes spatial information: in egocentric representations locations are encoded relative to the observer, whereas in allocentric representations locations are encoded relative to the environment. Both inform spatial memory, but the extent to which they influence behaviour varies depending on the task. In the present study, two preregistered experiments used a psychophysical approach to measure the precision of spatial memory while varying ego- and allocentric task demands. Participants were asked to detect the changed location of one of four objects when seen from a new viewpoint (rotated by 0°, 5°, 15°, 45° or 135°). Experiment 1 used a Same/Different task and Experiment 2 used a 2AFC task. Psychophysical thresholds were calculated, showing that in both experiments, spatial change detection thresholds showed a monotonic but non-linear increase as viewpoint change increased. This was consistent with a preregistered model including distinct parameters corresponding to egocentric and allocentric contributions that change lawfully as a function of viewpoint shift. Our results provide a clearer understanding of how underlying memory representations interact to inform our spatial knowledge of the environment.

3.2 Introduction

Spatial memory allows mobile organisms to act adaptively in complex environments that extend, temporally and spatially, beyond the range of their senses. The way in which the environment and salient objects and locations within it are represented in memory can take different forms, each of which may be suited to specific situations. Extensive neuroscientific and neuropsychological evidence indicates a distinction, in humans and other mammals, between egocentric and allocentric representations (see e.g., Burgess, 2008; Ekstrom et al., 2014; Galati et al., 2000; Hartley et al., 2014). Egocentric representations are supported by neurons in parietal, sensory and motor cortices, which have spatial receptive fields that are anchored with respect to the body or parts of the body. Allocentric representations are supported by neurons in the medial temporal lobe with spatial receptive fields that are anchored to the environment, and thus insensitive to changes in viewpoint. It is argued that these different forms of spatial representation have distinct functional advantages, with egocentric representation being well-suited to guiding action in the immediate environment over short timescales, and allocentric representation being well-suited to long-term memory and navigation (Burgess, 2006; Goodale & Milner, 1992;).

Although they may make distinct functional contributions to spatial behaviour, ego- and allocentric representations are unlikely to be wholly independent; as spatial information is initially encoded by perceptual systems in egocentric form, it must be translated into allocentric form in order to represent environmental locations in memory. Conversely when motor systems need to draw on allocentric information in the control of behaviour, it first must be converted *back* to an egocentric representation before being used. Mechanistic neural-level models are able to explain this hierarchical interaction by postulating an intermediate transformation circuit (in retrosplenial cortex) which mediates the translation (Bicanski & Burgess, 2018; Byrne et al., 2007).

From a psychological perspective, the coexistence and likely interdependence of ego- and allocentric systems makes it challenging to interpret the source of information present in spatial behaviour since it can be difficult to isolate their contributions to a given task. Yet understanding, distinguishing and quantifying these contributions is potentially important, not just for theories of spatial memory, but also in a clinical context where failures of spatial memory can be symptomatic of underlying disease processes. Notably, because Alzheimer's Disease typically affects the medial temporal lobe in its earliest stages, a decline in allocentric spatial memory may give an early indication of disease pathology (Hartley et al., 2005; Moodley et al., 2015; Serino et al., 2014).

Previous work has established that both forms of representation can contribute to performance on laboratory tasks involving memory for objects' locations. A common approach has been to ask participants to remember the locations of an array of objects before being asked to complete a spatial task, such as indicating which object has moved, or pointing to an object's previous location, from a novel viewpoint (Burgess et al., 2004; Easton & Sholl, 1995; Greenauer & Waller, 2008; Holmes et al., 2018b; Holmes & Sholl, 2005; King et al., 2002; Mou et al., 2006b; Negen, Heywood-Everett, et al., 2018; Newcombe & Huttenlocher, 1992; Simons & Wang, 1998; Waller & Hodgson, 2006; Wang & Simons, 1999; Wraga et al., 2000, 2005; Zhang et al., 2011). These tasks are revealing due to the way viewpoint changes differentially affect different forms of memory representation. Specifically, in order to remain valid while the objects are obscured, egocentric representations of unseen targets must be updated to reflect participants' bodily movements. This spatial updating process relies on idiothetic self-motion cues from visual, vestibular and proprioceptive information.

By making these cues unavailable or invalid it is possible to selectively disrupt viewpoint-dependent contributions to spatial memory from the egocentric system. By contrast, because viewpoint-independent representations encode locations with respect to the environment, the information in the allocentric system should remain valid even in the absence of spatial updating. Studies show enhanced performance when the appearance of the array of objects is consistent with spatial updating cues (i.e., a viewpoint-dependent contribution to spatial memory; Wang & Simons, 1999), but they also indicate a viewpoint-independent contribution where spatial updating is unavailable (Burgess et al., 2004). Note that while in real-world tasks, the physical sensations of motion undoubtedly play an important role in spatial updating, the distinct contributions of viewpoint-independent and viewpoint-dependent memory can be inferred in purely visual tasks using virtual environments presented on computer displays (King et al., 2002; Lambrey et al., 2012; Spiers et al., 2001; Tu et al., 2017). For example, in King et al. (2002) a patient with damage to the hippocampus showed profound impairment of memory for object locations in a virtual environment that was apparent only when tested from a different viewpoint from that used at encoding.

Beyond the simple dissociation of viewpoint-dependent and -independent forms of memory, there are clear indications that performance in spatial memory tasks can be parametrically affected by the *degree* of viewpoint change: notably, the spatial accuracy of responses is reduced as viewpoint at test is rotated away from viewpoint at encoding (Diwadkar & McNamara, 1997; Shelton & McNamara, 2001) with increased viewpoint rotations having been found to encourage the use of allocentric cues and activation of allocentric-associated brain regions (Schmidt et al., 2007; Zhao & Warren, 2015). These results suggest that a more detailed understanding of the contributions of ego- and allocentric systems to behaviour can be gained by investigating the relationship between the degree of viewpoint change and the accuracy of spatial responses in memory tasks.

This parametric approach connects with recent developments in investigations of working memory where the precision of responses is increasingly used to investigate capacity constraints. Some authors argue that visual working memory capacity is limited by a common resource that, rather than being limited to a fixed number of items, can be allocated to the fidelity with which each item's properties (such as colour and location) are stored (Alvarez & Cavanagh, 2004; Awh et al., 2007; Bays & Husain, 2008; Bays et al., 2009; Ma & Wilken, 2004). In this context, psychophysical techniques provide a sensitive means to measure the precision of memory representations in visuo-spatial change detection tasks (e.g., Brady & Alvarez, 2015; Jiang et al., 2000; Luck & Vogel, 1997; Luck et al ; Luck & Vogel, 2013, see Brady et al., 2011 for review). However, spatial working memory tasks have typically been concerned with 2D stimulus arrays in which the participant's viewpoint remains fixed;

in most cases such tasks can shed little light on the way in which different forms of spatial representation contribute to working memory.

A notable exception is a recent study by Aagten-Murphy and Bays (2019), which investigated the precision of spatial memory responses to 2D arrays when a persistent landmark was present in the display. In this paradigm, although the participant's viewpoint is fixed, the landmark can be thought of as providing an allocentric cue (i.e., allowing for the encoding of each object's spatial relationship to its surroundings), which - when it is present and stable - can improve spatial precision, or – if it is shifted – can be placed in conflict with egocentric memory. Consistent with this, stimulus locations near to a landmark were remembered with greater precision than locations further away from a landmark when it remained persistent and stable, and responses were shifted when the landmark was moved between presentation and testing. The results were interpreted as indicating independent contributions of ego- and allocentric representation to the precision of spatial responses, and the authors developed a mathematical model describing how ego- and allocentric contributions are combined in the precision of responses in their task.

These intriguing results align with the neuroscientific and psychological research described above in suggesting that ego- and allocentric representations make distinct contributions to spatial memory. However, the precision of spatial memory has yet to be investigated with 3D arrays in more realistic situations in which viewpoint varies systematically between presentation and testing. In these circumstances, we might expect responses to be guided principally by viewpoint-dependent egocentric representations when the viewpoint remains fixed, with an increasing contribution of viewpoint-independent allocentric representation with larger unpredictable viewpoint shifts between encoding and retrieval.

To characterise the relationship between viewpoint change and the precision of spatial memory, in the current study we develop a change detection task in which participants remember the locations of an array of objects in a virtual environment over a brief delay. We adopt a psychophysical approach to measure spatial change detection thresholds while parametrically varying the degree of viewpoint change between presentation and testing. We then fit a preregistered model to participant data in order to distinguish and quantify ego- and allocentric contributions to spatial precision as it varies with viewpoint.

3.3 Methods

3.3.1 Overview

In two experiments, described in more detail below, participants were asked to make judgements about the locations of objects in an array in a virtual environment displayed on a standard PC monitor (see Figure 3.1A). In each task, participants first viewed the object array for six seconds (presentation) from a standard viewpoint and then, after a brief delay, saw the objects from a new viewpoint before making a spatial judgement (testing). To probe the effect of viewpoint changes between presentation and testing, the viewpoint used to render the scene at testing was manipulated across trials (see Figure 3.1B), with the degree of viewpoint change being selected at random on each trial. Thus, participants could not anticipate at presentation the viewpoint they would encounter at testing.

Experiment 1 (Figure 3.1B) used a Same/Different (S/D) spatial change detection task. Between presentation and testing, the array was left unchanged on half of trials, selected at random ('Same' trials), or (for the remaining 'Different' trials) a single object was moved by a distance controlled by an adaptive staircase (described below), with participants being asked to respond 'Same' or 'Different'.

Experiment 2 (Figure 3.1C) used a two-alternative forced choice (2AFC) spatial change detection task. Between presentation and testing, an adaptive staircase controlled the degree of spatial distortion (the distance a single selected object was moved) applied to an array of 4 objects.

Our objective in employing two experiments was to provide a degree of internal replication, to determine whether results would generalize to different spatial judgements and to compare tasks, thus both experiments were piloted and preregistered prior to any data collection and the order of reporting below is arbitrary (see https://osf.io/p6gzt/ and https://osf.io/j7a8p/).

3.3.2 Participants

All participants were recruited via the University of York's Department of Psychology participant recruitment system. For Experiment 1, twelve subjects were recruited and two participants were subsequently excluded (see exclusion criteria, below), leaving ten participants in the analysed data (10 female, mean age 19.7, SD 1.16). For Experiment 2, thirteen participants were recruited and three participants were subsequently excluded leaving ten participants in the analysed data (10 female, mean age 19.5, SD 0.71). Participants received course credit for participation and gave their written informed consent in line with guidance from the University of York Department of Psychology Ethics committee.

3.3.3 Task and Materials

All testing took place in a quiet testing room in Department of Psychology, University of York. Task and stimuli were rendered and controlled by custom software programmed in the Unity3d game engine in C# running on an HP Z400 computer and displayed on a 20” ‘Elite Display 232’ monitor. The Unity engine allows us to specify environment and object locations in realistic 3D coordinates that are then accurately rendered to the screen as if seen from a particular viewpoint (see Figure 3.1A).

The virtual environment (see Figure 3.1A) consisted of a 10m diameter, circular room lit with contrasting colour lighting and containing generic objects next to the walls to provide directional cues. Participants viewed a 4m diameter table-top from a distance of 1m (i.e., 3m from centre). When the participant viewed the array from a rotated angle, the degree of rotation was either 0, 5, 15, 45 or 135 degrees relative to the initial viewpoint which remained fixed across all trials. The viewpoint rotations were chosen to avoid rotations that might be susceptible to non-spatial or idiosyncratic strategies – (for example advantages at 90 and 180 degrees; Mou & McNamara, 2002).

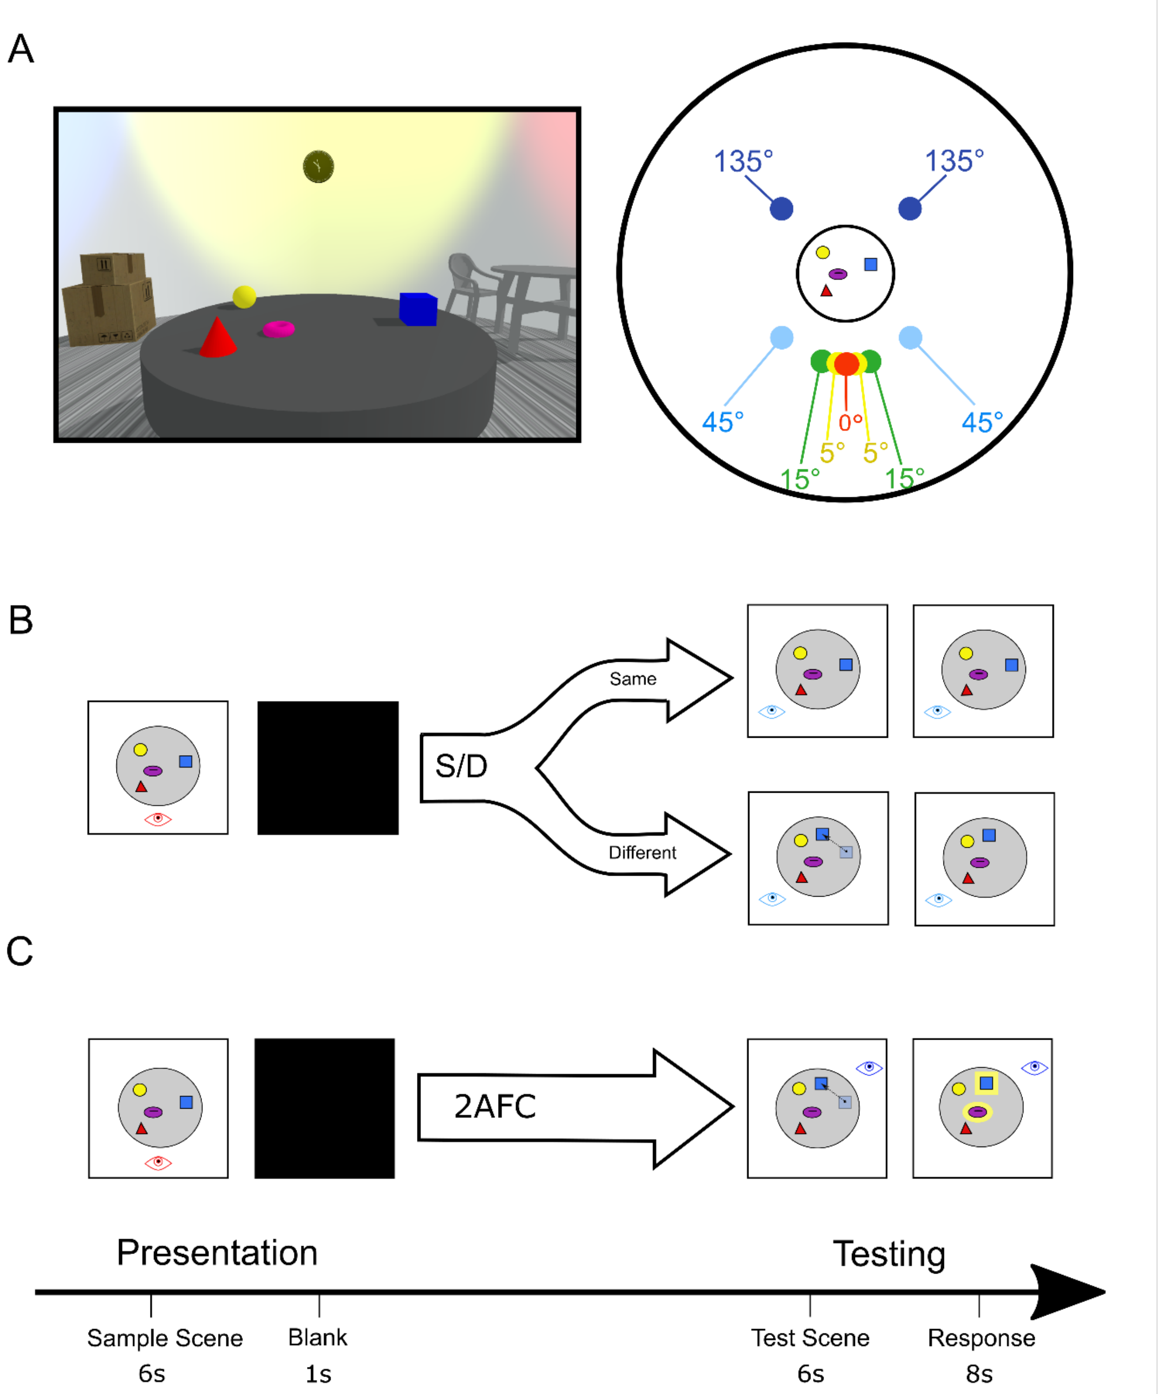

Figure 3.1 Chapter 3 task and Materials. A (left) The testing environment from participant perspective at 0° view rotation used at Presentation (right) environment layout and viewpoints; coloured circles mark the viewpoints used at Testing. Red: 0° rotation, yellow ± 5° rotation, green: ±15° rotation, light blue ±45° rotation, dark blue ±135° rotation. B Experiment 1, trial structure for Same/Different task, see text for detailed description. At Testing. participants view the object array from one of the viewpoints shown in A (right). In Different trials (50%) one of the objects in the array is moved between presentation and testing. Participants judge whether the array is the same or different to the one they encountered at Presentation. C Experiment 2, trial structure for Two Alternative Forced Choice task. At Testing participants view the object array from one of the viewpoints shown in A (right). On every trial one of the objects is moved between presentation and testing. The moved object is highlighted along with one other object selected at random. Participants indicate which of the highlighted objects has been moved.

3.3.3.1 Experiment 1: Same/Different

Experiment 1 (see Figure 3.1B) used a same/different spatial change detection task with an adaptive staircase controlling spatial change. At Presentation participants viewed an array of 4 objects (simple 3D geometric shapes with distinct geometry and contrasting colours selected from a pool of 10) arranged in a random spatial configuration for 6 seconds ('Sample Scene'). The screen dimmed over 0.5 seconds, remained black for 1 second before fading back over 0.5 seconds. At Testing, the participant viewed a new scene ('Test Scene') with their viewpoint shifted by 0, 5, 15, 45, or 135 degrees around the centre of the table for 6 seconds. Both the degree of view rotation and its direction (clockwise or anticlockwise) was selected randomly for each trial. In 'Same' trials – the object locations did not change between presentation and testing; in 'Different' trials, one randomly selected object was moved. The magnitude of this change was controlled using an adaptive staircase procedure described in more detail below. The participant then viewed the Test Scene for 6 seconds before being prompted to respond 'Same' or 'Different' by pressing the 's' or 'd' key on the keyboard.

3.3.3.2 Experiment 2: 2AFC

As in Experiment 1, at Presentation on each trial participants in Experiment 2 (Figure 3.1C) viewed a Sample Scene, table-top with 4 objects arranged in a random spatial configuration for 6 seconds. The screen dimmed over 0.5 seconds, remained black for 1 second before fading back over 0.5 seconds. At Testing one object was selected at random and moved, by a distance specified by an adaptive staircase forming a new spatially distorted configuration. Participants viewed this Test Scene with their viewpoint shifted by 0, 5, 15, 45, or 135 degrees around the centre of the table for 6 seconds. Two objects (including the target object and a foil, selected at random) were then highlighted and the participant was then allowed up to 8 seconds to choose which object had moved by clicking with the mouse.

3.3.4 Adaptive staircase procedure

In order to focus the majority of trials around participants' precision thresholds (corresponding to 75% correct responses), in each experiment, an adaptive staircase procedure was used to determine the amount of spatial distortion (i.e., distance moved by selected object in spatially distorted arrays) based on participant performance in each view rotation condition. Each level of the staircase referenced a specific distance in 15 logarithmically-spaced steps between 0.05m to 0.9m, with the largest distortion being limited by the edges of the display (see Appendix D for full details). The staircase level, specifying the current spatial distortion, changed following a 3 down 1 up procedure (the level of spatial distortion was reduced after 3 consecutive correct responses and increased after a single incorrect response). A reversal was recorded whenever responses led to a change in the direction of progression from a sequence of successively smaller distortions to a larger distortion or vice versa. Staircases completed when they reached 14 reversals (Experiment 1) or 7 reversals (Experiment 2). Independent staircases controlled each view rotation condition, such that the degree of spatial distortion on each trial was determined by the staircase level based on previous trials at the same view rotation condition.

In order to expedite the threshold estimation process, the number of steps moved after each response was initially large and gradually reduced over successive reversals: 6 steps until the first reversal, 4 steps until the second reversal, 2 steps until the third reversal, and single steps thereafter.

To avoid biasing participants' Same/Different criterion, in Experiment 1 a single staircase was used for each viewpoint condition starting at the largest level of spatial distortion; pilot work suggested that when beginning the task, trials involving very small spatial changes would be easily misinterpreted as a 'Same' trials. The experiment finished when all five staircases were completed or if the session duration exceeded 2 hours (two participants exceeded the time limit in Experiment 1).

In Experiment 2, a pair of staircases were used for each viewpoint condition. One staircase began with the smallest spatial distortion level and one began at the largest spatial distortion level. The experiment finished when all ten staircases were completed or if the session duration exceeded 2 hours (one participant exceeded the time limit in Experiment 2).

On each trial, matched spatial arrays were generated randomly. Initially, four of ten possible objects with distinct geometry and contrasting colours were selected randomly without replacement. Objects were placed at pseudo-random locations by first specifying evenly spaced locations on a circle of radius 1m around the centre of the table, and then applying a random global rotation of the array and individual radial (between 0.01m and 0.5m) and angular displacements (between 0° and 45°) to each object. A matched distorted array was then created by moving a single object, selected at random, by the distance determined by the staircase while avoiding distortions that resulted in the overlap of two objects or objects appearing beyond the table-top. Finally, the matched pair of arrays was shuffled so that either array could appear at Presentation, with its partner being used at Testing.

3.3.5 Data Analysis

To estimate each participant's spatial precision threshold (see A) we counted the number of correct responses made at each staircase level. We then fit a logistic function (Equation 3.1) to the data from each viewpoint condition using a maximum likelihood approach with the Nelder-Mead simplex algorithm (as implemented in 'fminsearch'; MATLAB ver. R2019a). The logistic function is defined as:

Equation 3.1

$$f(x) = 0.5 + 0.5 * (\frac{1}{1+e^{-k(x-x_0)}}) \tag{1}$$

Briefly, the logistic function (*f*) allows us to calculate the probability that a particular proportion of responses will be correct at each staircase level (x) given parameters governing its threshold (x_0) and slope (k). The fitting procedure systematically varies these parameters, calculating at each iteration the overall likelihood of the data from a given participant and condition (expressed as log-likelihood) and stopping when the likelihood is maximized. These fitted parameters, which describe the threshold and slope of the psychometric curve under which the observed data is most likely, were then recorded for further analysis.

For clarity below we have converted the k parameter to equivalent Weibull beta values, so that increasing values indicate steeper psychometric functions (which we refer to as "slope"). We also recorded the log-likelihood of the data given the fitted curve and calculated the log-likelihood of the data under a baseline model in which the observed proportion of responses was constant across all levels of spatial distortion. These likelihoods were used to calculate a Bayes Factor indicating the relative likelihood of the data under the two different models, which could be used as a measure of fit across participants/conditions where different numbers of observations were available (Wagenmakers, 2007). We report Bayes Factors (BF_{10}) where applicable below and in full in the appendices (Appendix B). Change detection thresholds were converted from meters to decibels (see Appendix A for more information).

To allow for a clear comparison between experiments 1 and 2 our main analyses (reported below) focus on the parameters derived from the psychometric curves and the underlying proportion of correct responses. For experiment 1 we additionally carried out a d' analysis (Appendix C) to ensure our main results were not affected by response biases.

3.3.6 Procedure

Before beginning the experiment, participants undertook a 10-minute practice session. Performance on the 0° condition at the easiest staircase level, was recorded and analysed, and exclusion criteria applied. They then completed the experimental protocol (i.e., completed all staircase reversals) in two sessions, each approximately 1.5 hours including breaks (around 3 hours total for each data collection). Data from each of these sessions was combined for analysis (an average of 514.8 trials per participant).

3.3.7 Participant exclusion criteria

We preregistered our intention to exclude participants who performed below 66% percent correct in 0° trials during the practice session from further participation. We determined, based on pilot work, that this level of performance would be indicative of exceptionally poor spatial memory, failure to understand task instructions, or a lack of motivation or cooperation. However, all recruited participants passed this exclusion criteria.

3.3.8 Data Exclusion Criteria

If a precision threshold was estimated far outside the testing range (above 10dB), this indicated that we were unable to effectively model a participant's performance in that condition and all data associated with this participant was excluded. Three participants failed to meet these criteria and we were unable to effectively estimate their threshold in the 135° viewpoint condition. Their data were excluded from the analysis. Note that this limit is a slight deviation from our preregistered plan in which we had specified a more conservative cut off, corresponding to the maximum distortion distance in our testing range (0.9m, 5.12 dB). This change was based on unexpectedly poor performance in the 135° condition (which would have resulted in excessive data exclusion), and because we found acceptable fits of the psychometric function were nonetheless obtained in all participants.

3.4 Results

Participants' responses in both experiments were well-described by the logistic function (see Appendix B for individual participant Bayes factors). In each case, the accuracy of responses increased with the degree of spatial distortion applied to the studied array forming a sigmoid curve as illustrated, for a representative participant, in Figure 3.2 A. In all cases (i.e., in each participant and viewpoint condition) the fitted curve described the data better than the model described by the null hypothesis, with the large majority of Bayes Factors classified as 'Strong' or 'Very strong' (Kass & Raftery, 1995; $BF_{10} > 20$ for 69/100 comparisons).

Using the fitted sigmoid we obtained each participant's spatial change detection thresholds, corresponding to the degree of spatial distortion required to obtain 75% accuracy at each level of viewpoint change. The average parameters for the best fitting sigmoid curves are summarized in Table 3.1.

Table 3.1 Average psychophysical sigmoid parameters in each condition in Experiment 1 and 2. Thresholds are measured in decibels standard deviations are following in brackets. Slopes are expressed in equivalent Weibull beta.

	Experiment : Same/Different		Experiment 2: 2AFC	
Viewpoint	Mean Threshold dB (SD)	Mean Slope (SD)	Mean Threshold dB (SD)	Mean Slope (SD)
0°	-8.81 (2.73)	2.10 (1.59)	-12.11 (2.99)	1.81 (1.19)
5°	-5.49 (2.53)	2.50 (1.54)	-9.94 (1.63)	1.99 (0.70)
15°	-3.73 (2.05)	1.76 (1.31)	-5.99 (2.96)	1.57 (0.72)
45°	-1.02 (2.84)	1.55 (2.09)	1.14 (3.1)	0.98 (0.64)
135°	3.01 (2.03)	1.66 (0.97)	4.91 (1.68)	1.81 (1.29)

3.4.1 Spatial Precision Thresholds and View Rotation

3.4.1.1 Differences between conditions

We predicted that, across participants, thresholds would a) depend on the degree of view rotation and b) would be higher for larger view rotations (greater distortion required to correctly detect moved object). In line with these predictions, in both experiments, the average threshold showed a monotonic increase in precision thresholds (poorer performance) as viewpoint increased as seen in Figure 2.2. This pattern was also seen when analysing the results from Experiment 1 as d' statistics (see Appendix C) indicating no evidence of response bias. Separate repeated measures ANOVA (carried out in R-Studio with the 'ez' package; Lawrence, 2016) for each experiment showed a significant effect of viewpoint on thresholds in both experiments (Experiment 1 [$F(4,36) = 44.2$, $p < 0.001$], Experiment 2 [$F(4,36) = 102.3$, $p < 0.001$]).

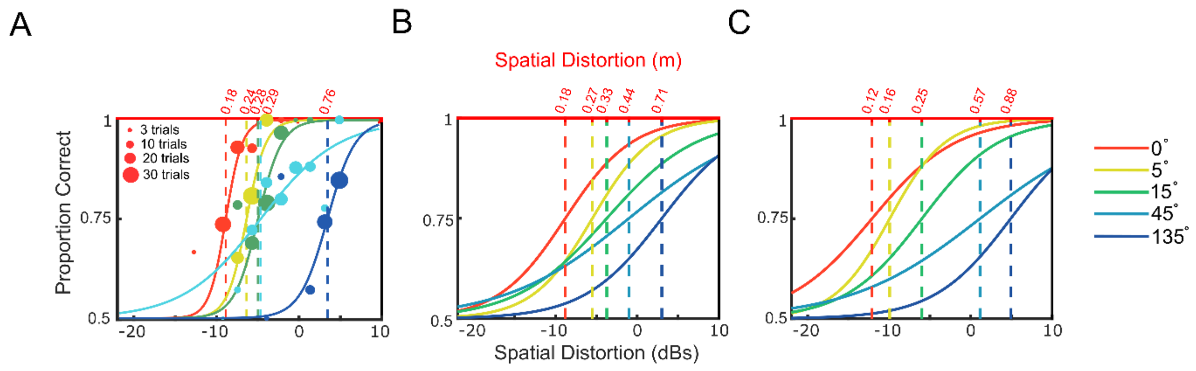

Figure 3.2 Psychophysical sigmoid curves for each view condition. Different viewpoint conditions are shown as with different coloured curves as shown in the legend. A. Representative data from a single participant (Experiment 1). Circle symbols indicate proportion correct at a specific staircase level corresponding to spatial distortion applied (in dB relative to 0.5m, black x-axis, bottom), with the size of each circle indicates the number of trials recorded at that level. Dotted lines show the estimated threshold level at 75% performance (with the corresponding spatial distortion in metres, red x-axis, top). B, Experiment 1, Same/Different task, group average sigmoid curves (created by averaging fitted threshold and slope values for each condition (see also Table 2.1). C. 2AFC task, group average sigmoid curves.

3.4.1.2 Comparing monotonic ordering to chance

Since we had anticipated that these effects would be monotonic but likely non-linear in nature, we carried out an additional pre-planned permutation-based analysis to determine, for each experiment, whether there was a monotonic increase in threshold across increasing viewpoint shifts. Specifically, to determine whether the average Spearman's correlation (over participants) was statistically significantly greater than would be expected under the null hypothesis (no monotonic relationship), we shuffled the threshold data from each participant (i.e., randomizing viewpoint labels) 10,000 times, calculating the cross-participant mean Spearman's correlation for each permutation, using the 95th percentile (single-tailed) of this distribution as the critical value in a comparison with the experimental results. We found that the observed non-parametric correlations were statistically significantly above chance in both experiments (Experiment 1: mean ρ observer = .94, mean ρ critical = .26, $p<0.0001$; Experiment 2: mean ρ observed = .95, mean ρ critical = .26, $p<0.0001$) indicating a consistent monotonic increase in thresholds with increases in viewpoint change.

3.4.1.3 Threshold differences between experiments

Using independent 2-tailed t-tests, we compared average threshold levels at each viewpoint between experiments (Table 3.2). We can see that there are significant differences between the thresholds measured in each method at the 0°, 5° and 135° conditions and no significant differences at 15° and 45°. Participants in the 2AFC task significantly outperformed those in the S/D task in the low viewpoint change conditions but performed significantly worse in the high viewpoint condition.

Table 3.2 2-tailed t-tests comparing thresholds obtained in Same/Different and 2AFC tasks for each viewpoint condition.

	Fitted Threshold (dB) Experiment 1 Same/Different	Experiment 2 2AFC			
Viewpoint	Mean (SD)	Mean (SD)	df	t	p
0°	-8.81 (2.37)	-12.11 (2.99)	18	-2.73	0.014*
5°	-5.49 (2.53)	-9.94 (1.63)	18	-4.68	0.000**
15°	-3.73 (2.05)	-5.99 (2.96)	18	-1.98	0.063
45°	-1.02 (2.84)	1.14 (3.1)	18	1.62	0.122
135°	3.01 (2.03)	4.91 (1.68)	18	2.27	0.036*

3.4.1.4 Effect of viewpoint on slope of psychometric function

Although we had no strong prior hypothesis concerning the effect of viewpoint on the slope of the psychometric function, we recorded and analysed these data (Figure 3.3) which give an indication of how the consistency of responses changes around the threshold and might be interpreted in terms of participants' certainty or uncertainty about judgements under different degrees of spatial distortion. Comparing these values using an exploratory repeated measures one-way ANOVA we found no statistically significant differences in Experiment 1 ($F(4,36) = 1.91$ ($p = 0.13$)) but we did find a significant difference in the average slopes of Experiment 2 ($F(4,36) = 2.96$ ($p = 0.03$)). However, the non-monotonic numerical pattern is strikingly similar in showing the shallowest slopes for the 45° condition in both experiments (a pattern which is evident in Figure 3.2).

A

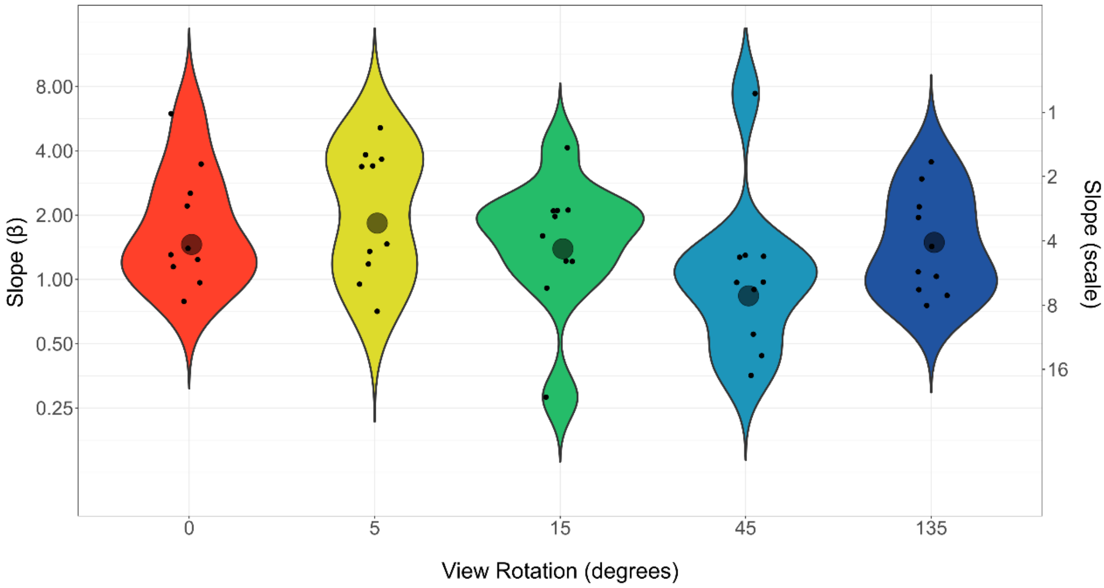

B

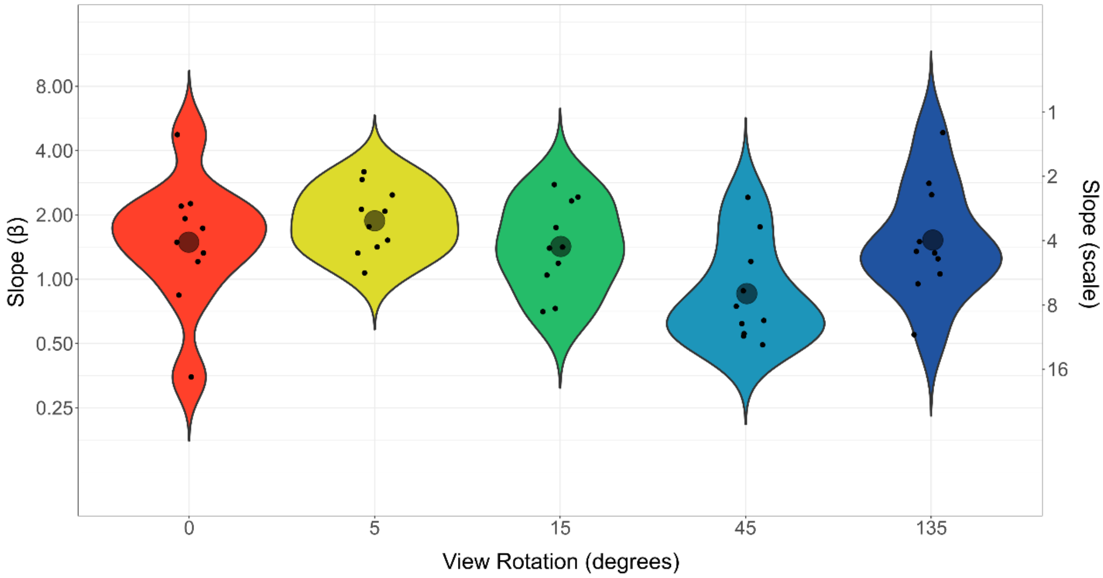

Figure 3.3 Distribution of estimated psychometric slope parameter as Weibull beta (left axis) and logistic scale (right axis). High beta values indicate steeper sloping sigmoid curves. Small points indicate individual participant estimates, larger grey circle indicates group geometric mean estimate for each condition. A Experiment 1. B Experiment 2. Violin plots created in R-Studio with the GGplot 2 package (Wickham, 2016)

3.4.2 Modelling spatial precision as a function of viewpoint shift

Across view conditions in both experiments, each participant reliably produced a similar pattern of results (see Figure 3.4). Spatial precision thresholds were lowest (i.e., performance was best) in the 0° viewpoint shift condition and increased monotonically as the angle between viewpoints at presentation and testing increased. Thresholds seemed to plateau beyond the 45° viewpoint shift. To quantify the relationship between view-dependent and view-independent contributions to the task, we modelled the relationship between spatial precision thresholds and viewpoint change conditions with a single function which we had preregistered (see https://osf.io/p6gzt/ and https://osf.io/j7a8p/). The model is defined as:

Equation 3.2 , viewpoint change model describing spatial memory precision as a function of viewpoint change using 3 parameters relating to aspects of the viewpoint rotation process.

$$f(\theta) = a\left(1 - \frac{\theta}{\pi}\right)^{b} + c \qquad (2)$$

and describes a family of curves relating spatial change detection thresholds to viewpoint change (see Figure 3.4). Its three parameters distinguish viewpoint specific, viewpoint-dependent and viewpoint-independent features; *a* governs the intercept with the y axis which represents the spatial precision in the 0° rotation condition (i.e., viewpoint-specific performance).; *b* determines the rate of change as viewpoint at test is rotated away from viewpoint at encoding for smaller viewpoint changes (i.e., viewpoint-dependent performance); *c* is the asymptote of the curve as the viewpoint change tends to 180°, (i.e., the viewpoint-independent limit of precision under large and unpredictable changes in perspective). Note *that* parameter *a* is always negative, reflecting the *decrease* in threshold seen at smaller viewpoint changes.

Although the model is intended primarily as a descriptive one, the arrangement of Equation 3.2 is intend to highlight the way the model characterizes view-threshold function in terms of the sum of two theoretically-motivated components, one, $a\left(1 - \frac{\theta}{\pi}\right)^{b}$, arising from view-sensitive representation and the other, *c,* arising from a view-invariant representation.

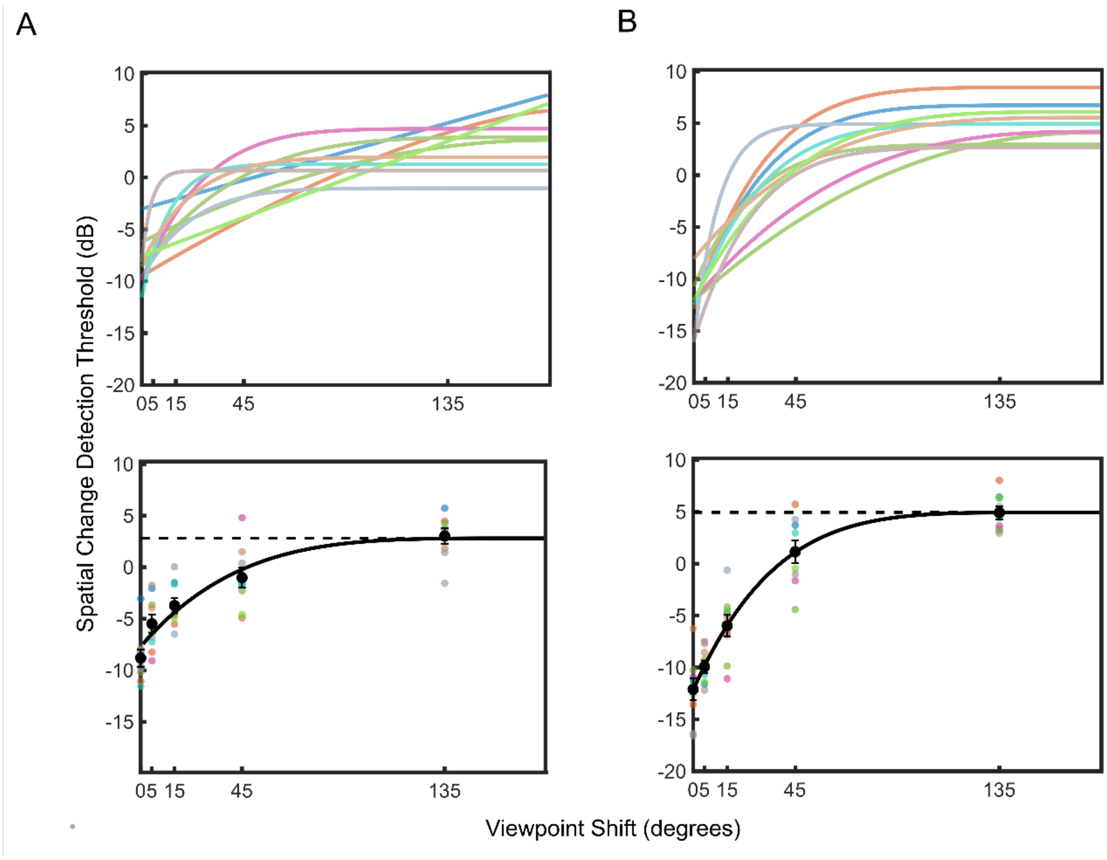

Figure 3.4 Relationship between viewpoint shift and threshold in individual participants, data and fitted model (Equation 3.2) . A Experiment 1 (Same/Different task) B Experiment 2 (2AFC task). Upper graphs show the fitted model for each individual participant. Lower graphs show group average model fit (i.e., taking the mean parameter values over participants). Coloured points on the graph represent individual participant thresholds at each viewpoint, with the colour scheme corresponding to the curves in the upper panel. The dotted line represents the asymptotic plateau, the level of ability in the largest view rotation condition. This is described by parameter c in the fitted model.

Figure 3.4 shows the fit of the model to each individual participant's thresholds (using 'lsqcurvefit' as implemented in Matlab ver. R2019a). Briefly, this involves varying parameters *a*, *b* and *c* systematically to find the curve that minimizes the sum (over viewpoints) of the squared distances between the observed data and the model.

In both experiments the model provides a reasonably good fit to the threshold data, although the fits in Experiment 1 are numerically poorer (Experiment 1: mean RMS error = 1.14 dB , Experiment 2: mean RMS error = 0.97 dB) and qualitatively, for three participants the fitted model suggests a nearly linear change, with the value of the *b* parameter approaching 1, deviating from the anticipated curve with asymptote.

In direct comparison (Figure 3.5) we see statistically significant differences in the *a* parameter value between experiments (*a*; $t(18) = -4.76$ $p < 001$). However, there were no significant differences between experiments for parameters *b* or *c* (*b*; $t(10.71) = -0.715$ $p > 0.05$, *c*; $t(18) = 0.23$ $p > 0.05$).

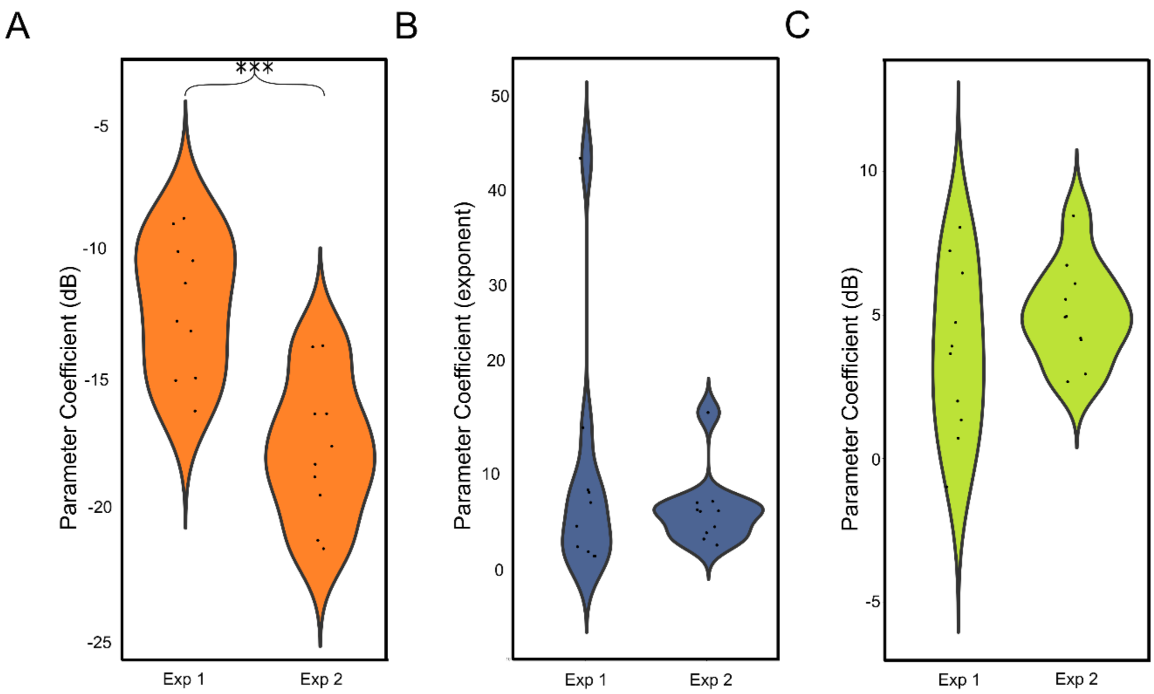

*Figure 3.5 Distribution of fitted model parameters . Small black dots represent individual participant coefficients. From left to right graphs show parameters a, b and c. Parameter a determines the intercept – the threshold spatial ability when the position at encoding is the same as the position at testing. Parameter b, describes the rate of change between precision in low viewpoint rotation conditions to precision in high viewpoint rotation conditions. Parameter c defines the asymptotic threshold representing the viewpoint-independent limit of spatial precision. *** indicates a significant difference in the a parameter for experiments 1 and 2, p<0.001.*

3.5 Discussion

Our aim in the current study was to elucidate the relationship between changes in viewpoint and the precision of spatial memory. In order to do this, we first needed to develop a new experimental paradigm with which we could measure spatial change detection thresholds for an array of objects in a virtual 3D space while manipulating the degree of viewpoint change between presentation and testing. We found that a psychometric function was an appropriate model for estimating performance in both a Same/Different (Experiment 1) and 2AFC task (Experiment 2). Using psychometric functions fit to participant data, we were able to estimate spatial change detection thresholds in all conditions, showing that as viewpoint shift increased from 0° to 135°, thresholds increased monotonically but not linearly. In each task, participants required increasingly large changes to the location of a single object within an array before they could reliably detect the distortion. This pattern of results, present in all participants, and the curvilinear relationship between viewpoint shift and spatial change detection threshold was well described by a preregistered model that separated and parameterised viewpoint-specific, viewpoint-dependent and viewpoint-independent components.

In this account, the high precision of responses in conditions where the viewpoint is fixed between presentation and testing (0° rotation, described by parameter *a* in our model) results from the unimpeded operation of viewpoint-specific egocentric representations (such as view matching, see e.g., Waller & Hodgson, 2006). When the viewpoint remained unchanged participants could detect the movement of a single object within a four-item array as small as 17.4cm on average. As the viewpoint at testing rotated away from viewpoint at presentation, thresholds increase towards a plateau for viewpoint changes beyond 45° (rising to 78.7cm on average). The viewpoint-dependent decay in precision across increasing shifts in perspective (governed by parameter *b* in our model), is perhaps consistent with the operation of a mental rotation process for small viewpoint shifts which introduces an accumulating error. However, the plateau itself (described by parameter *c* in the model) does not appear to result from a process that transforms and manipulates purely egocentric representations, since in that case we might expect to see error continue to accumulate linearly for larger viewpoint shifts. Instead, we interpret it as the limit of viewpoint-independent performance supported by allocentric representations, when little to no valid contribution from egocentric spatial representations is available.

This task and analysis adds to the evidence that ego- and allocentric systems make distinct contributions to spatial working memory (e.g., Aagten-Murphy & Bays, 2019). Together they offer a means to distinguish and quantify these contributions, which may have applications in identifying and monitoring disorders that differentially affect the distinct brain systems that underpin memory for object locations, as well as in understanding individual differences in healthy people.

3.5.1 Comparison between tasks

Had there been no difference in the precision of memory representations in the two tasks then under assumptions in standard signal detection theory, we might still have expected to see a constant difference in the thresholds observed at each viewpoint. This is because the 2AFC task involves the comparison of two locations (each subject to noise/imprecision which is combined when generating responses) whereas the S/D task involves only a single location (for example see Makovski et al., 2010). However, our results show a different pattern: thresholds show a greater dynamic range as viewpoint is varied in the 2AFC task than in the S/D task. This we believe could be due to a combination of measurement issues (the S/D provides a noisier estimate of thresholds) and ecological factors (the S/D task is more effective in eliciting allocentric encoding strategies).

On the first point, we note that the 2AFC task is better suited to the use of a psychophysical staircase design, since it is possible to make a spatial change on every trial (each yielding valid information about the psychophysical function), whereas the interpretation of responses is ambiguous for Same trials in the S/D task.

On the second point, the S/D task has potentially greater ecological validity when estimating allocentric memory, in that participants are instructed to judge whether the entire array has changed (i.e., an explicitly allocentric judgement), whereas the 2AFC task encourages participants to pay attention to the individual objects (potentially at the expense of their allocentric relationships).

The comparison of Experiments 1 and 2 is consistent with a role for both factors in determining the threshold-viewpoint function. First, we were better able to fit performance in the 2AFC task (Experiment 2) with our model. This might reflect the benefit of using a less noisy method to elicit unambiguous psychophysical judgements. Second, the S/D task showed significantly higher spatial precision (i.e., lower change detection thresholds) in the most viewpoint-independent condition (135°), and significantly lower spatial precision in the two most viewpoint-dependent conditions (0° and 5°). Together with a significantly higher a parameter in our model (indicating somewhat reduced sensitivity to spatial changes in the 0° condition), these findings that are consistent with the idea that the S/D task encouraged greater use of allocentric encoding strategies. Overall, while the S/D task may afford strategies that allow for somewhat better performance under view-shift conditions, the 2AFC task appears to be better suited to rapid and reliable estimation of spatial precision and the greater dynamic range of the resulting threshold-viewpoint function may make it more suitable for distinguishing ego- and allocentric contributions in the presence of noise, although thresholds may not indicate maximum possible sensitivity to allocentric change. In similar vein, we should note that both tasks may favour allocentric strategies relative to everyday settings, in that the large majority of trials involve an unexpected change of viewpoint, which the participant can anticipate even if the size and direction of the change is unpredictable.

3.5.2 Additional findings

One interesting and unanticipated quirk of our results emerged from the analysis of the slope of psychometric functions across different viewpoint shifts. The non-monotonic pattern of slopes is striking and in both experiments the shallowest slopes are seen in the 45 degree viewpoint condition (rather than the more extreme 0 degree or 135 degree conditions, see Figure 3.2 and Figure 3.3). We speculate that this might reflect increased uncertainty resulting from conflict between ego- and allocentric systems, perhaps resolved through some form of Bayesian cue combination (see e.g., Deneve & Pouget, 2004; Xu et al., 2017; Zhao & Warren, 2015); in this account, at much smaller viewpoint shifts responses would be dominated by high-precision egocentric information, and at much larger viewpoint shifts allocentric responses would be dominated by low-precision allocentric information, but at intermediate shifts, around 45° in the current tasks, the two systems would be in tension, providing conflicting influences on the response, and increasing uncertainty.

Another difference between view conditions was increasingly poorer fits to the psychophysical curve as viewpoint rotation increased. This noise indicates that our manipulation of spatial distortion did not account for as much variance in task performance. This may be due to the fact that with increasing viewpoint change, reorientation becomes increasingly reliant on cues related to the spatial configuration of the array (Sargent et al., 2008, 2010). Depending on the random configuration from one trial to the next, the configuration itself could impact on the success of accurate reorientation and recognition, independently of memory for the individual object locations.

3.5.3 Theoretical implications and conclusions

Overall, our data indicate that spatial change detection thresholds arise through a process combining a fixed, viewpoint-independent contribution and a view-dependent contribution that increases sensitivity to spatial changes for smaller shifts in viewpoint. In our descriptive model this combination is achieved by simply *adding* components corresponding to egocentric and allocentric contributions, but it is important to recognize that this additive description, while capturing the behavioural outcome, may not reflect the complexity of the underlying neural mechanisms. In particular, our model might suggest that ego- and allocentric representations contribute *independently* to behaviour, but at a mechanistic level this is unlikely because allocentric representations must ultimately be derived from egocentrically-encoded sensory information (Bicanski & Burgess, 2018; Byrne et al., 2007). This implies any task requiring viewpoint-independent representation can never be entirely process pure, and that that limitations (e.g., noise, biases, spatial resolution) of the egocentric code must be inherited by the allocentric code whereas the converse does not apply for viewpoint-specific tasks. Any such limits will thus impose an "*egocentric bottleneck*", at least affecting the initial encoding of allocentric information. It is conceivable that allocentric representations could be cleaned-up by integrating information from different cues (for example, in an allocentric representation, the information about the location of one object might serve to constrain the representation of another). A further consideration is that the process of translating between ego and allocentric representations at encoding and retrieval may itself impose limits on the precision of the underpinning representation. At a mechanistic level, the resolution of cue conflicts in behaviour (i.e., between responses driven by allo- and egocentric- representations) is likely to involve complex dynamic interactions between neurons representing different ego- and allocentric locations (Jeffery et al., 2016).

In short, a more complete understanding of ego- and allocentric contributions to spatial memory will depend on more detailed, mechanistic modelling of tasks like those we have developed. Our results provide some useful constraints on the behavioural predictions of such models. The methods provide a useful empirical technique for quantifying the precision of spatial memory for the locations of objects in virtual stimulus arrays and distinguishing ego- and allocentric contributions to precision which, our model indicates, respond lawfully to changes in 3D viewpoint.

Chapter 4. Exploring how stimuli of different scales are remembered

4.1 Abstract

The use of allocentric spatial memory is often associated with tasks that would naturally involve large scale environments such as finding shortcuts or recognizing environments from a novel perspective. Allocentric neural representations are characteristically anchored to the environment in fixed and absolute terms. Conversely, egocentric spatial memory is associated with tasks that unfold over short distances and timescales and egocentric representations are by definition dependent on the observer's viewpoint and body position. In this chapter the quality of spatial memory representations will be further investigated by asking whether there are differences in how egocentric and allocentric representations encode small- and large-scale arrays and environments and whether the physical size of a stimulus affects the precision of spatial memories. In two preregistered spatial change-detection experiments I (1) manipulate the size of a spatial array and (2) use immersive VR to manipulate the perception of environmental scale while controlling pictorial cues. In both viewpoint-dependent and -independent tasks, spatial change-detection thresholds increase as scale is increased and that the increase scales with the environment. I interpret this as evidence of the constraints initial egocentric perception and memory has on a subsequent allocentric encoding.

4.2 General introduction

The previous chapter established that viewpoint-independent memory has a poorer spatial precision than viewpoint-dependent memory. It was possible to model spatial memory precision as a function of viewpoint rotation and to estimate participant thresholds. These thresholds correspond to a physical distance in the environment. If an object moves more than this distance, participants detect the change on more than 75% of trials. However, it is unclear from the findings of the previous chapter whether these thresholds are highly specific to the task and stimuli we used, and to what extent viewpoint-dependent and -independent thresholds would scale to larger and smaller array configurations or environments. That is, are change-detection thresholds fixed, absolute values, or are they in some way proportional to the scale of the stimuli or to perception? Furthermore, does this vary between viewpoint-dependent and -independent tasks? Our current understanding of the nature of egocentric and allocentric representations suggests that changes in spatial scale may impact viewpoint-dependent and viewpoint-independent representations very differently.

Allocentric neural representations in the hippocampal formation are anchored to the environment and, at least in some circumstances (as reviewed in section 4.2.2), are associated with absolute physical locations and scales which are independent of the observer. So, it might be anticipated that the precision of viewpoint-independent memory, to the extent it relies on such allocentric representations, is unaffected by manipulations affecting the scale of stimuli or environment. On the other hand, egocentric representations are more directly associated with perception and in these systems the precision of responses is typically governed by Weber's law; a rule seen in a wide variety of perceptual judgements where change detection thresholds scale with the magnitude of the stimuli. In addition, as reviewed below, evidence suggests that allo- and egocentric codes and strategies may be systematically favoured at different spatial scales. With these differences in mind this chapter sets out to investigate whether and how viewpoint-dependent and -independent change detection thresholds are differentially affected by manipulations of spatial scale.

4.2.1 Behavioural evidence for egocentric and allocentric dissociation at different scales.

There is a behavioural distinction in how ego- and allocentric representations are used. Recent human cell studies (Kunz et al., 2021) and neuroimaging showing parietal cortex activation provides evidence for the employment of egocentric representations in tasks such as following a familiar route and small scale, short term movement tasks (Colby, 1998; Committeri et al., 2004; Galati et al., 2010; Schindler & Bartels, 2013). With high accuracy and automatic spatial updating, egocentric representations seem particularly well suited to immediate action and small scale environments such as moving around a room or picking items up from a table (Wang & Spelke, 2000). This scale of environment, known as 'vista space' (Montello, 1993) due to the characteristic of all associated parts of the environment being within view from a single vista, has been (for practical reasons) the most common type of environment used in laboratory-based spatial memory testing for both egocentric and allocentric spatial memory.

However, beyond vista space, into so-called 'environmental space', egocentric spatial updating alone falters. A familiar example of this is people who get lost in deserts, without landmarks or orientations cues, tending to walk in circles without realising (amazingly the circles can be as tight as <20m diameter, Souman et al., 2009). This is due to accruing additive noise during automatic spatial updating via visual and idiothetic cues when the observer is not provided with constant reference points in the form of proximal or distal landmarks (Benhamou et al., 1990; Riecke et al., 2002; Wang & Spelke, 2000).

In these instances, when there is a gap in automatic spatial updating (such as when the observer approaches a familiar place from a different route or is transported to a new place) an environmentally-anchored, but potentially lower precision, allocentric representation is more useful. Thus, longer timescales and by larger environments are characteristic of tasks that rely on viewpoint-independent memory (as reviewed in Wolbers & Wiener, 2014). For instance, hippocampal activation is associated with 'wayfinding', 'short cut finding' and self-location from a novel point of view (Wolbers & Büchel, 2005). These are demanding tasks that unfold in environmental space and require viewpoint-independent memory.

Furthermore, in some smaller scale environments (room sized spatial arrays) it has been shown that allocentric cues are not used at all. For example, 5 year olds did not use landmark information in small rooms (4ft x 6ft) but did in large rooms (8ft x 12ft) (Learmonth et al., 2002). From these studies we can intuit that changing the scale of the stimulus or environment, even within vista space, may impact the extent to which an observer relies on allocentric or egocentric spatial information to complete the task.

It is likely that for many real world examples an observer would make use of any and all spatial information available and a number of experimental results (Aagten-Murphy & Bays, 2019; Fiehler et al., 2014), computational and theoretical models (Bicanski & Burgess, 2018; Ekstrom et al., 2014) and also findings in the previous chapter suggest that for most tasks there is a shared contribution of egocentric and allocentric spatial representations. We have previously reviewed the way that task demands such as the viewpoint shift between presentation and testing, object-object proximity and availability of idiothetic cues can influence the relative contributions from these separate representations (Burgess et al., 2004; Hartley et al., 2007; King et al., 2002; Morris, 1984). What is currently unclear is whether the scale of the stimulus also affects the balance of egocentric/allocentric information that is available to the observer for use in the task. It may be the case that for physically larger stimuli, allocentric information is more effective and may be selectively used over egocentric information i.e., larger scale stimuli may favour the use of allocentric spatial information.

In summary, a substantial body of evidence that suggests ego- and allocentric representations are suited to different spatial scales. Egocentric representations seem better suited to short timescales and smaller stimuli while allocentric representations seem better suited to longer timescales and tasks that can take the observer past the sensory horizons and environmental boundaries. If this is the case, we might expect to see distinct effects on spatial change-detection thresholds when spatial scale is varied together with viewpoint.

4.2.2 Neuroscientific evidence for egocentric and allocentric dissociation at different scales

Beyond the behavioural suitability of viewpoint-dependent and -independent memory to small and large scale, there are also neuroscientific considerations which suggest that the effect of changing the scale of spatial stimuli might impact egocentric and allocentric memory spatial precision differently.

Spatial cells within the hippocampal formation, particularly place and grid cells, are characterized by their encoding of space, independent from the location of the observer (see Chapter 1.2 for more detail). For example, in the case of grid cells, when an observer moves across an environment, the grid cell fires when the observer enters that cell's grid field, locations that are tessellated across the environment in an equilateral triangular pattern. The firing field of a grid cell thus has an absolute scale, corresponding to the fixed spacing of the grid fields. Although hippocampal place cells typically have single fields in a small environment, it can be argued that the extent of the field and its location within the environment also provides an absolute scale. Indeed, the intrinsic scale of place and grid cells is organized into modules (Brun et al., 2008; Kjelstrup et al., 2008) and arranged topographically along the long-axis of the hippocampus, suggesting that it has some significance for local computations (Jung et al., 1994; Lyttle et al., 2013).

Where it can be most easily measured (in grid cells) the scale of the representation remains stable over time and in different environments. When transferred to a new environment the phase of the grid cell, the distance between grid fields, remains the same (Hafting et al., 2005). While there are some notable instances of grids distorting when a familiar environment is altered and in some cases expanding in novel environments (Barry et al., 2012; Hägglund et al., 2019), it is characteristic that in unchanged environments, grid fields remain spaced in relation to a fixed, physical distance uninfluenced by the perspective of the viewer. Generally, the fixed, absolute nature of this code is integral to the allocentricity of spatial cells in the hippocampal formation. With this in mind, we might expect the precision of judgements based on this representation to be unaffected by stimulus or environmental scale.

However, this contrasts how we might expect egocentric spatial memory thresholds to change with scale. Weber's law describes the perceived change in stimulus intensity to be relative to the original magnitude of the stimulus and is evidenced in almost all parts of human perception. Weber's law also applies to the intensity of spatial change, known as 'Weber's law for position' (e.g., Levi et al., 1988). In egocentric/visual discrimination tasks, it has been shown that object separation negatively effects observers' ability to estimate positions and that the accuracy of distance estimation also decreases with object separation, obeying Weber's law (Levi et al., 1988). Weber's law applies in the context of spatial scale whether conveyed by the size of objects or object separation (Levi & Klein, 1992; Whitaker & Latham, 1997). In a physically larger array, in which the objects are further spread or are at a greater distance from the observer, it is more difficult to perceive spatial changes. We would therefore expect similar principles to apply to egocentric spatial change-detection tasks.

4.2.3 General summary

In summary, measuring the precision of egocentric and allocentric memory while varying scale has the potential to reveal more about the nature of these representations. For one, allo- and egocentric spatial memory seem, respectively, better suited to larger and smaller scale tasks and environments in terms of real-life use. Secondly, characteristics of the firing fields of allocentric spatial cells suggest that to the extent that viewpoint-independent behaviour depends on such neural representations, precision may be fixed and absolute, anchored to the environment, whereas viewpoint-dependent precision is likely to obey Weber's law and scale relative to the stimulus.

In this chapter, two experiments investigate the way spatial memory precision, as measured by change detection thresholds, changes as the scale of the stimuli is manipulated. Experiment 4.1 addresses the way viewpoint-dependent and -independent spatial change detection thresholds are affected by manipulations of the overall size of the array (its 'footprint') while the sizes of objects and environment are fixed. Experiment 4.2, using immersive VR, compares the way viewpoint-dependent and -independent spatial change detection thresholds are affected when the perceived scale of the entire stimulus (i.e., environment, array and objects) changes while 2D pictorial cues are preserved.

4.3 Experiment 4.1

4.3.1 Introduction

Perceptual experiments show that Weber's law of position applies to the vast majority of egocentric perceptual tasks. When the magnitude of distances between objects increases, spatial changes (individual object movements) are more difficult to detect (Levi & Klein, 1992; Whitaker & Latham, 1997). This is thought to be necessary in order to allow the human perceptual system to operate at an appropriate level of sensitivity across a near-infinite range of stimulus magnitudes. In contrast, the nature of allocentric encoding is viewpoint-independent and there is evidence that allocentric encoding is anchored to distances irrespective of viewpoint. This might suggest that allocentric precision is fixed and absolute and if so, the size of a stimulus array would not be expected to have any effect on allocentric spatial precision.

Experiment 4.1 will measure the threshold spatial precision in a spatial change-detection task in both viewpoint-dependent and viewpoint-independent contexts (using the degree of viewpoint rotation between presentation and test to manipulate viewpoint dependence as in previous chapters). To test the effect of scale on spatial representations of the array, the size of the array footprint will be manipulated while keeping the size of the objects and environment intact.

4.3.1.1 Experiment 4.1 Hypotheses

As in previous chapters we expect to see an effect of viewpoint change, such that spatial change detection thresholds increase when the viewpoint at test is different from the viewpoint seen during presentation. It is further predicted that a change in the spatial scale of the array will negatively affect precision of viewpoint-dependent memory in line with Weber's law. We will also investigate the effects of scale on the precision of viewpoint-independent spatial memory. If there is an interaction between viewpoint change and array scale, such that change detection thresholds in the viewpoint-independent (view rotation) condition do not change (or that any increase is reduced relative to viewpoint-dependent performance), this would be consistent with the suggestion that allocentric spatial representations are absolute in nature. It would follow that their greater involvement in the viewpoint-independent task makes performance (relatively) insensitive to changes in the scale of the array.

4.3.2 Methods

4.3.2.1 Overview

In Experiment 4.1, building on the methods established earlier in the thesis (Chapter 2, 3), participants completed a spatial change-detection task, making 2AFC judgements about the locations of objects in a virtual environment. Between presentation and testing, an adaptive staircase controlled the magnitude of movement of one array item. Both the scale of the array footprint (Small or Large Array, see 4.3.2.5) and the viewpoint angle shift (0° and 135°) were manipulated in a 2x2 factorial design.

This experiment was piloted and preregistered prior to any data collection (https://osf.io/v78kx/).

4.3.2.2 Participants

All participants were recruited via the University of York's Department of Psychology participant recruitment system, 12 subjects (9 Female) who each completed around 500 trials (minimum: 486, maximum: 546, mean: 506 trials). Participant recruitment size was based on a power analysis using the effect sizes from our previous experiment and estimates from the current Experiment's pilot study. A power calculation in Gpower indicated that a sample size of 12 allows for a detection of main effects (at a $p<0.05$ level of significance) with an effect size of 0.8, $F = 0.4$. This corresponds to a change in the spatial change-detection threshold of approximately 2dB between conditions (in pilot data we observed an effect of scale of approximately 4-5 dB).

4.3.2.3 Participant Exclusion Criteria

Participants were excluded if they performed below 66% correct on 30 'easy trials' (set at the highest change factor present in the staircase) during the tutorial phase. However, no participants were excluded based on this criterion.

4.3.2.4 Data Exclusion Criteria

If our threshold estimation technique provided a single participant with a threshold of above 15dB, this indicated that we were unable to accurately determine a psychophysical threshold. All recorded participant results were reported and included within the OSF data folder but excluded from subsequent modelling and analysis (https://osf.io/v78kx/). Data was also excluded if curve fitting failed to converge. In total 3 participants were excluded due to failed curve fitting and threshold estimations beyond the allowed range. These participants will be excluded from further analysis, but their data will be included in the associated data repository.

4.3.2.5 Materials

All testing took place in a quiet room in Department of Psychology, University of York. Task and stimuli were rendered and controlled by custom software programmed in the Unity3d game engine in C# running on an HP Z400 computer and displayed on a 20" 'Elite Display 232' monitor. As in previous studies, virtual environments and experimental procedure were implemented with the Unity3D game engine.

To accommodate a wider range of array sizes, the virtual environment was changed from the one used in Chapters 2 and 3 (see Figure 4.A). It consisted of a 10m diameter, circular space with a mountainous skybox, providing distal orientation cues at the horizon (no distance cues). Participants viewed objects in a random configuration in the centre of the array, placed onto a floor plane with a striped noise material. This material was generated procedurally before testing using a repeating Perlin noise pattern adding orientation cues without providing distinct landmarks with which participants could use as location cues. Spatial arrays were generated by placing objects in pseudo-random location explained fully in Chapter 2 (section 2.3).

In this experiment, we manipulated the scale of the array footprint. Note that as object displacements are distorted unpredictably (using random numbers subject to constraints) and include target objects whose location is affected by the staircase procedure, it is not possible to simply scale arrays of fixed and predetermined size. Instead, array scale manipulation was achieved by adjusting the procedure used to generate the arrays (see section 2.3 and Figure 2.2); changing the radius of the initial circle placement of the objects and scaling the radial displacements of the objects by the same factor. We tested two scale conditions, a ‘Small Array’, generated using an initial radius and displacement factor 0.75x that used in Chapter 3, and a ‘Large Array’, generated using an initial radius and displacement factor 1.5x that used in Chapter 3. In the Small Array objects were thus placed at an initial radius of 0.75m, each object location was then perturbed resulting in objects a maximum of 3.26m from the environment centre (mean final distance = 0.45m, SD = 0.11). The initial radius used for the Large Array was 1.5m, with individual object locations then being perturbed resulting in objects a maximum of 4.1m from the environment centre (mean final distance=0.83m, SD = 0.12). Note that final mean distances of objects from the array centre are not exactly doubled due to both to randomness in the array generation procedures and amount of distortion applied to the target object changing due to performance.

When the participant viewed the array from a rotated angle, the degree of rotation was either 0°, or 135°. This included the extremes tested in the previous experiment, the most viewpoint-dependent condition (0°) and the most viewpoint-independent condition tested (135°). In the previous chapter we saw that spatial change-detection thresholds declined and then plateaued as viewpoint rotation increased. A view change of 135° was past the point of plateau, indicating that it was a reliant on largely viewpoint-independent information to solve the task.

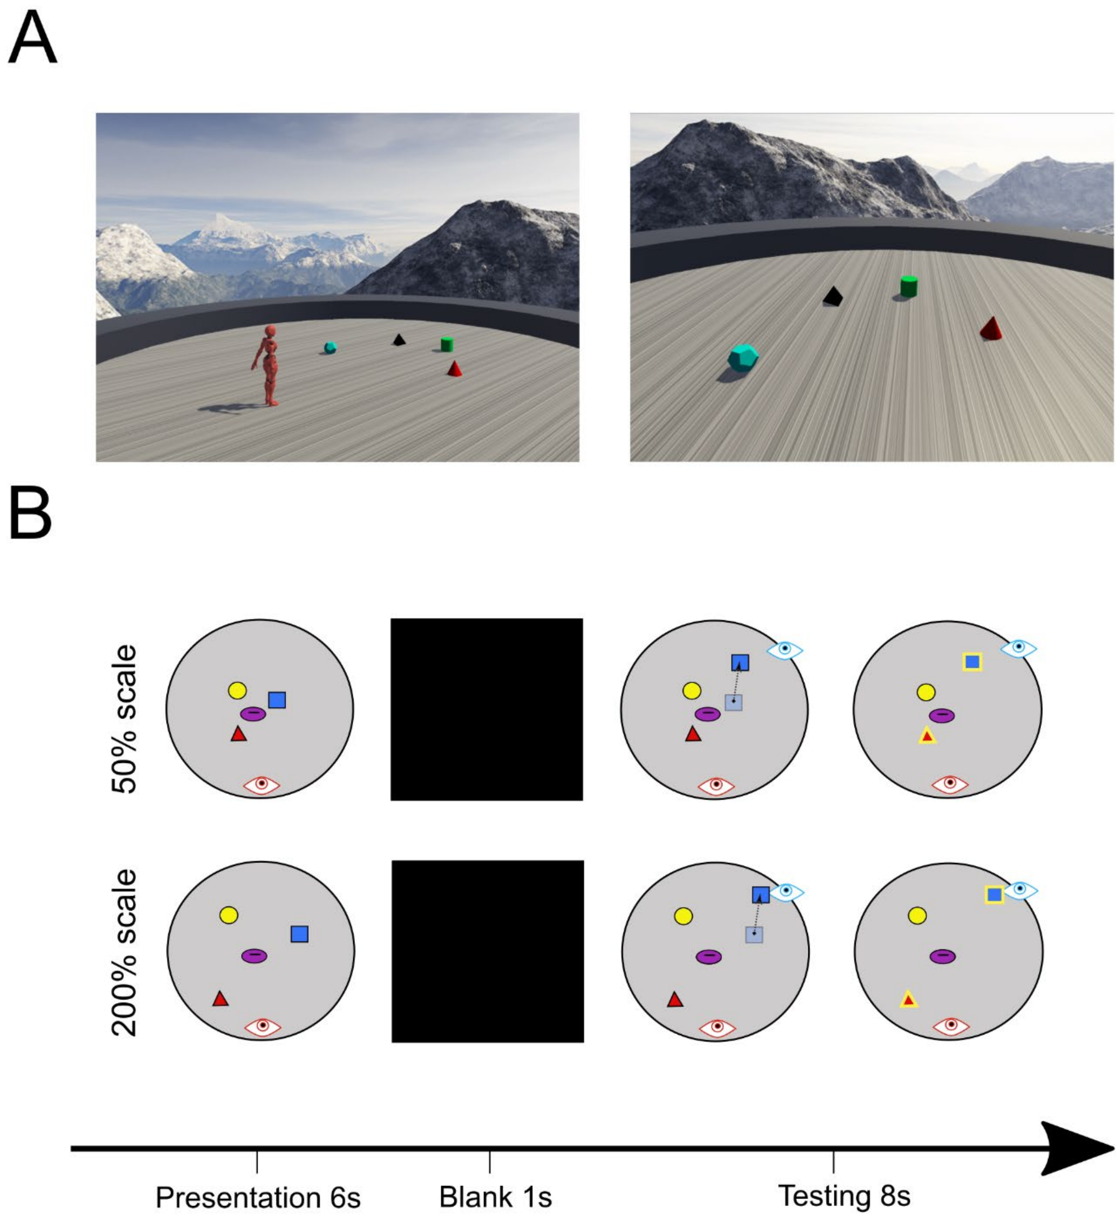

Figure 4.1 Experiment 4.1: Task and Materials . A shows an example of the virtual environment used. Left depicts the participants location relative to an example array of objects (a view never seen by the participant) and right shows an example array from the participants point of view at presentation. B shows the trial procedure used in Experiment 4.1. In the 'Presentation' phase, participants re shown an array of object in either the Small Array or Large Array condition. The 'scale' refers to the average array footprint. One object is moved while the screen is blank and participants are moved around the array by 0 or 135 ° (denoted by the red and blue eye symbol). In the 'Testing' phase, the target object and a random foil is highlighted and participants are prompted to select which object had moved. They have 8 seconds to respond by clicking the mouse pointer on the object that they thought had moved.

4.3.2.6 Procedure

Participants first undertook a tutorial phase during which they completed a short version of the task. In the tutorial phase, 0° trials were presented with increased likelihood and both array scales were used. The block was completed when participants finished thirty 0° trials. Performance in the tutorial trials was checked against the exclusion criteria ensuring that participants understood the task.

Participants then progressed to the main task. As in previous studies, on each trial, they viewed 4 objects arranged in a random orientation from a starting position (Presentation). They then viewed the array from a rotated position of either 0 (not moving) or 135° (rotated left or right around the array) and were prompted to select which of two objects had moved (Test). One change in this experiment from Chapter 3 is that participants were able to respond immediately when the changed stimulus was presented. This change was included for efficiency and to reduce the overall length of the experiment.

Each participant completed 4 separate conditions in a 2×2 factorial design (Viewpoint: 0 or 135° × array size: Small/initial radius 1m × 0.75 or Large/initial radius 1m × 1.5). A separate adaptive staircase modified the degree of spatial distortion applied to the array for each condition (4 total). Trials were presented in a random sequence, conditions for trials were randomly drawn from a list of incomplete staircases.

Each staircase began at a distortion level equivalent to the highest level to spatial distortion in our previous experiment (0.9m, Chapter 3) although for this experiment the testing range was extended, based on poor performance in pilot work, to include larger distortions. Each adaptive staircase followed a 3 down 1 up procedure (in which participants must answer 3 consecutively correct responses for spatial distortion to decrease and 1 incorrect response for spatial distortion to increase). Each staircase completed when it had reached 14 reversals (see Chapter 2.3.2). The experiment finished when all staircases completed or when the test session exceeded 1 hour.

4.3.2.7 Analysis

Spatial change detection thresholds were estimated with the same Maximum Likelihood curve fitting procedure implemented using MATLAB's fminsearch and as described in previous chapters (see Chapter 2.5).

To investigate effects of array scale and viewpoint on precision as well as interactions between scale and viewpoint, a 2x2 repeated measures ANOVA was carried out.

4.3.3 Results

4.3.3.1 Psychometric Curve fitting

Participant responses in all conditions (an average of 506 trials per participant) were well described by the psychometric function, as in previous experiments, the accuracy of responses increased with the degree of spatial distortion applied to the array, forming a sigmoid curve.

To quantify the goodness of fit of the psychometric curve, a Bayes factor was calculated for each sigmoid, comparing the fit of the psychophysical curve to a null hypothesis (i.e., the likelihood of the data under a uniform model where the proportion of correct responses, the mean across all trials, is unaffected by the degree of spatial distortion). In all cases the fitted curve described the data better than an alternative model suggested by the null hypothesis, with all Bayes Factors giving evidence for H1 (based on interpretations from Kass & Raftery, 1995, see Appendix E).

Four spatial change detection thresholds were estimated at 75% performance for each participant, in each of the test conditions (Small Array, 0° rotation; Small Array, 135° rotation; Large Array, 0° rotation; Large Array, 135° rotation).

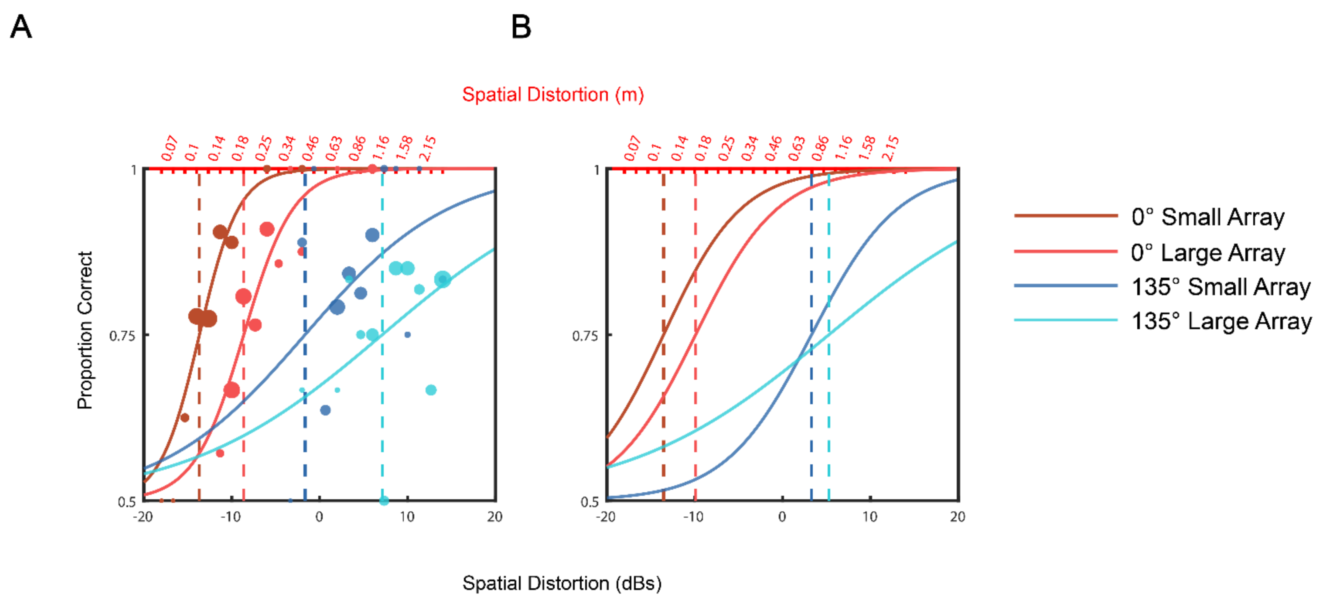

Figure 4.2 Psychophysical sigmoid curves for each experimental condition . Different conditions (viewpoint x array scale) are depicted with different colours. Red denotes 0° viewpoint shift and blue denotes 135° viewpoint shift. A: Representative data from a single participant. Circle symbols indicate proportion correct at specific staircase levels (in dB relative to 0.5m) and the size of each marker representing the number of trials undertaken at that level. Dotted lines show the estimated threshold level at 75% performance. B: group mean sigmoid curves (created by averaging fitted threshold and slope values for each condition).

The average parameters for the best fitting logistic curves are summarized in Table 4.1.

Table 4.1 Experiment 4.1, means of best fitting slopes for each condition and participant . Each value is the mean of separately fit psychophysical functions for each condition and each participant.

Condition (viewpoint/scale)	Mean Threshold dB (SD)	Mean Slope (SD)
0° Small Array	-13.54 (3.23)	4.43 (3.01)
135° Small Array	3.29 (3.29)	4.93 (2.96)
0° Large Array	-9.89 (2.69)	4.67 (3.79)
135° Large Array	5.27 (5.37)	11.49 (7.28)

4.3.3.2 Effects of viewpoint and array scale on spatial change-detection thresholds

To assess the effects of manipulations of viewpoint and scale, a 2x2 factorial ANOVA (Viewpoint x Scale) was carried out. This showed a significant main effect of viewpoint with significantly higher spatial change-detection threshold at higher viewpoint rotations, ($F(1,10) = 207.56$ ($p < 0.01$) indicating reduced accuracy and replicating results seen in previous experiments in this thesis. There was also a main effect of array scale ($F(1, 10) = 9.08$ ($p = 0.01$) showing significantly higher spatial change-detection thresholds in the Large Array. There was no significant interaction between viewpoint and scale $F(1, 10) = 0.53$ ($p = 0.38$)).

Numerically, in each viewpoint condition, the change in thresholds between array scales was somewhat less than the two-fold scaling used to generate the array (i.e., 1.5:0.75). For the 0° view conditions: a 3.65dB difference in thresholds corresponds to a factor of 1.6 and in the 135° view conditions: a 1.98dB mean difference corresponds to a factor of 1.4. Interestingly the ratio between the average threshold values (i.e., 3.65:1.98 dB) is closer to the ratio of actual observed mean target distances from the centre of the array (0.83m : 0.45m) than the initial scaling factor for the overall array (1.5:0.75).

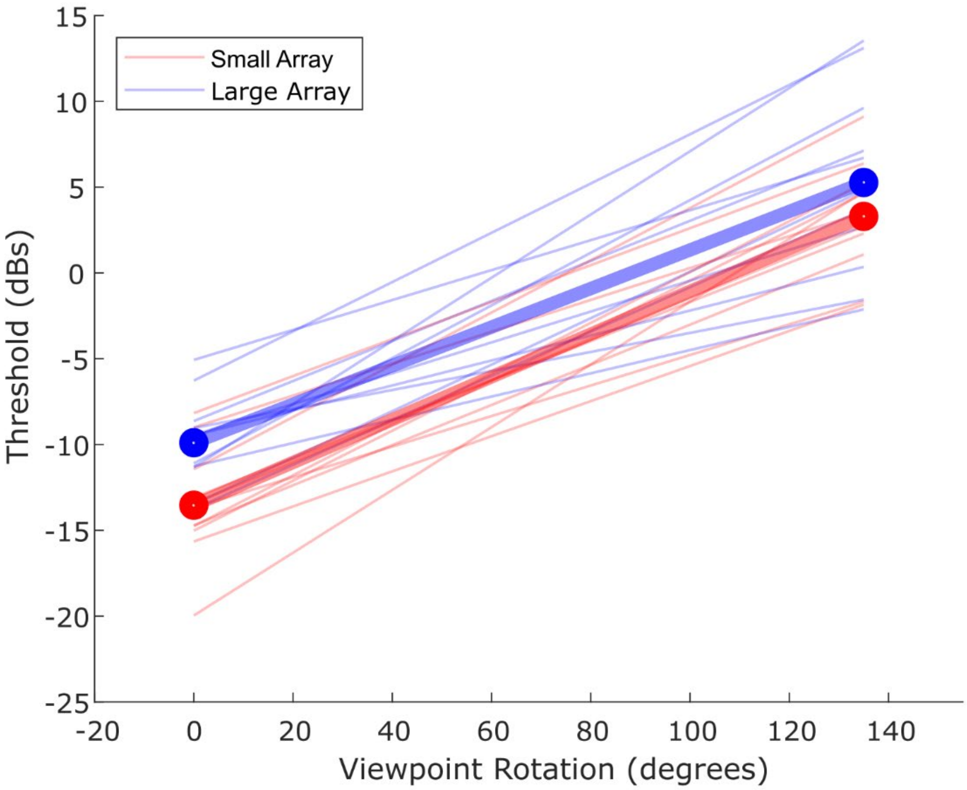

Figure 4.3 Experiment 4.1- participant thresholds for each condition . Small array, 0° and 135° conditions are connected by a thin red line, Large array, 0° and 135° conditions are connected by a thin blue line. The large circles and thick red and blue lines are mean thresholds for each of the 4 conditions.

4.3.3.3 Psychometric Slope Parameter: Effects of View Rotation and Array Size

We also conducted a repeated measures factorial ANOVA on the mean slope value for each condition. We found a significant effect of viewpoint $F(1,10) = 9.51$ ($p > 0.01$), a significant effect of scale $F(1,10) = 7.93$ ($p = 0.01$).

These effects seem to be driven by the slope value in the Large Array, 135° condition which is much shallower that slopes in the other conditions (see Figure 4.2 B). While there is no significant interaction between viewpoint and scale, the value is approaching significance $F(1,10) = 3.57$ ($p = 0.09$).

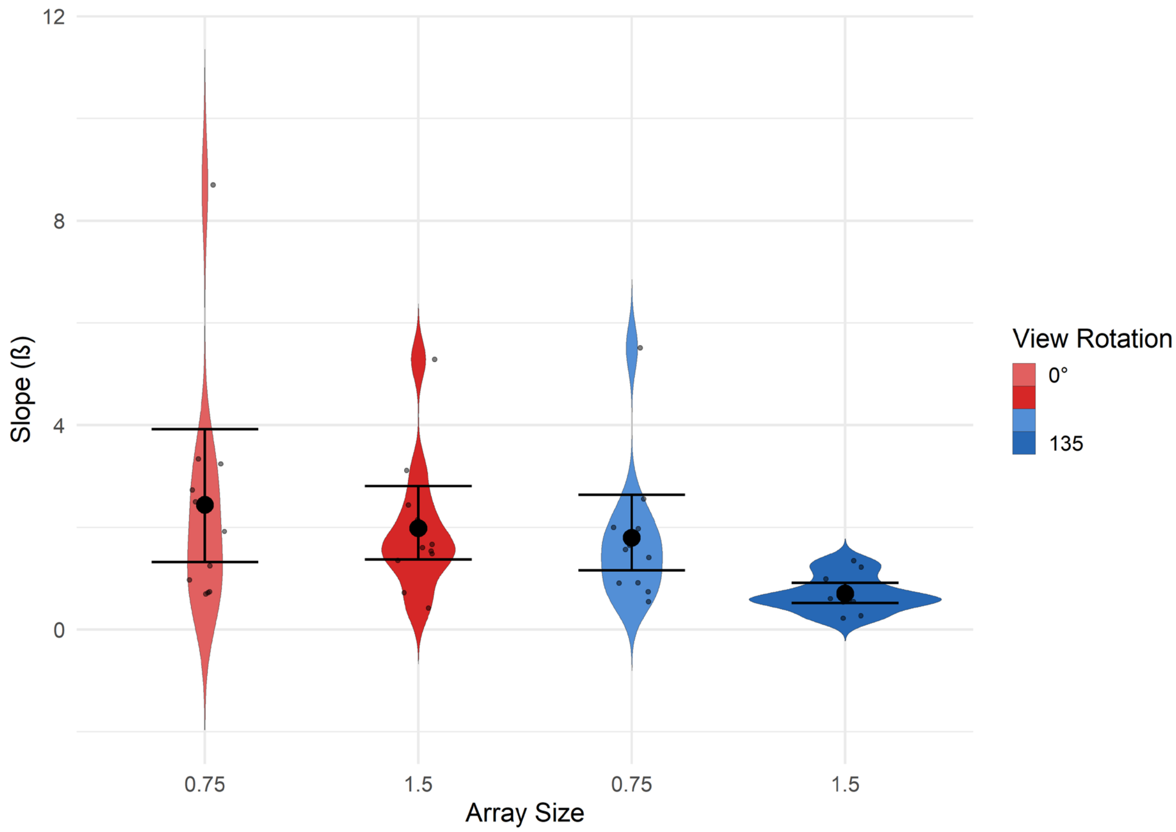

Figure 4.4 Comparing psychometric slope parameters as Weibull beta values (see section 3.3.5) between conditions . Small black points represent individual slope parameters from each participant in each condition and larger black circles represent mean slope parameters in each condition. Each violin plot describes performance in a separate condition. Array scaling factor is demarked on the x axis and colour demarks view rotation conditions (red for 0°, blue for 135°). A smaller beta parameter represents a shallower slope. Error bars showing standard error

4.3.3.4 Response time

In the experiments of the last chapter the trial structure meant that it was not meaningful to measure response times. In the current experiment, we are able to calculate response times for participants as they are able to respond immediately after the testing stimulus appears. However, this is an exploratory analysis as the experiment was not designed to measure response times and, for example, participants were not instructed to respond as quickly as possible.

Median response times were calculated for each participant and condition and a repeated measures factorial ANOVA was conducted on these values (see Figure 4.5). We found a significant effect of viewpoint change on response times $F(1, 10) = 19.5$ ($p < 0.01$), a significant effect of array size on response times $F(1, 10) = 21.4$ ($p < 0.01$) and also a significant interaction in the effects of viewpoint and array size $F(1, 10) = 8.7$ ($p = 0.02$).

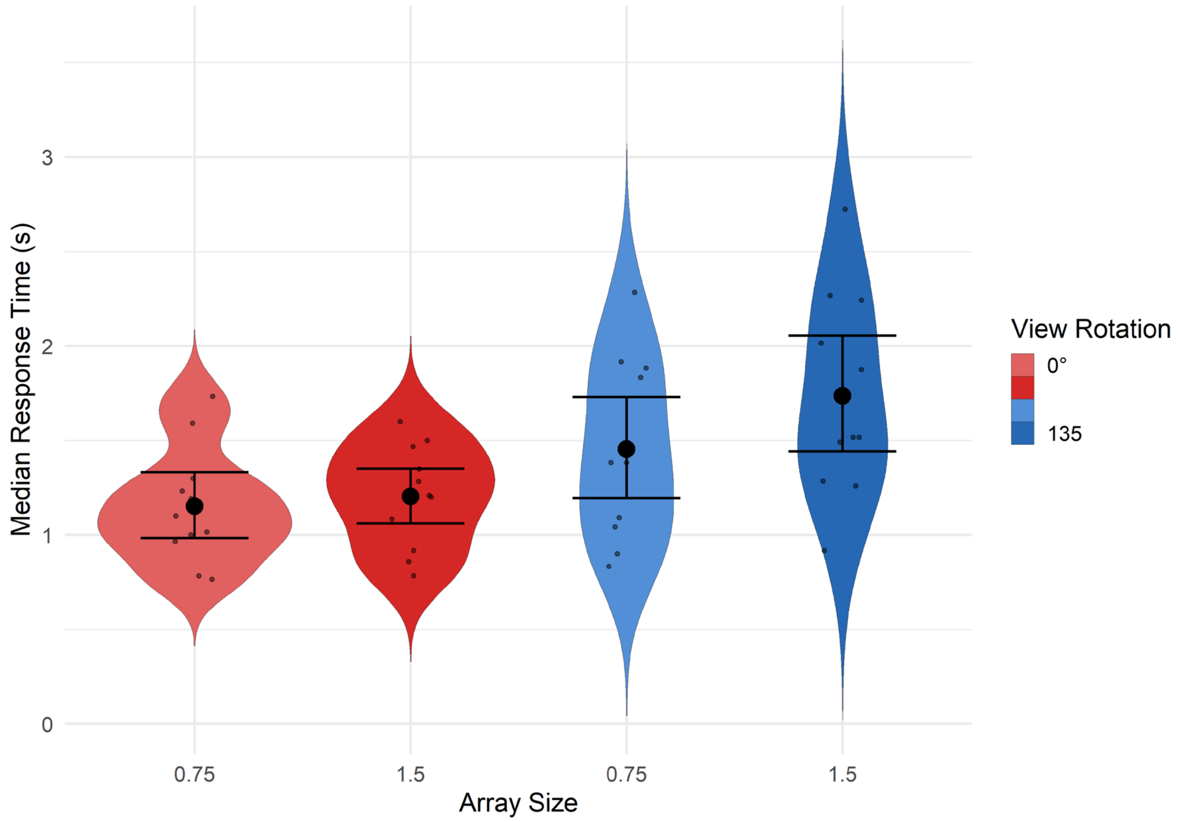

Figure 4.5 Violin plots showing median response times for each participant in each condition . The black points represent each participants individual median response times which are measured on the y axis. The violin plots show the distribution of these points in each condition. The viewpoint variable is shown by the violin colour, red shows 0° viewpoint shift and blue shows 135° viewpoint shift. The array scale conditions are also separated and plotted discretely on the x axis.

4.3.4 Discussion

The results show a significant effect of viewpoint on change detection thresholds. This replicates the findings in the previous chapter and demonstrates higher change detection thresholds in viewpoint-independent tasks.

Regarding array scale, responses to the Large Array condition showed significantly lower precision than to the Small Array condition. The effect of array scale on performance was similar in both the 0° viewpoint-dependent conditions and the 135° viewpoint-independent conditions and there was no significant interaction between array scale and viewpoint.

The main results, showing effects of scale and viewpoint but no interaction, suggest that both viewpoint-dependent and -independent memory are compatible with Weber's Law regarding perceived spatial changes to the array footprint and the scale of the initial stimulus, in that they both show comparable effects of scale. These results initially seem incompatible with the idea, formulated in the introduction, that viewpoint-independent precision is an absolute distance anchored to the environment (and therefore unchanged under manipulations of the scale of the array). Although the change in precision associated with the effect of scale was appreciably less than the expected 2:1 ratio, it was more similar to the actual mean ratio when comparing Small and Large Array footprint.

Comparing slope values we found similar findings, main effects of viewpoint and scale on mean slope values but no significant interaction. However, while non-significant, a clear numerical difference can be seen in the mean slope of the in the viewpoint rotated Large Array condition. Here the slope is noticeably shallower than in other conditions which can be interpreted as indicating less certainty in responses on these trials.

One consideration is that the reason for the similarity between the effect of scale on viewpoint-dependent and viewpoint-independent thresholds may be due to the information available for the formation of allocentric representations. Specifically, that allocentric representations are 'downstream' from egocentric perception and spatial memory. The process of creating and transforming allocentric representations from an initial egocentric perspective and also transforming back to an egocentric encoding upon retrieval (Bicanski & Burgess, 2018), may have implications on the resolution of data that we see in an allocentric type task. If the precision of the information provided to allocentric memory degrades according to Weber's law, then this 'egocentric bottleneck' could also determine limits for the precision of the allocentric spatial representation.

Another reason we did not see any interaction between viewpoint and scale on change detection thresholds, and in sigmoid slopes, might be due to some limitations in our methods. First, the change in array scale was quite small and so the experiment may have been underpowered to detect any correspondingly small effect of any interaction. It should be noted that there are practical constraints that limit our ability to create larger scale changes in this type of display; scaling the array any further would have resulted in objects appearing beyond the virtual test area, off screen, or required much smaller objects and arrays that would have been hard to see.

Second, the method of investigating spatial scale using a desktop virtual environment has several possible limitations. The scale cues available to participants in this desktop task with image stimuli are much less rich than would be available in a real-world task (which would include binocular vision and motion induced parallax depth cues). In addition, while the scale of the array footprint was manipulated, this inevitably affected the size that the array took up on the screen and in the participant's field of view. It is therefore possible that the lower precision seen in the larger scale array was due to factors such as increased saccadic eye movement in the encoding phase or extended search time due to the array of objects subtending a larger visual angle (Abrams et al., 1989; Kowler & Blaser, 1995). Finally, another possible confound relating to our manipulation of scale was that as well as changing the overall footprint of the array, it inevitably affects the spacing of individual objects, and it has been shown that this can influence the usefulness of viewpoint-dependent/independent information. Specifically, participants are able to make more effective use of static landmarks when then are in closer proximity to the remembered item (Aagten-Murphy & Bays, 2019; Clark et al., 2007).

The exploratory analysis of response times also showed main effects of viewpoint and scale and an interaction. With the caveat that this study was not designed to measure response times, this is an interesting result. The interaction appears to driven by participants' longer response times in the view-rotated Large Array condition. While this does not inform us about allocentric precision, it does indicate that there are differences in the way that Small and Large Arrays were processed by viewpoint-dependent and -independent memory systems. One possibility is that the process required to compare Test and Presentation arrays takes longer for physically larger stimuli when dealing with a viewpoint-independent task.

In sum, these results suggest that both viewpoint-dependent and viewpoint-independent precision changes with the scale of the array and that this change is a perceptually based one, associated with Weber's law. There was no evidence for a fixed level of precision between scale conditions in the viewpoint-independent task, suggesting that the threshold limits we see in the current task are unlikely to be solely due to the underlying resolution of a fixed scale allocentric representation. However, limitations to the desktop task mean make it difficult to rule out the possibility that an interaction involving stimulus scale might be detected in a more ecologically valid task, or one involving a much larger scaling manipulation, and it remains conceivable that Weber-like effects result from the interaction of allocentric representations with egocentric information during encoding or the generation of responses.

4.4 Experiment 4.2

4.4.1 Introduction

In Experiment 4.2 we addressed these issues by using immersive virtual reality to present the stimuli and a novel approach to manipulate participants' perception of scale independently of pictorial cues. This involved systematically changing the scale of the entire environment and the objects while at the same time moving participants away from the array by the same scaling factor. This meant that we could increase the perceived scale of the array and environment while the visual image (the relative size and location of the objects and environment and the visual angle subtended by the array and objects in the participants field of view) stayed the same (see Figure 4.5).

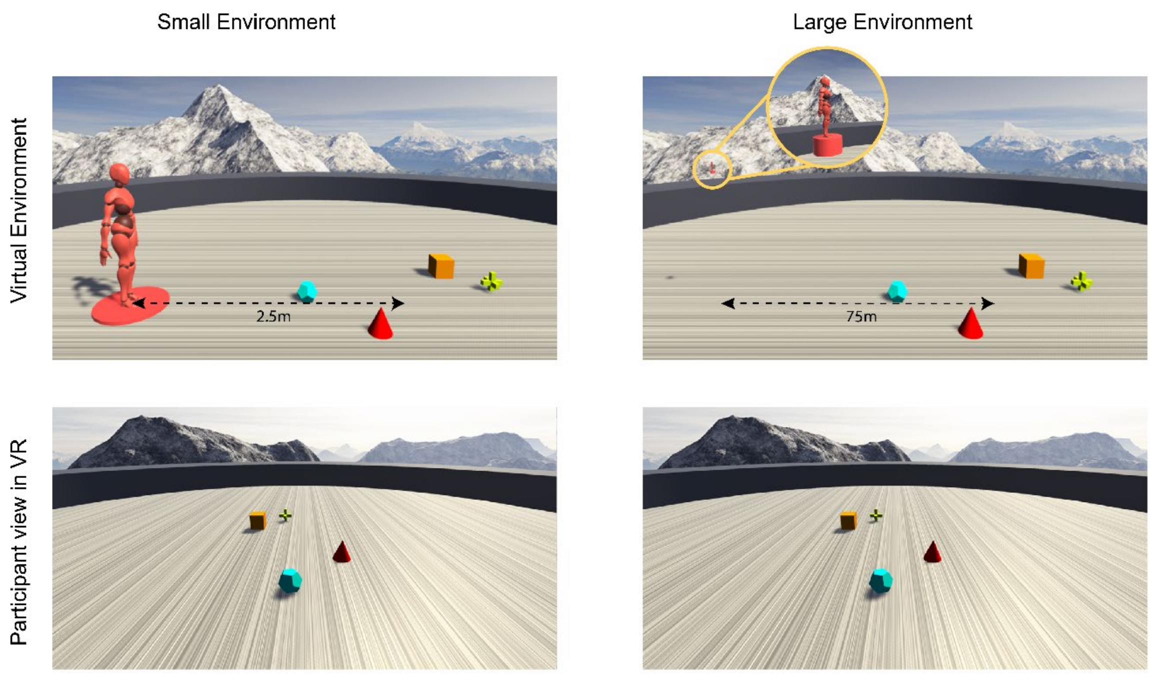

Figure 4.6 *Experiment 4.2 task and materials . Showing the small- and large-scale conditions from an adjacent view and the scene and the participant's view in VR. In the Small Environment, the participant is standing at ground level, at their natural height looking at the array of 4 objects. In the Large Environment, everything but the participant and the red pedestal has been enlarged by 30x. The participant is moved backwards and upwards to preserve the visual image of the array which is identical to the small scale. However, due to binocular vision and depth cues, the experience is that of a larger scale.*

As using this method means that the visual image of the stimulus is identical across any scale manipulation, the scale of the environment is communicated with the additional spatial cues that are available in immersive virtual reality e.g., binocular depth cues and parallax due small head and body movements. This means that participants have a much more natural and realistic sense of environmental scale and distance than the ambiguous scale and distances inferred from an image on a 2D display. This has been shown experimentally with factors such as binocular cues, natural movement and high-resolution headsets being contributing factors to improve accuracy of size/distance judgements in VR (Eggleston et al., 1996; Interrante et al., 2006).

This approach also means that the manipulation of scale is not limited by the size of the display and in Experiment 4.2 environments will be scaled up by 30x in the large condition. Note though, that as the effectiveness of cues based on binocular vision and parallax drops sharply after just a few meters distance from the body we do not expect participants to have a wholly veridical sense of scale. Therefore, as a manipulation check, to ensure that this very large scale change was clear to participants we also include an additional Scale Assessment Task.

As in Experiment 4.1, we would expect viewpoint-dependent spatial precision to be poorer, in absolute terms, at larger scales – as predicted by Weber's law. If we also see this decrease in viewpoint-independent precision, it would replicate the findings of Experiment 4.1 and suggest that viewpoint-independent precision is not fixed and absolute but affected by the scale of the stimulus. However, with the much larger scaling factor used in the current experiment, it might be possible to detect an interaction that was not detectable in Experiment 4.1.

The methods used in Experiment 4.2 allows us to control the 2D pictorial cues in the stimulus/environment while changing the perceived scale. They also allow us to query whether the effects of scale on precision seen in Experiment 4.1 were due to a change in the absolute size of the array or due to a change in the visual angle that the array subtends in the participant's field of view. If spatial change detection thresholds scale 1:1 with the environmental scaling factor, this would further suggest that the change detection thresholds are dependent on the visual image rather than the absolute scale of the environment.

4.4.1.1 Experiment 4.2 Hypotheses

As seen in previous experiments, viewpoint change is expected to be associated with increased spatial change-detection thresholds (replication of previous pre-registered experiment). When the perceived scale of the environment is changed but the pictorial cues remain the same, we expect that viewpoint-dependent spatial memory precision will change according to Weber's law and we will investigate whether viewpoint-independent memory precision also changes in line with the viewpoint-dependent memory, or whether we see an interaction with viewpoint change such that precision is less affected by changes of scale following a shift of viewpoint between presentation and testing.

We will test whether changes in spatial change detection thresholds are proportional to the (perceived) changes in in environmental scale. If so, the most parsimonious interpretation would be that spatial precision is correlated with the preserved visual angles in the stimulus as it appears in the participants' field of view, rather than its absolute size.

Finally, we predict that our manipulation check will show the scale manipulation is clearly perceived by all participants, although we do not expect participants' estimates to be veridical (due to the imprecision of parallax and stereoptical cues in very large environments and the relative poverty of non-pictorial cues in our simplified experimental environment).

4.4.2 Methods

4.4.2.1 Overview

In Experiment 4.2, participants completed a spatial change-detection task, making 2AFC judgements about the locations of objects in an immersive virtual reality environment. As in Experiment 4.1, a 2x2 factorial design was used to manipulate both the degree of viewpoint shift between presentation and testing and also the array scale. However, unlike in Experiment 4.1 where only the array footprint changed scale, in this task everything in the environment including scenery, array footprint and object size was changed. Only the scale of the participant's own body and the red pedestal on which the participant stands remained the same size (for example see Figure 4.5). While environmental scale was increased, the participant was moved upwards and backwards (away from the centre of the environment) so that relative spatial relationships between objects and the observer were preserved and the visual image of the stimulus remained the same. The scale of the environment was communicated though binocular and combined parallax and proprioceptive cues (resulting from small postural changes) using a VR headset.

Importantly, scales were varied using a counterbalanced block design such that, in a given session, participants would experience only a single scale. Viewpoints were varied randomly at either 0 or 135° view change. Viewpoint conditions were randomised on a trial-by-trial basis throughout each scale block.

For each participant, a spatial change detection threshold was estimated for each of the four experimental conditions based on the proportions of correct/incorrect response as detailed in Chapter 2.

This experiment was piloted and preregistered prior to any data collection (https://osf.io/hdwve/).

4.4.2.2 Participants

Sample Size:

All participants were recruited via the University of York's Department of Psychology participant recruitment system, 15 participants (12 female) who each completed a mean of 273 trials.

Sample size was based on a power analysis using the effect sizes and the same power calculation from Experiment 4.1. This indicated that a sample size of 12 would be sufficiently powered to allow for the detection of effects (at a $p<0.05$ level of significance) with an effect size of 0.8, $F = 0.4$.

Participant Exclusion Criteria:

Participants were excluded if they performed below 66% correct on 30 'high intensity trials' (set at the highest intensity available to the staircase) in the tutorial phase. No participants were excluded based on this criterion.

Participants were also excluded if they failed the stereopsis test (see below, 1 participant was excluded due to this criterion).

4.4.2.3 Data Exclusion Criteria

As in Experiment 4.1, if the threshold estimation technique provided a single participant with a threshold of above 15dB, this was significantly beyond our testing range (max 14.0dB) and indicated that we were unable to successfully model performance. The participants' results were recorded and included within the OSF data folder (https://osf.io/hdwve/) but excluded from subsequent modelling and analysis. Data was also excluded if curve fitting failed to converge. In total 3 participants were excluded due to these criteria.

4.4.2.4 Materials

To ensure that all participants were able to make use of immersive VR cues to environmental scale we used a Randot test of stereopsis (Randot Stereotests, Stereo Optical Inc.) as a screening tool.

As in previous experiments stimuli were dynamically created for each trial. Task and stimuli were rendered and controlled by custom software programmed in the Unity3d game engine in C# running on an HP Z400 computer. Stimuli/environments were presented to participants using a HTC Vive virtual reality head-mounted display (HMD; an OLED display with 1080x1200 pixel resolution per eye, 90 Hz refresh rate and 110° field of view). Virtual environments and experimental procedure were implemented with the Unity3d game engine. All testing took place in a quiet room in Department of Psychology, University of York. The physical location of the HMD and controller is tracked with 0.3mm accuracy such that small changes in the participant's pose and movements are accurately reflected in their 1st person view of the environment at all times. The same virtual environment and objects were used in Experiment 4.1 and Experiment 4.2.

As in Experiment 4.1, participant's viewpoint could be varied between presentation and testing to either 0°, or 135°.

In the environment there was a mountainous backdrop (skybox), set to a draw distance of infinity which did not change with scaling. All aspects of the environment were scaled including object size, array footprint, floor plane and environmental boundary. Only the small red circular platform (1m diameter) upon which the participant stood and the participant themselves (i.e., the virtual interpupillary distance provided to each screen in the VR headset, and the apparent height of their head above the virtual platform) and remained the same absolute scale across both scale sessions. In the Small Environment (1x), participants were 2.5m from the centre of a 10m diameter circular area and in the Large Environment (30x) participants were 75m from the centre of a 300m diameter circular area.

In the Large Environment participants were also elevated to maintain the same angle of inclination from their eyes to the centre of the array as they would naturally while standing up in the Small Environment (participant eye level × 30m). Without binocular vision cues, parallax and idiothetic body cues indicating the scale of the environment, the small and large environment would be identical (i.e., the lower images in Figure 4.5).

It is important to reiterate the qualitative experience of the stimuli in VR as the experience of scale is impossible to convey with images. When experiencing the stimuli, it was strikingly clear in the Large Environment that one was in a vast space and standing ~50m above the ground.

As a manipulation check and to ensure that all participants perceived the scale manipulations, an additional Scale Assessment Task was included involving estimation of real-world distances in VR to allow us to estimate participant perception of scale (see below for full procedure). This task involved the same environment and a 1m wooden reference rod attached to the controller (Figure 4.8). The reference rod was also accurately represented in the virtual environment, so that participants could both feel and see an accurate representation throughout the task.

4.4.2.5 Adaptive Staircases

The same adaptive staircase procedure as in Experiment 4.1 were used to target the participants spatial change detection thresholds (see Appendix F).

4.4.2.6 Spatial distortion range

The amount an object was distorted was governed by the current staircase level and the block (scale). Each staircase had 25 'steps' or testing levels distributed across a testing range. The testing range was centred on a 0 dB reference value (0.5 m Small Environment, 15m Large Environment) and spanned a range of -18 to 14dB. An equally spaced decibel scale was used which, when converted into meters provides a log-spaced test range centred on the applicable dB reference value (see Appendix F for details of staircase parameters).

4.4.2.7 Procedure

The experiment followed a two-block design, separating trials from each scale (Small Environment and Large Environment) condition into separate sessions. The experiment was counterbalanced and the order in which participants undertook these blocks was randomised based on participant number (odd participant recruitment ID undertook the Small Environment block first and vice versa).

Overall, participants experienced a trial structure and a sequence of trials very similar to that used in Experiment 4.1 and in earlier studies in the thesis. The main differences are that the stimuli are presented in immersive VR making it possible to represent the same pictorial image at different environmental scales, and that participants experienced the different scale conditions in separate sessions.

Introduction to VR

After a stereopsis screening using a Randot Stereotest, participants were instructed to stand in the centre of the testing area and were helped to put on the VR headset. They were given the controller-rod assembly (referred to as the ‘reference rod’) to hold in their dominant hand and a separate HTC Vive controller in their non-dominant hand.

In VR, participants were ‘placed’ at ground level in the centre of the array and presented with a row of array objects in front of them. This was to acclimatise them to the experience of virtual reality and also to more clearly demonstrate the scale of the environment. In the Small Environment, objects were rendered at approximately 15cm in height whereas in the Large Environment they were rendered at approximately 7.5 m in height, physically dwarfing the participant. The experimenter recorded the participant’s height by asking them to stand up straight and automatically recording the vertical axis position of the headset. This height measurement was used to determine the position that the participant would be raised to in the Large Environment in order to match their natural angle of inclination from their eyelevel to the array centre when standing on the ground in the small condition.

After looking around the environment, participants were prompted to press a button on the controller to begin the spatial change-detection task. This had the effect of moving their virtual position to the testing position at ground level. In the Large Environment only, they were then slowly raised upwards on a small red disk-shaped platform until their view of the environment was equivalent to their natural height scaled up by a factor of 30. The slow movement raising participants upwards to this height was to reduce any sense of vertigo (moving approximately 50 virtual metres upwards on a small platform from ground level), while also serving to reinforce the scale of the environment.

Spatial Change-detection Task

The trial timings (e.g., presentation and response time) were identical to Experiment 4.1. In both the Small Environment and Large Environment blocks, participants followed the same procedure. On each trial, as in Experiment 4.1, they were presented with an array of objects before the display faded to black. One object was moved a distance specified by an adaptive staircase and, as in previous experiments, participants were either moved around the array by 135° (on viewpoint change trials) or remained in the presentation position. The display faded back from black to the new viewpoint and two objects were highlighted. Participants were prompted to select the highlighted object that they thought had moved. In the response phase, a virtual "laser pointer" emitted from the controller in the participant's hands. To respond, they pointed the beam at the object that they thought had moved which highlighted it. To confirm their response, they pressed the trigger on the controller. After responding, a correct or incorrect sound provided feedback to participants based on their response.

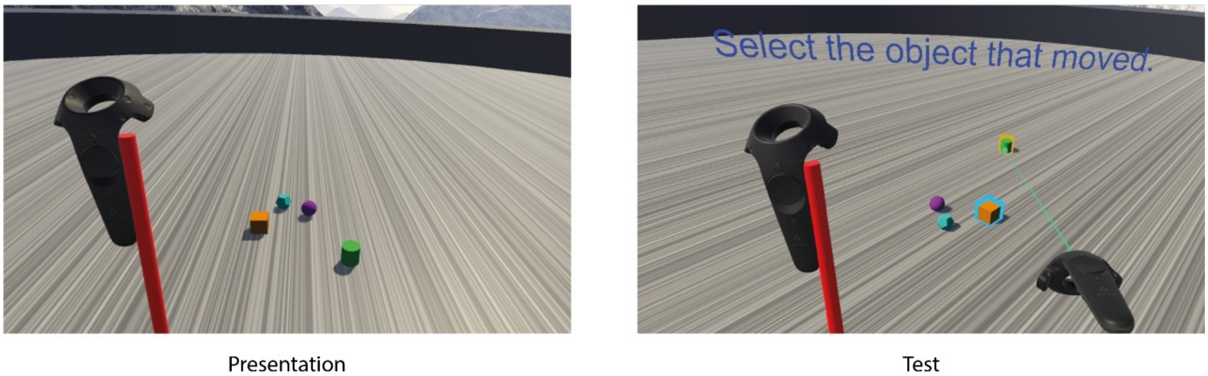

Figure 4.7 Example of response in Experiment 4.2 . At test, participants are prompted to 'Select the object that moved'. The right hand controller begins to emit a pointer with a small ball at the point the beam makes contact with object or environment. To select an object, participants aim the pointer at one of 2, blue glowing objects. When selected, the object's outline turns from blue to orange. To confirm selection the participant pulls the trigger on the controller.

Scale Assessment Task

The additional Scale Estimation Task took place every 30 trials in the same environment. The Scale Assessment Task was included to test how accurately participants' sense of scale matched the actual scale of the environment. Participants saw a rod in the environment which appeared at a random location (within the array area) and at a random size (between 1 and 30 meters, see Figure 4.8 for illustration). They were instructed to adjust the length of the distant rod in the environment to match the length of the reference rod in their hand. In both the Small and the Large condition, the size of the rod in their hand was 1m, exactly matching the size and position of the real wooden dowel attached to the controller, providing coherent haptic cues. Participants adjusted the size of the environment rod with buttons on the control pad until they judged that it matched the length of the rod in their hand and to confirm their size estimate, participants pressed the trigger button.

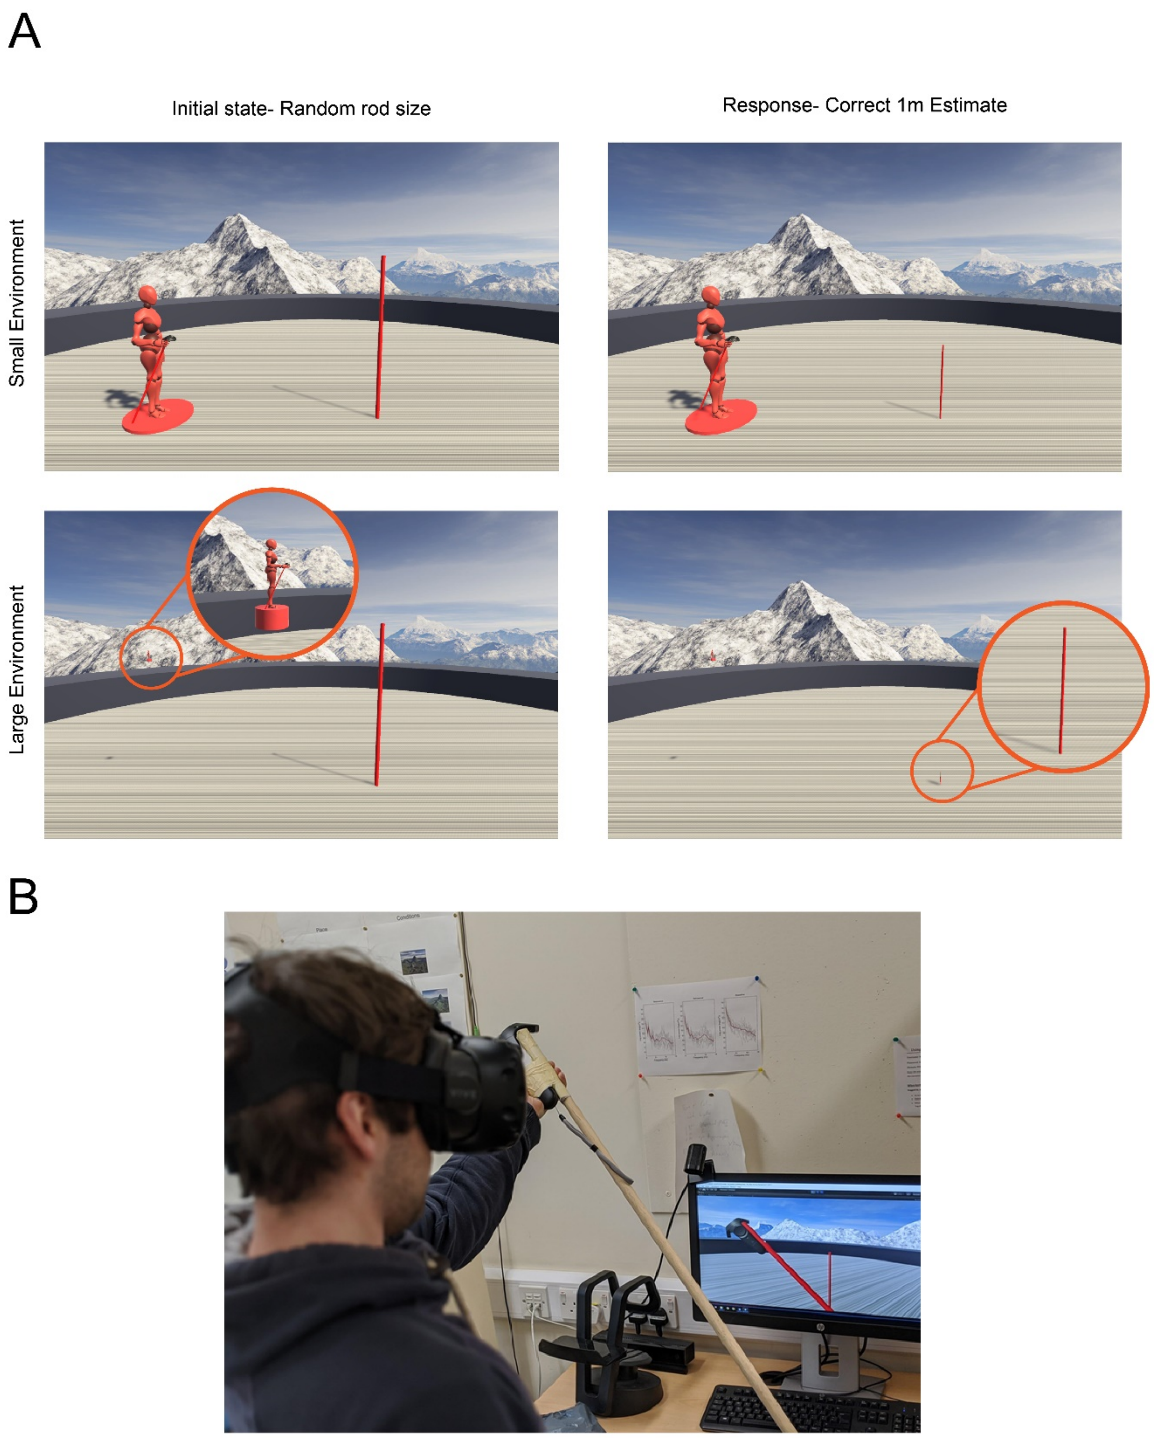

Figure 4.8 Diagram of the Scale Assessment Task . A shows the initial state, when a random sized rod is presented to participants in the Small and Large Environment conditions (2 images left) and the response state when the participant has responded with a correct estimate of 1m (2 images right). B shows a participant using the 1m rod to complete the Scale Estimation Task in VR. Here the participant holds a 1m rod attacked to the control pad. It has a digital analogue in VR which can be seen on screen in red and identically matches the dimensions of the real rod. The participant must use the control pad buttons to change the length of the red rod in the environment to match that of the one in the participant's hand.

4.4.2.8 Analysis

As in Experiment 4.1, thresholds were estimated with a Maximum Likelihood curve fitting procedure as implemented in MATLAB using fminsearch and as described in the previous chapters (see Chapter 2, section 2.3).

Comparing the effects of scale and viewpoint on change-detection thresholds

To analyse the effect of scale, we carried out two complimentary analyses. We first compared change-detection thresholds as we have done in previous experiments, using decibels to represent the absolute distances in which the object moved (which we will refer to as dB Abs, Equation 4.1). We then separately compared the thresholds after they had been scaled by the environmental scaling factor (dB Scaled, Equation 4.2). A magnitude of *n* dB Scaled would subtend the same visual angle in the participant's field of view in both Small and Large Environment. For example, if this task took place viewed on a desktop monitor without the VR scale cues an *n* dB Scaled movement in any scale environment would be identical.

Equation 4.1

$$dB\ Abs = 20 \times\ Log10(\frac{spatial\ change(m)}{reference\ value})$$

Equation 4.2

$$dB\ Scaled = 20 \times Log10(\frac{spatial\ change(m)}{reference\ value \times scale\ factor})$$

To compare the performance in both viewpoint conditions and both scales as well as any interactions between them, two 2x2 factorial ANOVAs (viewpoint x scale) will be used to compare dB Abs and then dB Scaled thresholds.

If we see no main effect of scale in our second analyses (i.e., dB Scaled thresholds are not significantly different across scale conditions) this means that thresholds do not differ significantly from what would be expected if they had scaled with the environment in Weber-like fashion. This would tend to support an interpretation in which precision is constrained by the (preserved) pictorial properties of the visual image and is not affected by the perceived absolute scale communicated via the additional VR depth cues. Note that these pre-planned analyses are not independent (dB Scaled is a simple transformation of dB Abs), and the aim is to provide a fairly intuitive test in which each potential outcome is compared with suitable null hypotheses.

Scale Assessment Task: Manipulation check

Size estimates given by participants in the Scale Assessment Task in both environmental scales were also analysed in order to check the validity of the scale manipulation. We wanted to ensure participants had experienced a change in the perceived scale of the environment.

In both conditions participants were required to alter the size of a distant (‘target’) rod to be 1m (the size of the ‘reference rod’ which is always 1m). We can use this as a “yardstick” with which to measure the perceived scale of the environment. For example, overestimating the size of the rod is equivalent to underestimating the scale of the environment.

Equation 4.3

$$perceived\ environment\ diameter = actual\ diameter\ \times\ perceived\ scaling\ factor$$

$where$

$$perceived\ scaling\ factor = \frac{reference\ rod\ size}{target\ rod\ estimate}$$

To check that participants perceived a sense of scale arising from the VR cues we used a paired samples t-test to compare perceived environment diameters in the Small and Large Environment. If participants experienced a change in scale we would expect this t-test to show a significant difference in their perceived environment diameter.

4.4.3 Results

4.4.3.1 Psychometric curve fitting

We estimated four spatial change-detection thresholds based on an average of 819.25 trials per participant. These corresponded to the degree of spatial distortion associated with 75% accuracy (Small Environment, 0° rotation; Small Environment, 135° rotation; Large Environment, 0° rotation; Large Environment, 135° rotation). Threshold values we also scaled by the environmental scaling factor.

Table 4.2 showing mean of best fitting slopes for each condition and participant . Each value is the mean of separately fit psychophysical functions for each condition and each participant.

Condition (viewpoint/scale)	Mean Threshold dB Abs (SD)	Mean Threshold dB Scaled (SD)	Mean Slope (SD)
0° Small Environment	-15.04 (2.02)	-15.04 (2.02)	3.75 (3.35)
135° Small Environment	-0.49 (5.47)	-0.49 (5.47)	4.4(2.12)
0° Large Environment	14.66 (2.50)	-14.89 (2.50)	3.12 (2.06)
135° Large Environment	31.84 (3.15)	2.3 (3.15)	3.57 (2.46)

To measure how well we could apply a psychophysical curve to the data, a Bayes factor was calculated for each fit sigmoid comparing the fit of the psychophysical curve to a null hypothesis (i.e., the likelihood of the data under a uniform model where the proportion of correct responses, the mean across all trials, is unaffected by the degree of spatial distortion). As in Experiment 4.1, participant responses in all conditions were well described by the psychometric with all Bayes Factors giving evidence for H1 (based on interpretations from Kass & Raftery, 1995, see Appendix E).

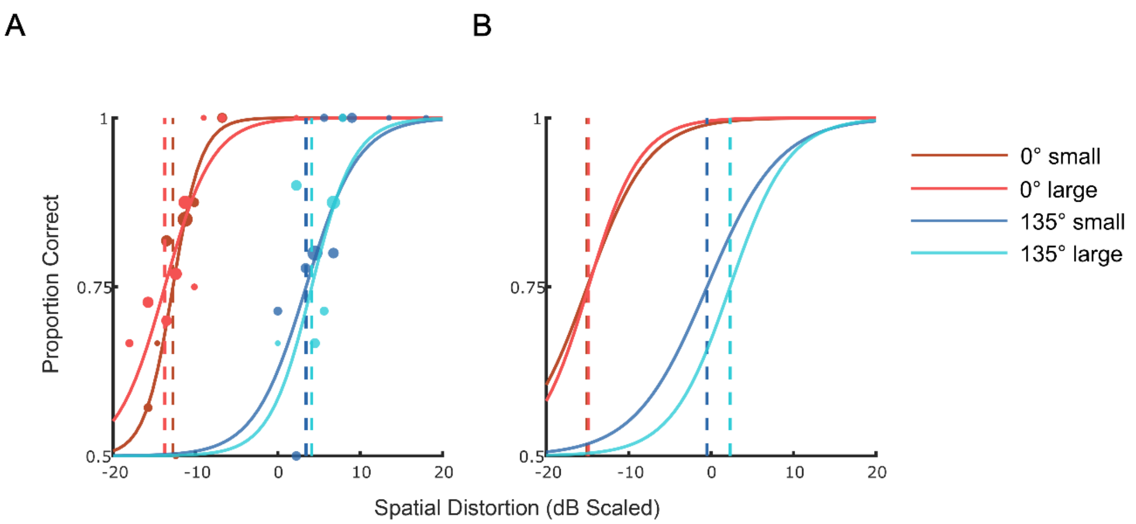

Figure 4.9 Psychophysical sigmoid curves for each experimental condition . Different conditions (viewpoint x array scale) are depicted with different colours. Red denotes 0° viewpoint shift and blue denotes 135° viewpoint shift. A: Representative data from a single participant. Circle symbols indicate proportion correct at specific staircase levels (in dB relative to 0.5m) and the size of each marker representing the number of trials undertaken at that level. Dotted lines show the estimated threshold level at 75% performance. B: group mean sigmoid curves (created by averaging fitted threshold and slope values for each condition). Values are in 'dB Scaled' i.e. scale corrected (see methods).

4.4.3.2 Effects of viewpoint and scale on spatial precision

Absolute change-detection thresholds

Using absolute spatial change detection thresholds (dB Abs, see Table 4.2) and conducting a 2x2 factorial repeated measures ANOVA we found that, as expected, as thresholds increased in the 135° viewpoint change condition. This change was statistically significant (mean threshold of -14.96 at 0°and 0.91 at 135°; $F(1,11) = 190.00$ $p < 0.01$). We also found a very large effect of scale $F(1,11) = 2028.58$ $p < 0.01$. This is unsurprising as the testing range as absolute distance had only a very small overlap (see coloured rectangular areas in Figure 4.10 A). We found no significant interaction between viewpoint and scale although these values were approaching conventional statistical significance criteria $F(1,11) = 3.85$ $p = 0.08$. Inspecting the individual data, the difference in thresholds in the different scale conditions appears to be driven by a small number of people who showed a smaller effect of viewpoint shift in the Small Environment condition (Figure 4.10 B).

Scaled change-detection thresholds

Converting change-detection thresholds into dB Scaled (i.e., adjusting for the scale factor used to generate the VR stimuli) and conducting a 2x2 factorial ANOVA, we necessarily see the same effect of viewpoint ($F(1,11) = 190.00$ $p < 0.01$) and the same non-significant interaction between viewpoint and scale $F(1,11) = 3.85$ $p = 0.08$, the relationship between viewpoint means and variance is the same between dB Abs and dB Scaled.

However, comparing scaled values we find no significant effect of scale on change-detection thresholds (although we do see data that is approaching statistical significance; mean threshold of -7.76 in the Small Environment and -6.29 in the Large Environment), $F(1,11) = 4.56$ $p = 0.06$). This means that change-detection thresholds in both viewpoint-dependent and -independent conditions are not significantly different from those that would be expected had they scaled 1:1 with the cues in the VR environment and the preserved pictorial cues which subtend the same visual angle in the participant's field of view in both scale conditions.

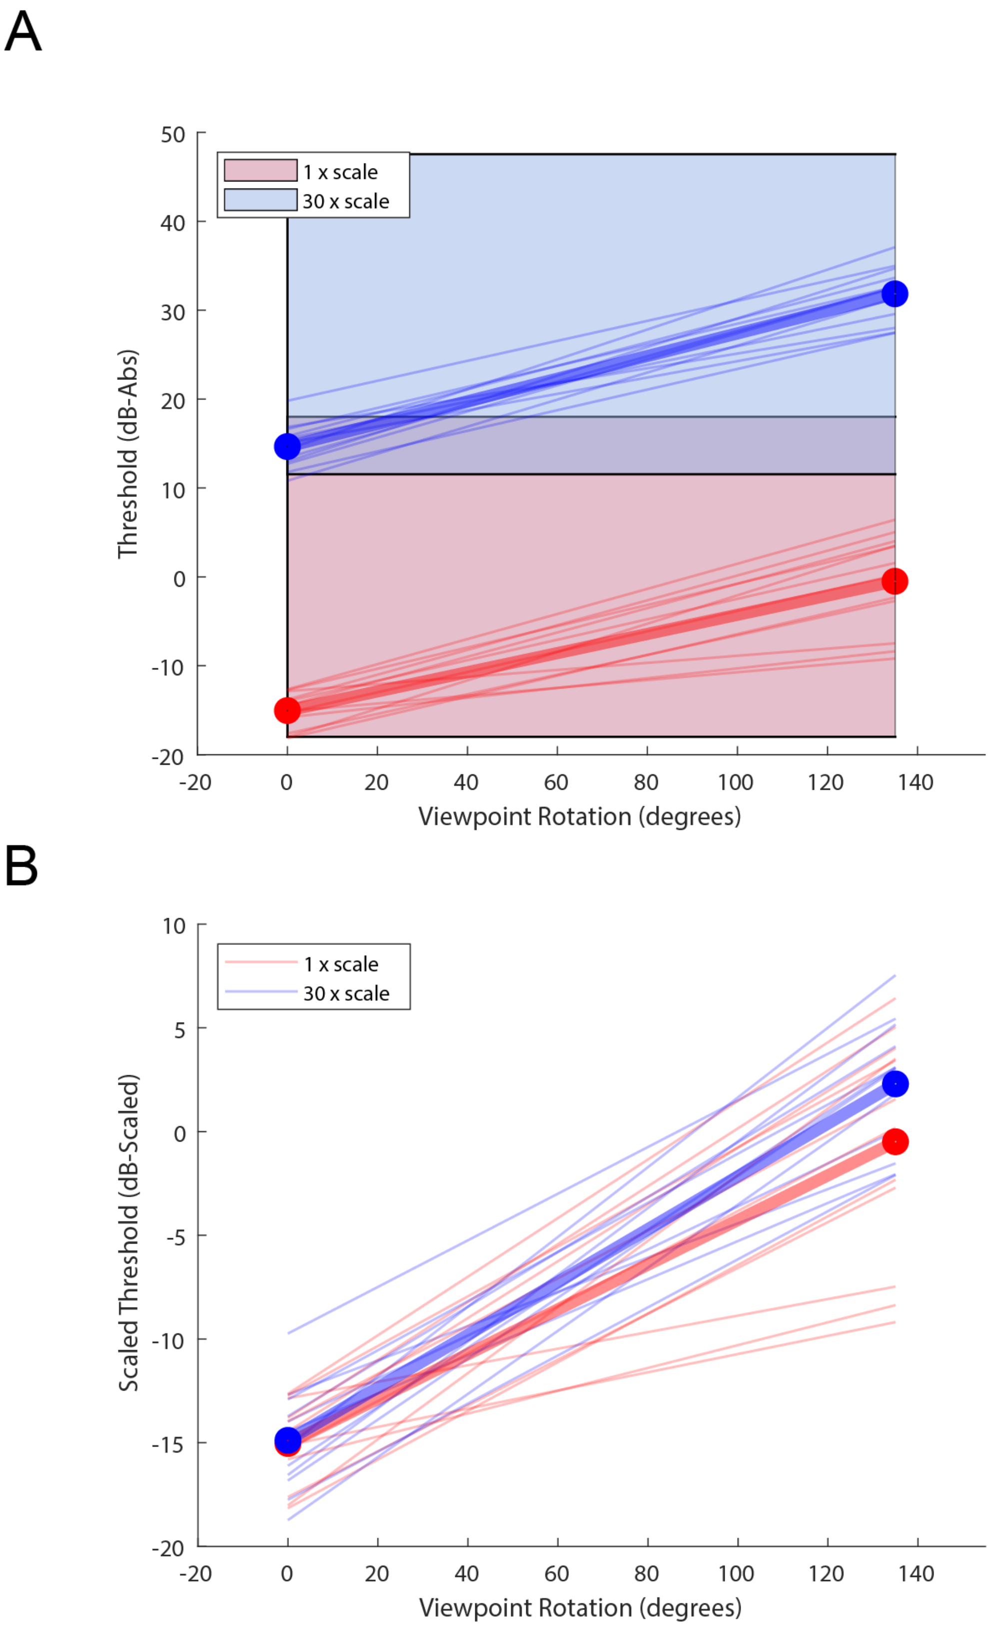

Figure 4.10 Participant thresholds in Experiment 4.2 expressed in dB Abs and dB Scaled . A and B show participant thresholds for each condition. 0°, Small and Large Environment conditions are connected by a thin red line, 135°, small and large conditions are connected by a thin blue line. The large circles and thick red and blue lines are mean thresholds for each of the 4 conditions. Figure A shows thresholds in dB Abs, using the actual and absolute Cartesian distances that objects moved in the environment. The blue and red boxes show the testing range, the adaptive staircases did not go beyond the

upper and lower bounds of the testing range. Figure B shows thresholds converted into dB Scaled (see Methods: Analysis for full conversion information). Here the absolute distances that objects moved are scaled with the condition scaling.

4.4.3.3 Psychometric Slope Parameter: Effects of View Rotation and Environment Size

As in Experiment 4.1, we conducted a repeated measures factorial ANOVA on the mean slope value for each condition. We found a no significant effect of viewpoint ($F(1,11) = 0.46$ ($p = 0.51$)), scale ($F(1,11) = 0.81$ ($p = 0.39$)) and no significant interaction ($F(1,11) = 0.02$ ($p = 0.88$)).

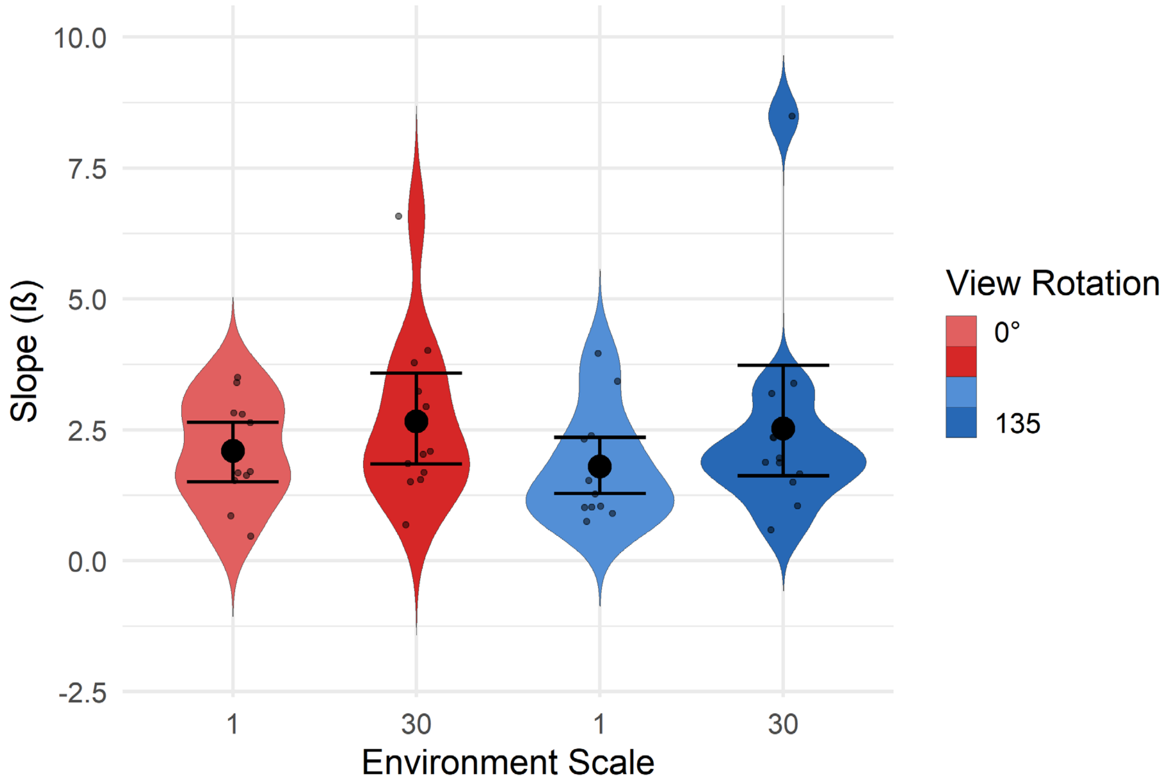

Figure 4.11 Comparing psychometric slope parameters as Weibull beta values between conditions . Black points represent individual slope parameters from each participant in each condition and larger black markers represent mean slope parameters in each condition. Each violin plot describes performance in a separate condition. Environment scale factor is demarked on the x axis and colour demarks view rotation conditions (red for 0°, blue for 135°). A smaller beta parameter represents a shallower slope.

4.4.3.4 Perception of scale and scale estimation task

Participants were asked to change the size of a rod in the environment to match 1m (the size of the reference rod).

Manipulation check

A value of perceived environment diameter was calculated by comparing participant target rod size estimates to the size of the reference rod (see Figure 4.12 and section 4.3.2.7).

To test whether participants perceived a change in environmental scale a paired samples t-test was carried out to compare mean perceived environment diameter in the Small and Large Environment. It was found that there was a significant difference in diameter size estimates between conditions $t = -4.02$, ($p < 0.00$). These results indicate that the manipulation of scale was very clear to participants.

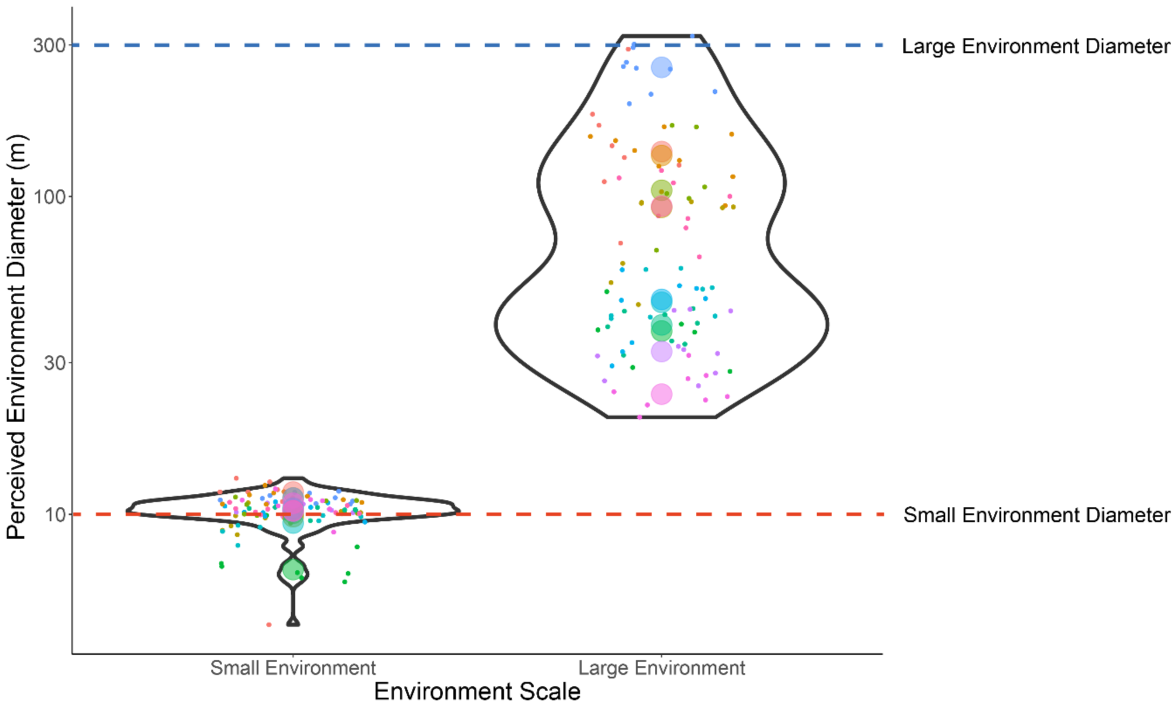

Figure 4.12 A bee-swarm, violin plot of perceived environment diameter for each participant . Perceived environment diameter (m) was calculated by comparing participant rod estimates to the correct rod length (1m); see section 4.4.2.7. An individual participant is represented by a specific colour. Small points are raw data, showing perceived environment diameter size in the Small (left) and Large (right) Environment. Large dots show individual participant means.

Individual perceived scale factor

While participants were clearly able to perceive the dramatic difference between scale conditions their estimates within the large environment were rather inaccurate overall (most participants adjusted the rod to be much larger than it would appear in an accurately scaled environment, underestimating the size of the environment), and they were highly varied and idiosyncratic. An exploratory analysis was conducted to test whether the accuracy of the Scale Assessment Task correlated with threshold accuracy as estimated by the 2AFC task. There was no sign of a correlation between accuracy in the Scale Assessment Task and spatial change-detection thresholds in any of the four conditions (all rs<0.28, all p>0.05).

One noteworthy point from the Scale Assessment Task data is that, in the Large Environment, while there is a lot of between-subjects variability (range = 13.9, SD = 3.56), there is much less within-subjects variability (max range = 5.36, max SD = 1.91). This suggests that participants' inaccuracy was systematic rather than random and that each participant had a consistent perception of environmental scale throughout the experiment (see distributions of coloured points in Figure 4.12).

4.4.3.5 Response times

Median response times were calculated for each participant and each condition, and a repeated measures factorial ANOVA was conducted to compared differences between conditions. It was found that there were significant effects of both viewpoint and environmental scale on median response times (for viewpoint; $F(1,11) = 26.3$, $p < 0.01$, and scale; $F(1,11) = 32.6$, $p < 0.01$). However, there was no significant interaction between viewpoint and scale, $F(1,11) = 0.7$ ($p = 0.4$).

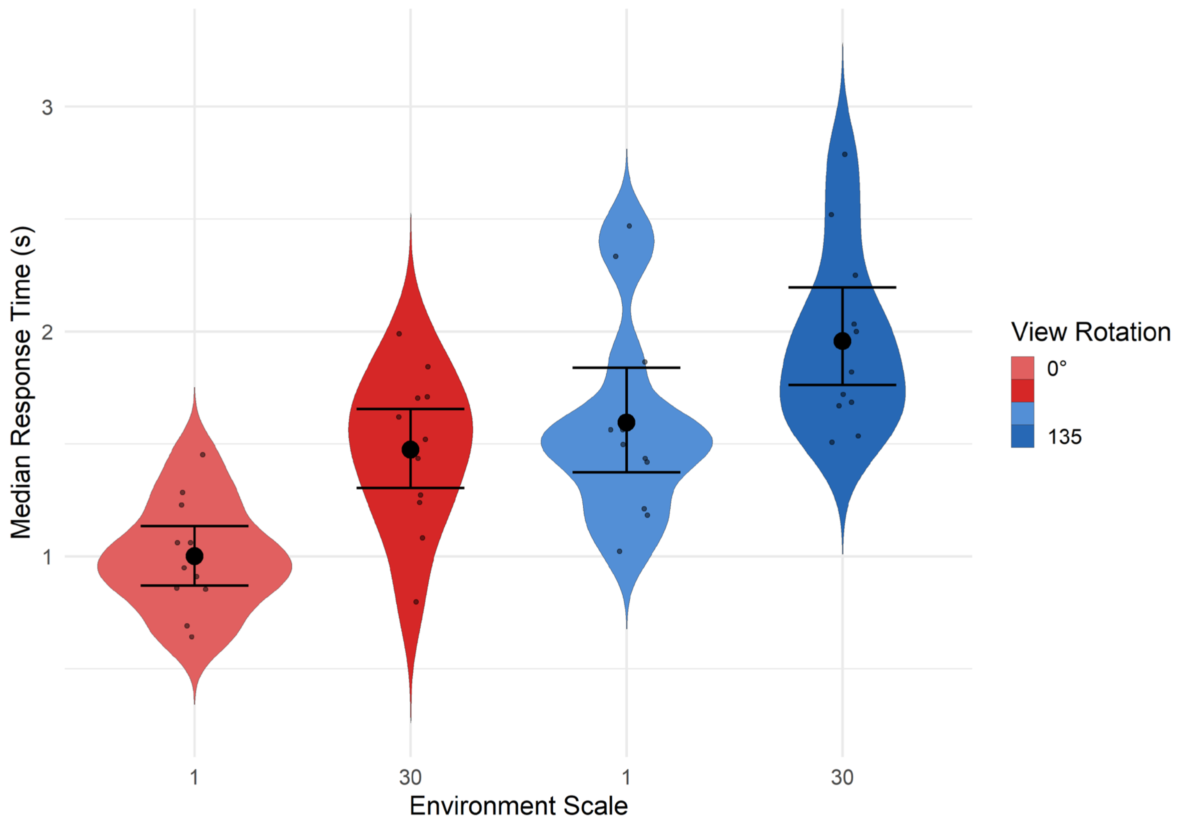

Figure 4.13 Median response times for each participant in each condition . The black points represent each participants individual median response times which are measured on the y axis. The violin plots show the distribution of these points in each condition. The viewpoint variable is shown by the violin colour, red shows 0° viewpoint shift and blue shows 135° viewpoint shift. The array scale conditions are also separated and plotted discretely on the x axis.

4.4.4 Discussion

Before exploring the main results, it is necessary to briefly cover the results from the Scale Assessment Task relating to hypothesis 4 which is necessary for the validity of the scale manipulation in other results. By extrapolating the perceived environmental scale from rod size estimates, it was found that a sense of scale was clearly imparted to participants as intended. All participants responded to the absolute scale of the environment and there were clear and significant differences between environment scale estimates in the Large and Small Environment conditions. The inaccuracy of participants' perception of distances in a large scale VR environment is consistent with the literature which shows that accurate perception of scales in immersive virtual reality does vary between peripersonal and extrapersonal space (Armbrüster et al., 2008). We found that participants tended to underestimate the scale of the large environment (Figure 4.12), likely due to the relative weakness of stereopsis and parallax as absolute depth cues beyond a few meters (Brenner & Van Damme, 1998). We also found no relationship between participants perception of scale as measured in the scale assessment task and spatial precision as measured in the 2AFC task. This is an interesting and enlightening result which we will return to later.

As in Experiment 4.1 and previous chapters we observed a main effect of viewpoint on change-detection thresholds with larger viewpoint rotations corresponding to lower spatial precision. Once again, we interpret this as evidence of the coarser, less spatially precise nature of viewpoint-independent encoding.

Although we saw a significant main effect of scale when comparing the absolute change-detection thresholds, the average difference between thresholds corresponded closely to the factor of 30 used in the scale manipulation. When thresholds were scaled by this factor and reanalysed, we found no significant effect of scale. Change-detection thresholds reliably scaled with the pictorial cues in the stimuli, which subtended the same visual angle when the environment was room-sized to when it was amphitheatre-sized, even though participants perceptions of absolute scale were highly varied as noted above.

While non-significant, the interaction between viewpoint and scale was approaching significance. This could imply that perhaps there are differences in how scale is encoded in viewpoint-dependent and -independent tasks but that they are much more subtle that expected, and the current task is not suited to measure them. Critically, the direction of this interaction would suggest that (if anything) larger environments may be encoded with *poorer* spatial memory precision in the viewpoint-dependent task, the opposite direction to the one we anticipated in the introduction. This might then indicate that viewpoint-independent resources are sensitive to changes in absolute scale whereas viewpoint-dependent resources are not, however, this result would not be immediately compatible with a representation in which locations are coded in a reference frame of fixed resolution.

There are several possible reasons for these results. One explanation that we must entertain is the possibility that this task simply does not use allocentric information at all. It is possible to explain the drop in precision due to viewpoint as associated with some internal egocentric rotation process. If this were the case, we might expect a larger environment, and therefore further imagined movement to have a larger impact on spatial precision, as the data showed. However, there are strong reasons not to believe this, beginning with the large body of spatial memory research that show links between the novel view paradigm and allocentric representation and processing (Burgess et al., 2004; Holmes & Sholl, 2005; King et al., 2002; Lambrey et al., 2012; Motes et al., 2006; Negen, Roome, et al., 2018; Schmidt et al., 2007; Waller & Hodgson, 2006). Secondly, in the previous chapter when systematically increasing view rotation, precision did not degrade linearly as might be expected if the drop in precision was associated wholly with egocentric view rotation processes. Finally, we see significant (Experiment 4.1) or very close to significant differences (Experiment 4.2) for interactions between viewpoint and scale when comparing change-detection thresholds, and in some response times (Experiment 4.1). This observation indicates that scale is affecting performance in the two viewpoint conditions differently and that allocentric representation is involved in the task, but that out initial predictions were overly simplistic in treating ego-and allocentric constraints on precision as independent. The results suggest that this is not tenable, and that their interactions affect precision in a more subtle way and not in the direction we had anticipated.

A more likely explanation of similar impact of scale on viewpoint-dependent and -independent thresholds, is that the scale affected an initial egocentric encoding which due to the order of processing, led to constraints on the allocentric representation. This is because the viewpoint-independent spatial precision as measured in this task is inherited from an egocentric perception and initial spatial egocentric encoding. The precision of this initial egocentric representation which is subject to Weber's law and the scale of the visual image rather than the absolute scale of the environment. It is *this* information which is then transformed into an allocentric code. While allocentric mechanisms may not themselves be subject to principles such as Weber's law, the information received by allocentric systems has already been shaped by systems that are, and so this constraint is still present, at least at the scales tested in the current experiments.

This interpretation is also supported by the evidence from the Scale Assessment Task. While it was clear that the manipulation of scale had significant impact on participant's perception it was found that participants were able to complete the task with varying degrees of success, with the poorest performance underestimated the scale of the large environment by a factor of 12.1. This scaling inaccuracy was not correlated with change-detection thresholds which nonetheless scaled with the environment. This suggests that participant's perception of the absolute scale of the environment did not directly determine their precision. What remained constant across both conditions were the extent and relative locations of objects and environment in the participant's field of view and it seems that this (as well as viewpoint rotation) was the basis for participant's spatial memory precision.

With these points in mind, if allocentric representations are always abstracted from initial egocentric representations, it is debatable whether it is possible to measure a process pure limit of allocentric precision, untainted by the influence of egocentric encoding, matching or response processes. Some implications of this will be discussed below.

In summary, this experiment found that spatial change-detection thresholds changed in line with manipulations affecting the perception of the scale of the environment. However, the effect due to scale was very close to the overall environmental scaling factor suggesting that across scales, precision is limited principally by pictorial cues. In this experiment, spatial precision for objects preserves the same angular extent in the observer's field of view regardless of the perceived scale of the environment. This effect was seen in both viewpoint-dependent and -independent tasks and is interpreted as a consequence of constraints on the initial egocentric representation formed at encoding.

4.5 General discussion

This pair of experiments has provided an interesting insight into the quality of the information available in the spatial change detection viewpoint rotation task and poses some interesting questions about the quality of allocentric spatial memory. Before running the experiments, it seemed plausible that spatial change detection thresholds would be constrained by the fixed 'resolution' of a reference frame, anchored to the environment as seen the firing of cells in the hippocampal formation. Instead, it appears that spatial memory precision in both viewpoint-dependent and viewpoint-independent tasks changes in a way that is in line with Weber's law, with thresholds decreasing relative to the size of the array.

The results from Experiment 4.2 showed that the significant effect of scale in Experiment 4.1 was likely not due to the absolute scale of the remembered environment but rather the scale of the stimulus in the field of view. In both experiments, the level of precision in the viewpoint-independent task was lower than that of in the viewpoint-dependent task but this decrease was at a fixed offset relative to the level of precision measured in the viewpoint-dependent condition.

4.5.1 Applying findings to models of spatial memory.

While somewhat counterintuitive due to the nature of allocentric information being anchored to the environment, the results are broadly consistent with Bicanski and Burgess' model of spatial memory (Bicanski & Burgess, 2018). This model conceptualises the allocentric representation as a conversion of the sensorily based egocentric representation within the 'parietal window'. Before information becomes allocentric and abstracted from the viewpoint, it is first subject to any constraints as egocentric information. This includes Weber's law, the effect of the shape and size of the array and also any ceiling or floor levels of precision derived from the egocentric spatial encoding. This view of the role of information transfer and interdependency of egocentric and allocentric representations in determining the limits of spatial memory is an important overarching theme across experiments and findings in this thesis. Our data show that the precision of spatial memory as measured in a change detection task is relative to the size of the array within the field of view and not the absolute scale. This was the case in both a viewpoint-dependent and viewpoint-independent task and provides useful new constraints on computational models.

4.6 Conclusion

This chapter investigated the effect of scale on spatial memory precision in both viewpoint-dependent and -independent spatial memory tasks and resulted in 3 main findings. Firstly, I found that increasing the scale of the array increased spatial change-detection thresholds in both viewpoint-dependent and -independent tasks. Secondly, I found that the effect of scale on change-detection thresholds was likely not due to the absolute scale of the remembered space but instead due to the scale of the stimulus in the observer's field of view (an egocentric, rather than allocentric property). Finally, we can infer from this that the precision of viewpoint-independent spatial memory is not a fixed absolute value as we have might expected based on a simplistic view of the general characteristics of allocentric neural representations, but that it is – at least in the current studies – scale invariant, perhaps tracking pictorial cues that are preserved across scale conditions in the current experiments. I interpret this result to be evidence of the interdependence of ego- and allocentric memory systems and egocentric perception in these tasks. While allocentric representation may not be directly be constrained by factors – such as Weber's law – that effect egocentric representation, it is subject to the quality of the information it receives.

Chapter 5. Measuring capacity in spatial memory

5.1 General Introduction

So, far this thesis has investigated the limits of spatial memory in terms of precision: in Chapter 3 we modelled the way the rotation of viewpoint between presentation and testing affected spatial memory change detection thresholds and used the function relating viewpoint rotation to precision to gain insight into how underlying ego- and allocentric processes interact. Chapter 4 focussed on the question of whether the scale of the stimulus affected change detection thresholds and tested whether the precision of viewpoint-independent spatial memory corresponded to a fixed absolute distance in the environment. In this chapter we will investigate the limits of spatial memory in terms of capacity. We will ask if egocentric and allocentric representations have strict capacity constraints as has been seen in Working Memory, whether there is a trade-off between capacity and precision and if so, whether there are distinct resources that are available to ego- and allocentric spatial memory.

5.1.1 Capacity of visuospatial Working Memory

The term Working Memory describes the ability to hold and manipulate information required for ongoing cognition and Working Memory research has covered a wide range of different tasks across a range of sensory modalities from verbal Working Memory and digit spans to visuospatial patterns, colours and orientations (Baddeley, 2000, 2007). Visuo-spatial Working Memory (VSWM) refers to a dissociable sub-component of Working Memory related to tasks concerning visual features and spatial locations (Brady et al., 2011). These tasks generally use 2D, view-dependent stimuli and require participants to remember visuospatial elements over a short delay before being tested, for example, recalling the specific colour, rotational orientation or location of a cued item from a larger set. Changing viewpoints and 3D space are typically not included in VSWM tasks and accordingly much research in this area has not drawn a distinction between ego- and allocentric contributions. However, VSWM tasks are very relevant to the current thesis as the viewpoint-dependent (0° view rotation) condition seen in Chapters 2, 3 and 4 is very similar to classic fixed viewpoint 2D VSWM tasks; it requires participants to remember the spatial location of an array of objects from a fixed position and then tests participant memory after a brief delay. Although the thesis began with a consideration of the spatial characteristics of representations in more general memory and behaviour, properties of Working Memory are particularly relevant to their capacity limits in the current tasks.

A defining characteristic of Working Memory is a limited in terms of capacity; the amount of information that it can hold (Baddeley, 1992). Miller's (1956) seminal work suggested that Working Memory had an item-based limit of 'seven, plus or minus 2'. Although the number of items is now thought to be around 4 (Luck & Vogel, 1997), the notion of an item-based limit has remained a prominent feature of the literature and has been extremely influential in shaping some current models which describe Working Memory capacity being based around 4 discrete 'slots' regardless of modality (Cowan, 2001; Cowan et al., 2005; Rouder et al., 2008).

In this view, each 'slot' can hold a single features or element such as digits, words, colours etc. or related chunks of information such as, word pairs, complex constellations of chess pieces or numerical sequences (McNamara & Scott, 2001; Rouder et al., 2011). For example, whether participants are asked to recall lists of either single words random previously learned word pairs, they are able to recall 3-5 items despite the learned word pairs containing twice as many pieces of information (Cowan, 2001). While strategies such as rehearsal or mnemonics be deployed to use the available capacity more efficiently, methods can be devised to disrupt these strategies and, it is argued, reveal the true limits of Working Memory. These methods often employ dual tasks, such as participants repeating a word or listening to distracting audio during encoding or ignoring a spatial distractor cue (Gao & Theeuwes, 2020), prevent the use of grouping and rehearsal strategies and allows experimenters to reliably observe the limit of 3-5 items (i.e., complex span tasks Schmiedek et al., 2009).

An alternative model describes the capacity limit of Working Memory as 'resource based' (Alvarez & Cavanagh, 2004; Bays, Catalao, et al., 2009; Bays & Husain, 2008; Ma & Wilken, 2004). Instead of a limit in terms of a fixed number of items, the resource based model describes a system that is informationally limited with resources being shared between items in Working Memory. The model explains memory deficits due to set size in terms of a trade-off between capacity and precision, a larger number of items means a more thinly distributed resource, resulting in less precision for each item. For example, Bays et al., (2009) asked participants to estimate the movement direction of a cued object and participant accuracy declined as set size increased. Similarly, when participants are asked to recall the colour of cued blocks from larger sets, increasing set size results in a poorer, less precise estimate of colour (Ma & Wilken, 2004). This trade-off between capacity and precision also works the other way and it has been observed that Working Memory capacity is reduced for more complex items (Alvarez & Cavanagh, 2004).

The resource model interpretation of the precision-capacity trade off remains contentious, and some authors have interpreted similar results in terms of discrete slot models (i.e., see Awh et al., 2007 as reviewed in Brady et al., 2011). The present study does not aim to answer the difficult question of whether capacity limits in VSWM are due to discrete slots or the flexible allocation of resources, and in either case, we would expect these limits to have a similar effect on spatial change detection thresholds. In a slot-based model, in trials which probe an item which happens to have been encoded within a 'slot' participants would answer correctly, but this would occur on a diminishing proportion of trials as the set size increased resulting in a higher threshold. Alternatively, within a resource-based model in which a finite resource was shared between capacity and precision we would also see a decline in the precision of spatial memory as set size increases due to the lower fidelity representation of each item.

In sum, VSWM is limited in terms of its capacity. This limit can be described in terms of a fixed number of items that when exceeded, memory becomes much poorer. It can also be described in terms of a limited informational resource that is shared between the number of items and the precision at which they are remembered. In either case we would expect spatial change detection thresholds to increase with larger set sizes. The view-dependent condition in the spatial change detection task used throughout this thesis bears many similarities to VSWM tasks. With this in mind, we might expect performance a view-dependent spatial memory task to also be severely limited in terms of capacity.

5.1.2 Separate allocentric and egocentric resources

However, a recent study has provided evidence that egocentric and allocentric spatial memory resources are distinct in such tasks. Aagten-Murphy and Bays (2019) asked participants to learn 2D spatial configurations of 1, 3 or 4 array objects in a condition with a (separate) stable visual landmark present either being present (encouraging an allocentric strategy in which the relative locations of objects and landmarks contribute to performance) or absent (demanding an egocentric strategy in which each object's location is encoded relative to the observer). They found that participants were better able to remember object locations that were nearer the stable landmark, thus apparently making use of the allocentric cue of the distance and direction between the landmark and the object. Importantly, the precision for objects that were *not* in the vicinity of the landmark did not benefit when a landmark was present. This means that although the location of the landmark was encoded within memory, it did not seem to take up additional resources. This suggests that the stable landmark did not take up a 'slot' or drain the same informational resources that were being used by the to-be-remembered objects (as this would have presumably impacted memory accuracy for all items). These findings imply separate memory resources for egocentric and allocentric spatial memory.

It may be that VSWM uses both ego- and allocentric memory as each may fulfil different functions beyond the lab setting and that when these discrete processes are used, they rely on separate resources. There is evidence to show that when participants experience a change in viewpoint between encoding and retrieval of spatial information, even over brief periods, the hippocampal formation and the same allocentric machinery is activated. For example, Schmidt et al (2007) measured fMRI BOLD signal while asking participants to remember the location of a lamp in a 3D virtual environment. After a brief delay and a viewpoint rotation of between 0 and 180° participants were asked whether the lamp had moved in a same/different style design. They found that increased view rotation and successful completion of the task correlated with increased hippocampal activation. Even over brief periods, viewpoint-independent spatial memory relies on specifically allocentric encoding of space within the hippocampal formation which is distinct from the neural regions and resources involved in fixed view 2D VSWM (Baker et al., 2018; Schöning et al., 2007; Suchan et al., 2006; Zago & Tzourio-Mazoyer, 2002).

While VSWM has generally studied spatial memory capacity in a sparse, 2D unnatural context, spatial memory is applied in a diverse range of tasks; following an almost forgotten route, recalling events over a lifetime, and it is required to remember the locations of countless densely situated objects around the house. Indeed, there are other forms of memory that can potentially contribute to spatial tasks, some with potentially far greater capacity than seen in typical VSWM studies. In the following sections we will review evidence from a number of sources to understand why it could be expected that that egocentric and allocentric spatial representations rely on distinct resources that have different capacities and how that might be shown in behaviour.

5.1.3 Visual recognition memory

Visual recognition is a core component of self-location and navigation (Aguirre, 1999). Experiments investigating long-term visual recognition memory often involve showing participant many images and then after a delay of days of weeks they are shown a selection of both novel and previously shown images, and they are asked whether they have seen them before. Konkle et al., (2010) used this paradigm applied to spatial scenes, showing participants 2,912 images from 160 difference scene categories. Participants had an excellent ability to differentiate familiar and unfamiliar images and performed with accuracy of between 96% and 76% correct depending on the overlap of scene categories (i.e., how similar the images were). Similar results, showing impressively large recognition memory capacity have been shown for sparse and complex images, objects and objects embedded within scenes (Brady et al., 2008; Brady & Oliva, 2008; Hollingworth, 2006; Konkle et al., 2010a; Shepard, 1967; Standing, 1973). These results show that, in principle, when a spatial scene is unchanged between presentation and testing, performance can be supported by a form of visual memory that has a very large capacity and can endure a long period of time.

These tasks bear both similarities and differences to elements of the spatial change detection task seen in this thesis. While there are well documented conceptual differences between the feeling of *familiarity* needed for deciding whether one has seen a single image before (required for visual recognition tasks) compared with the involved and active process of holding, manipulating and comparing the location of multiple objects (required to complete the spatial change detection task and many VSWM tasks; Yonelinas, 2002)- there are also similarities. For example, one well known egocentric spatial cue and strategy for self-location is 'visual pattern matching', simply matching one's current view to a previous and familiar view, perhaps exploiting a form of visual recognition memory similar to that seen in the above studies. Visual pattern matching is seen in animals and is used in such complex spatial tasks as reorientation and route following. (Cheung et al., 2008; Sheynikhovich et al., 2009; Stürzl et al., 2008; Wystrach & Graham, 2012). While research with human navigators indicates the use of distance and geometry of environmental features as in addition to the 2D image to match spatial features (Lee et al., 2012; Spiers et al., 2001), something akin to visual pattern matching also likely provides spatial information in 2D viewpoint-dependent spatial tasks (Spiers et al., 2001).

5.1.4 Episodic and allocentric spatial memory

Episodic memory describes memory for 'personal events within a spatiotemporal context' (Burgess, Maguire, et al., 2002) and provides another example in which memories with a spatial component show an extremely high capacity.

The spatial context of episodic recollection and imagination has been theorised to involve specifically allocentric spatial memory due to shared neural bases within, and shared activation of, the medial temporal lobes and hippocampus (O'Keefe & Nadel, 1978; Vann et al., 2009; Vargha-Khadem, 1997). This link has been shown experimentally with hippocampal patients showing severe deficits in episodic memory tasks as well as viewpoint rotation and topographical spatial memory but not recognition memory or spatial tasks which can be solved with purely egocentric cues (Holdstock et al., 2000; King et al., 2002; Spiers et al., 2001). Furthermore, place responsive spatial cells (in humans) within the MTL have been shown to activate when the episodic context of memories are retrieved even without navigation (Miller et al., 2013). While their precise role is debated, allocentric spatial representations are thought to be a key component in reconstructing the spatial context of events from episodic memory and also in imagination (Byrne et al., 2007; Hassabis et al., 2007; Lambrey et al., 2012; Mullally & Maguire, 2014; Smith & Mizumori, 2006).

Again, there are clear differences between referring to the passive 'capacity' of episodic memories which relates to how many discrete events one can recall accurately and referring to the capacity of VSWM which relates to how many items one can hold in active cognition (Baddeley & Hitch, 1974). Nevertheless, the shared neural machinery between episodic and allocentric spatial memory offers the possibility that there may be capacity limits for allocentric memory representations which are distinct from those seen in 2D view-dependent VSWM tasks.

5.1.5 Memory capacity and viewpoint rotation

Several studies (reviewed below) posit that computational differences in how egocentric and allocentric representations encode and update spatial locations should be evident in their interaction with different set sizes. This is based on the notion that 'Egocentric representations encode each object location with respect to the observer rather than with respect to an external reference frame…' and that ' …updating requires representing the change in the relationship between the observer and each object in the environment.' (Wang et al., 2006). As in an egocentric encoding item locations are updated separately and by piecemeal, a larger number of items should effect the response time and possibly the spatial precision on encoding and retrieval (for example, shown in the case of path integration, Klatzky et al., 1999). Conversely, 'Allocentric representations encode object locations in an external reference frame, and thus remain relatively stable as the observer moves…' '…updating involves computing changes in the observer's position and orientation with respect to a constant external reference frame.' (Wang et al., 2006). As allocentric representations are thought to update holistically and reorientate as one (Amorim & Stucchi, 1997; Klatzky et al., 1999), it has been theorised that the number of items stored should have less impact on the precision and response time involved in encoding and retrieving allocentric information.

However, superficially similar studies, exploring the relationship between (viewpoint-dependent and -independent) spatial memory precision and set size have provided very conflicting results. In some cases, when participants are asked to replace objects from familiar or novel viewpoints a significant negative effect on accuracy can be seen as a result of increasing set size from 1 to just 3 items (Wang et al., 2006). Taken on its own, this would indicate that both viewpoint-dependent and -independent memory have very limited set size capacity, similar to that seen in VSWM.

However, there is also evidence showing little or no effect of set size (Hodgson & Waller, 2006; Rieser & Rider, 1991). In a study by Hodgson and Waller (2006), when participants were asked to point to the locations of between 1 and 15 objects from within the array (while blindfolded and after rotation), there was no significant change in pointing accuracy due to an increase in set size, however, there was a significant effect on response time which increased linearly as set size increased. The authors argued that the absence of an effect on accuracy is due to the formation of a long-term memory representation.

In another study manipulating viewpoint and set size, King et al (King et al., 2002) tested a bilateral hippocampal patient, 'Jon' and healthy controls in a spatial viewpoint rotation task in which participants were required to remember the locations of between 3 and 7 items in a rotated viewpoint task and up to 10 items in a fixed viewpoint task. As set size increased beyond 1 item in the rotated view condition, Jon's performance quickly reached chance level. The healthy control group also performed somewhat less well overall in the rotated view condition but between set sizes of 3 and 7 items, there was little change in their performance. Jon was able to store multiple object locations in a viewpoint-dependent condition but unable to store multiple locations in a viewpoint-independent condition which is compelling evidence that there are different capacity constraints for these two systems. The authors suggested that the viewpoint rotated condition utilised, the allocentric system which can store more locations or more locations for longer due to its ability to store a single amalgamated representation of the environment rather than sensory snapshots of each object. This study provides evidence of a case in which, due to neurological damage, the capacity for allocentric encoding was at floor while the capacity for egocentric encoding was at ceiling. This dissociation is further evidence of separate and distinct resources responsible for viewpoint-dependent and -independent tasks.

The 4 Mountains test is another example of a task that uses viewpoint rotation to behaviourally bias viewpoint-independent memory (Hartley et al., 2007). In the task participants view a scene and then are asked to choose the same spatial scene from three similar alternatives. This test of topographical short-term memory has been shown to differentiate patients with mild cognitive impairment (MCI) from those with mild Alzheimer's disease. This is thought to be due to the early pathology of Alzheimer's affecting the hippocampus and specifically viewpoint-independent memory. For the current task, this is further evidence of separable and dissociable memory performance which could indicate reliance on separate memory resources.

The conflicting evidence on spatial memory and capacity may be an indication of the task dependent nature of capacity limits in memory. As we see in VSWM, different types of task can yield very different results due to a participants ability to economise and chunk information. While the above viewpoint rotation tasks all investigate the effect of set size on spatial memory precision, there are important task differences which may have affected their sensitivity to capacity limits. These differences relate to the nature and sensitivity of response measures, the availability of idiothetic and other spatial cues, and the structure and content of stimuli that can afford different opportunities for chunking or other strategies to conserve or manage memory resources.

To summarise, studying the capacity of memory has been an important and defining method in memory research. Studies have clearly established that visuo-spatial Working Memory, at least superficially similar to the type of memory under investigation in this thesis, has limited capacity. However, there is evidence to suggest that viewpoint-dependent and -independent memory may be reliant on separate resources which could mean that their capacity limits are different. The common dependence of allocentric spatial and episodic memory on the hippocampal formation (Nakazawa et al., 2004; O'Keefe & Nadel, 1978; Vargha-Khadem, 1997) and the central role that allocentric spatial representations appear to play in memory reconstruction and imagination (Hassabis et al., 2007; Hassabis & Maguire, 2009; Moscovitch et al., 2006; Mullally & Maguire, 2014) might lead us to believe that allocentric spatial memory capacity is much larger than the strict limits seen in Working Memory.

So far, studies explicitly investigating the effect of set size on spatial memory in viewpoint-dependent and -independent tasks have shown conflicting results and there remains considerable doubt over whether and how viewpoint changes influence the capacity of spatial memory. This provides a compelling rationale to investigate capacity limits of spatial memory using tasks similar to those developed in previous chapters. In two experiments, we will investigate the limits of capacity in viewpoint-dependent and -independent spatial memory using two distinct methods- change detection and object replacement.

5.2 Experiment 5.1

5.2.1 Introduction

When set size is varied in a viewpoint-dependent memory we might expect to see a drop in spatial memory precision due to the same system limits as seen in a variety of other visuospatial Working Memory tasks (Brady et al., 2011) which show that the precision with which we retain information about visuospatial stimuli gets poorer as the number of items/locations to be remembered increases. However, it is unclear whether this limitation will also be observed in viewpoint-independent spatial memory which has been found to be distinct from visual Working Memory processes in terms of its neural bases and processing requirements (Aagten-Murphy & Bays, 2019; Hartley et al., 2007; Schmidt et al., 2007; Suchan et al., 2006). With this uncertainty in mind, in Experiment 5.1 we intend to compare spatial change detection thresholds for varying set sizes in viewpoint-dependent and -independent tasks.

When initially developing the psychophysical spatial change detection task (Chapter 2) we piloted a version of the task designed to measure capacity. In this version of the task, rather than the stimulus-change being spatial distortion as in all other subsequent experiments (Chapters 3-4), the adaptive staircase was used to adjust the set size of the array. It was initially assumed that increasing set size would make spatial changes of fixed magnitude more difficult to detect, so that set size could be substituted for stimulus intensity in a psychophysical design. However, in these pilot studies we found that the psychometric response distribution gathered from participants were disorderly, often non-monotonic and were poorly fit by psychophysical curves, especially in the viewpoint rotation conditions. This was an interesting and unexpected result. One possible reason for the unusual pattern of results is that increased set size does not have a reliably negative effect on spatial change detection. In some circumstances it may be easier to notice spatial changes when there are more objects present to provide context, for example providing more spatial cues or more possibilities for chunking. There is also a larger proportion of the array that remains unchanged after the movement of a single object.

These interim conclusions were useful in guiding the direction of the first experiment in this chapter which opts to measure the spatial change detection thresholds for different set sizes separately. Rather than a single adaptive staircase varying set size (as in Chapter 2), each set size condition will be assessed with its own adaptive staircase controlling trial difficulty and resulting in separate change detection threshold estimates. This approach allows us to investigate the relationship between capacity and precision and determine whether different capacity limits affect view-dependent and view-independent memory and whether in each case, precision is affected by set size consistent with a fixed resource model (Brady & Alvarez, 2015; Cowan, 2001). If either form of spatial memory is constrained by a set size or informational resources shared between capacity and precision, the change detection threshold should be higher for larger sets as either objects will not be encoded within the memory (such that even very large changes cannot be reliably detected) or they will be encoded with reduced spatial precision.

5.2.1.1 Experiment 5.1: Hypotheses

We expect to see a drop in precision associated with the effect of viewpoint rotation as has been seen in other experiments through this thesis. We will investigate whether there is also an effect of set-size on change-detection thresholds. Finally, we will investigate whether there will be an interaction between viewpoint and set size. If so, this would provide evidence of distinct resources for egocentric and allocentric spatial memory.

5.2.2 Methods

5.2.2.1 Overview

In Experiment 5.1, as in experiments presented in Chapters 3 and 4, participants were asked to make 2AFC judgments about which object in an array had been moved between presentation and testing. The viewpoint angle change between presentation was manipulated. The array set size i.e., the number of items in the array, was also varied to investigate possible limits on the capacity of spatial memory. This experiment used an online (web-based) task because face-to-face laboratory testing was not possible during the COVID-19 pandemic.

5.2.2.2 Participants

In total, 8 were recruited via the University of York's Department of Psychology participant recruitment system, and 10 responded to advertisements on social media. Both connected to the online test by following a link. In total, 8 participants were excluded according to criteria detailed below, leaving 10 remaining participants (6 female, mean age = 20.7 SD = 4.3).

5.2.2.3 Task and Materials

The task was a web application version of the same virtual task used in previous chapters. The application ran on all web browsers and participants were asked to run the experiment in full screen. As In previous versions, task and stimuli were developed in Unity3d Game engine in C#.

The virtual environment consisted of a 10m diameter, circular space with a mountainous skybox providing distal orientation cues at the horizon (no distance cues). Participants viewed arrays of objects in a random configuration (see Chapter 2) in the centre of the arena, placed onto a floor plane with a striped noise material.

Similarly, to experiments in previous chapters, Experiment 5.1 uses an adaptive staircase paired with a 2AFC task to estimate participants' spatial change detection thresholds. On each trial (see Figure 5.1), at presentation, participants saw an array of objects (sampled without replacement from a pool of 10 simple 3D geometric shapes with contrasting colours) arranged in a random configuration for 6 seconds (‘Sample Scene’). The screen dimmed over 0.5 seconds, remained black for 1 second before fading back over 0.5 seconds. The initial viewpoint during the Sample scene was fixed across all trials. At Testing, the participant viewed a new scene (‘Test Scene’) either from the same viewpoint (a 0° viewpoint rotation) or a different viewpoint (viewpoint rotated 135° either clockwise or anticlockwise at random). Two objects were highlighted with a glowing yellow border and participants were immediately prompted to select the object they thought had moved, by clicking it with the mouse. Every 30 trials participants were offered a break of an undetermined length.

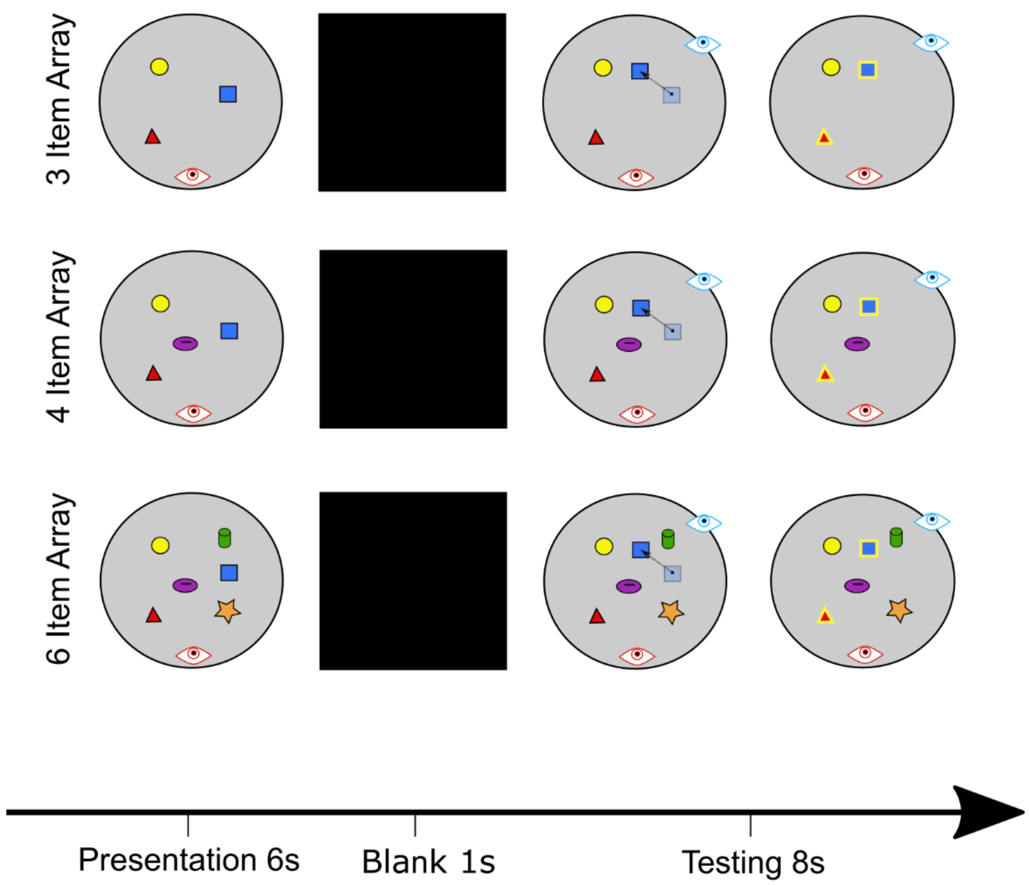

Figure 5.1 Experiment 5.1: Task and Materials . The above shows the trial procedure for the spatial change detection paradigm as used in Experiment 5.1. In the ‘Presentation’ phase, participants are shown an array of either 3, 4 or 6 Items depending on the condition of that trial. One object was moved while the screen is blank and participants were moved around the array by 0 or 135 degrees (denoted by the blue and red eye symbol). In the ‘Testing’ phase, the target object and a randomly chosen foil are highlighted and the participant much choose which they think has been moved by clicking with the mouse.

5.2.2.4 Set Sizes

Set sizes of 3, 4, and 6 were chosen as being around the midpoint of set size limits found in the visuospatial Working Memory literature, spanning the critical range of 3-5 items (as in Bays et al., 2009). While studies have shown set size effects with both larger and smaller arrays, we chose this range as an intermediary and predicted that due to the adaptive difficulty of the task, performance will never reach a ceiling or floor and constraints on memory and processing will be evident. Three levels (3, 4, 6) were chosen so that if an effect was seen, we would gain additional insight into possible capacity constraints. As a minimum, 3 items is the smallest practical set size for a 2AFC task, it also provides enough intrinsic spatial information for an observer to determine which item has moved through spatial triangulation without information from outside the array. An array of 4 items offers comparison with other change detection experiments in this thesis. Practically, the creation of arrays with larger set sizes than 6 is prone to involve more extensive computation (to produce valid arrays where all objects are visible, non-overlapping and so on), another consideration was that largest set size of 6 should be manageable within the memory and processing resources available to participants’ browsers running on unknown hardware.

5.2.2.5 Procedure

The entire experiment took between approximately 1.5 and 2 hours. In order to maintain participants’ cooperation, avoid boredom, fatigue or inattention in the online task, the total time was divided into two separate sessions, each between 40 minutes and 1 hour long. Participants were asked to undertake both sessions within a week of each other. Before session 1, participants completed an interactive tutorial which explained the task and allowed them to complete 20 example trials. Each session used exactly the same structure and content (the same adaptive staircase procedures and interleaved randomized trials of varying viewpoint and set size operated in the same way).

Each participant completed 4 separate conditions in a 2x2 factorial design (Viewpoint change: 0° or 135° x Set Size: 3, 4 or 6 items). A separate adaptive staircase modified the degree of spatial distortion applied to the array for each condition (6 total). Trials were presented in a random sequence, with the chosen condition for each successive trial being randomly drawn from a list of incomplete staircases. An opportunity for a break of indeterminate length was provided every 30 trials.

Each staircase began at a distortion level equivalent to the highest level to spatial distortion in our previous experiment (1m) and the testing range matched that used in Chapter 4 (see Appendix F). Again, each adaptive staircase followed a 3 down 1 up procedure (in which participants must answer

3 consecutively correct responses for spatial change to decrease and 1 incorrect response for spatial change to increase). Each staircase completed when it had reached 14 reversals (see Chapter 2.3.2). The session finished when all staircases were completed or the duration exceeded 1 hour.

5.2.2.6 Participant Exclusion Criteria

Participant performance was assessed dynamically, during the first session of the main experiment. When a participant had completed 1 reversal of each staircase, their proportion of correct responses was calculated. If a participant's average proportion correct was below 60% on the view change staircases, the experiment was cut short, and the participant was excluded. This cut-off criterion was based on information gained from poor-fitting data in previous experiments, in which it was observed that in cases where the proportion of correct responses in early view rotation trials was below 60%, there was often difficulty in fitting psychophysical curves to the data such that it would have to be excluded in analysis. One participant was excluded on this basis.

5.2.2.7 Data Exclusion Criteria

If a spatial change detection threshold was estimated outside of the testing range (below -18 dB or above 14 dB, Appendix F) this indicated that we were unable to effectively model a participant's performance in that condition and all data associated with this participant was excluded. Seven participants failed to meet these criteria and we were unable to effectively estimate their threshold in the 135° view conditions with 3 or 6 objects. This relatively high rate of exclusions compared to previous experiments is likely due to an increased noise in the data, and attributable to the online format in which, unlike in the laboratory-based experiments, it was impossible to fully control participants' immediate environment and ensure that they were free of distractions.

5.2.2.8 Analysis

General data analysis procedures for estimating threshold levels of performance with psychophysical curves were as detailed in Chapter 2. To assess the general fit of the psychophysical function to the data from each condition, we calculated individual Bayes Factors for each psychometric curve, comparing the fit of the psychophysical curve to a null hypothesis (i.e., the likelihood of the data under a uniform model where the proportion of correct responses, the mean across all trials, is unaffected by the degree of spatial distortion).

To investigate the effects of viewpoint changes and array set sizes on change detection thresholds, a 3x2 (set size x viewpoint change) repeated measures factorial ANOVA was used, and median response times from each condition were also compared, again using a 3 x 2 repeated measures factorial ANOVA.

5.2.3 Results

5.2.3.1 Psychometric curve fitting

A logistic function was fit to participant response data based on an average of 320 trials per participant. A separate function was fit to each separate staircase for each participant (6 conditions per participant). In all cases the fitted curve described the data better than an alternative model suggested by the null hypothesis (i.e., the data is better described by a psychophysical model than mean performance). Despite all but one Bayes Factors giving evidence for H1 (based on interpretations from Kass and Rafferty (1995), see Table 5.1) this data is clearly much noisier than that seen using similar methods in previous chapters.

Table 5.1 Frequency of Bayes factor statistics for each participant and condition of Experiment 5.1 . A Bayes Factor was calculated comparing the likelihood of the data as described by a psychometric curve (H1) and the likelihood of the data as described by a null hypothesis, a constant 'at mean' performance (H0). Bayes Factor interpretations from Kass and Rafferty (1995).

Bayes Factor Classification	Frequency
no evidence	1
Anecdotal evidence for H1	12
Moderate evidence for H1	17
Strong evidence for H1	7
Very strong evidence for H1	11
Extreme evidence for H1	12

5.2.3.2 Spatial Precision Thresholds: Effects of View Rotation and Set Size

In terms of main effects, our first hypothesis predicted that as viewpoint rotation increased spatial precision would decrease. A factorial repeated measures ANOVA showed that there was a significant main effect of viewpoint with, as expected, larger view rotations resulting in significantly higher spatial change detection thresholds ($F(1, 9) = 87.93$, $p<0.001$). We also set out to investigate whether larger set sizes would also be associated with a decrease in spatial precision. The factorial ANOVA showed no significant main effect associated with set size ($F2,18) = 0.18$, $p=0.85$) and there was no significant interaction between viewpoint change and set size $F(2, 18) = 1.18$, $p = 0.33$).

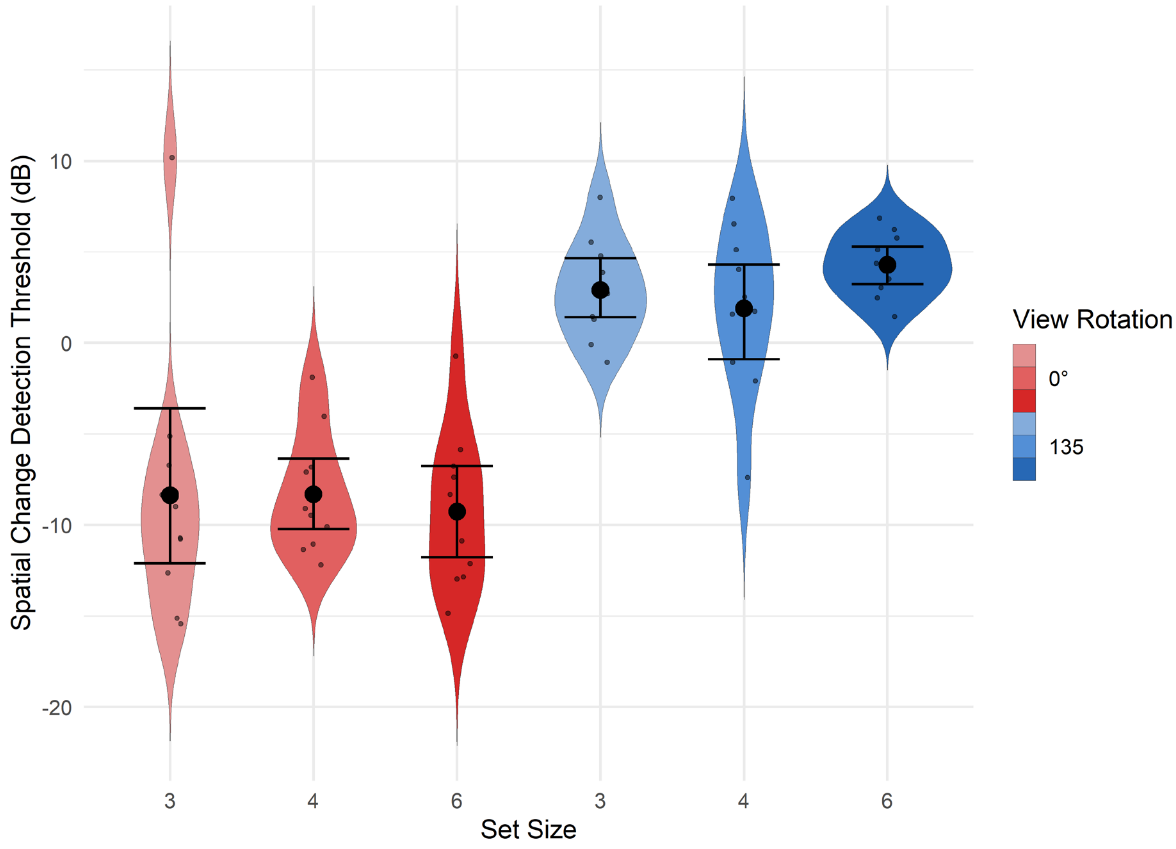

Figure 5.2 Experiment 5.1, mean estimated thresholds in each condition . Conditions are divided by set sizes on the x axis and plot colour distinguishes viewpoint. The red violin distributions show the mean estimated thresholds in the 0° rotation conditions and the blue distributions shows the mean estimated thresholds in the 135° conditions. Black points show individual participant thresholds. Larger black markers show group means with error bars showing standard error.

5.2.3.3 Psychometric Slope Parameter: Effects of View Rotation and Set Size

Comparing the psychometric slope parameters from different conditions with a factorial repeated measures ANOVA, we find no significant effect of viewpoint ($F(1,9) = 0.23$, $p = 0.64$), and no significant effect of set size ($F(2,18) = 0.74$, $p = 0.49$). However, there was a narrowly significant interaction between viewpoint and set size ($F(2,18) = 3.69$, $p < 0.05$). This appears to be driven by the 3-item, 135° condition.

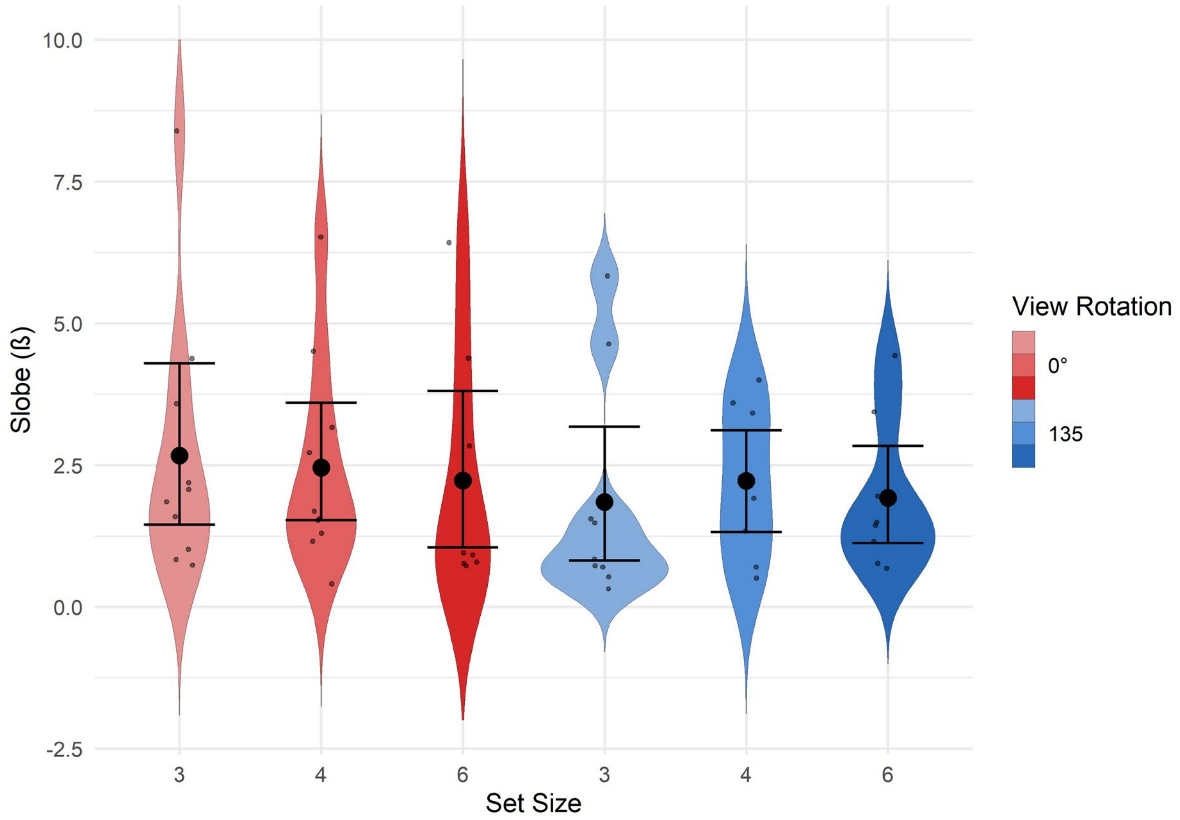

Figure 5.3 Comparing psychometric slope parameters as Weibull beta values between conditions . Black points represent individual slope parameters from each participant in each condition and larger black markers represent mean slope parameters in each condition. Each violine plot describes performance in a separate condition. Set size is demarked on the x axis and colour demarks view rotation conditions (red for 0°, blue for 135°). A smaller beta parameter represents a shallower slope.

5.2.3.4 Response time

Median response times were calculated for each participant and each condition, and a repeated measures factorial ANOVA was conducted to compare differences between conditions. There were significant effects of both viewpoint ($F(1, 9) = 37.38$, $p < 0.001$), and set size ($F(2, 18) = 21.88$, $p < 0.001$) on median response times and a significant interaction between viewpoint and set size ($F(2, 18) = 12.41$, $p < 0.001$). Numerically there appears to be a monotonic increase in response times with set size in the view rotated condition which is not apparent in the fixed view condition.

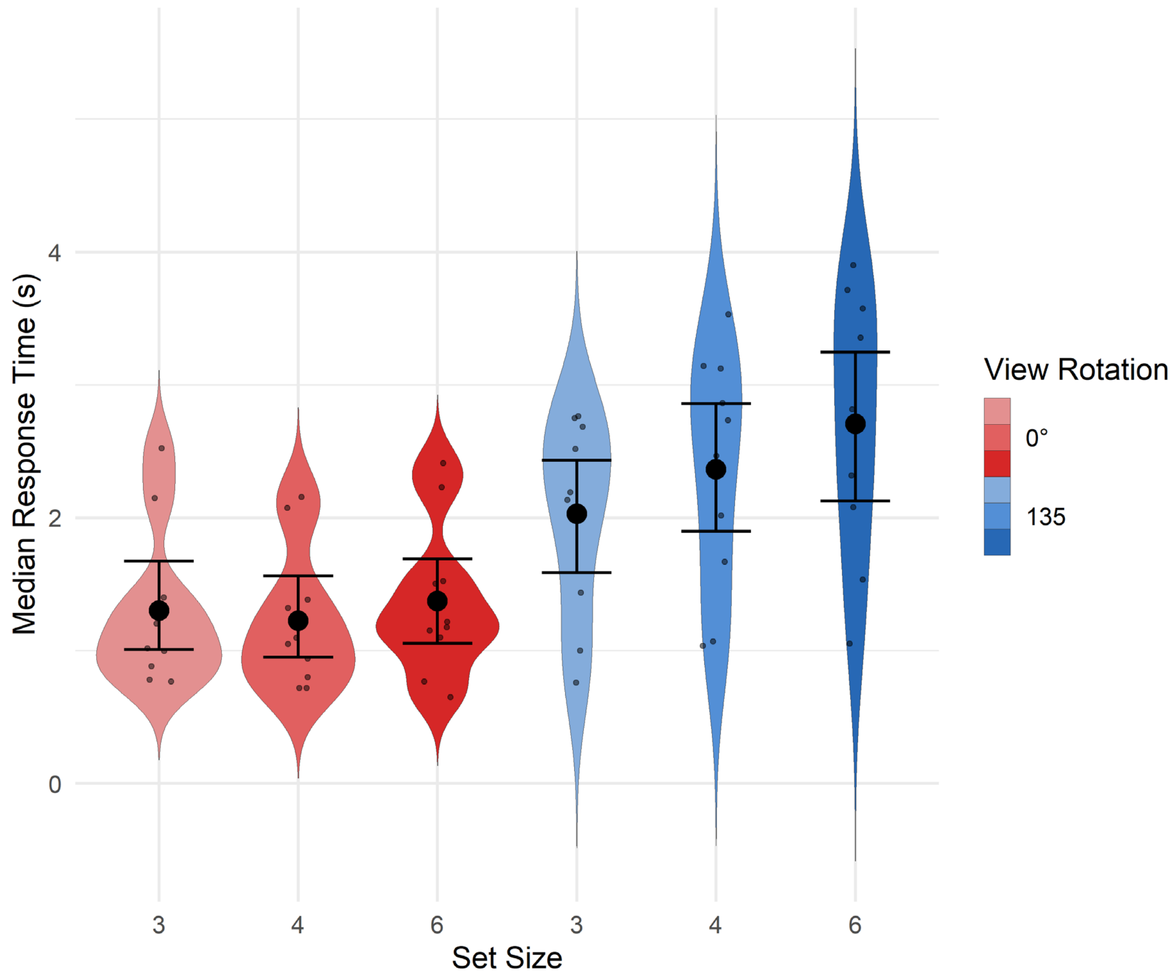

Figure 5.4 Experiment 5.1, median response times for each participant in each condition . Black points on each graph represent individual participant medians in seconds which is measured on the y axis. Violin plots show distribution of these points in each condition which are distinguished based on colour. Larger black markers show group means with error bars showing standard error. Red plots are 0° view shift conditions and blue show 135° conditions. Set sizes are separated on the x axis.

5.2.4 Discussion

The aim of Experiment 5.1 was to use psychophysical methods to estimate change-detection thresholds in a viewpoint-dependent and viewpoint-independent task. We varied the set size between 3 and 6 objects to investigate the possibility of distinct capacity limitations in viewpoint-dependent and viewpoint-independent contributions to spatial memory.

As predicted and consistent with results in previous chapters, we found a significant effect of viewpoint on threshold levels. This followed the same pattern as seen in all change detection experiments comparing 0° and 135° view rotations.

However, we found no significant main effect of set size, and no interaction between set size and viewpoint on change detection thresholds. This suggests that set size had no effect on spatial memory precision as measured by this task and it may indicate that no capacity limits were reached. That is, the largest set size we investigated (six items) may not have placed sufficient load on viewpoint-dependent or independent memory to tax the available resources. This could be due to two main reasons, firstly that this Experiment was not adequately powered to detect effects between 3-6 items (i.e., that this manipulation was too small) or (as described in Hodgson & Waller, 2006) it is possible that this task, in both viewpoint conditions, is exploiting some form of long term memory with a capacity limit greater than that of VSWM.

Regarding the possibility that our set size manipulation was too small, we can look to other studies that show effects of set size on spatial memory precision comparing arrays of as little as 1 and 2 objects. This effect was seen in tasks using a different paradigm and testing strategy to the psychophysical approach, requiring participants to recreate arrays by replacing objects in 3D space after walking through the environment (Wang et al., 2006). It is possible that this method of recording spatial memory accuracy is more sensitive than the current psychophysical approach using binary change detection decisions.

It is also possible that spatial updating of very small "arrays" (1-2 objects) represents a special case that does not benefit in the same way from allocentric memory; for example, recall that patient 'Jon', with bilateral hippocampal damage was able to encode single item arrays in the view rotated condition while approaching chance performance on multiple item arrays (King et al., 2002) and was able to do so with egocentric spatial updating. In Wang et al. (2006) participants walked through the environment to the novel viewpoint while objects were not visible. Therefore, they did not have an abrupt unexpected viewpoint change and may have been able to track the location of 1-2 objects with purely egocentric spatial cues without a significant loss in precision but was unable to use this strategy for larger sets.

While in the present study there was no effect on precision, analysis of response times did show an effect of set size and an interaction with viewpoint rotation. This is consistent with previous findings with similar set sizes showing little to no effect on precision but a significant effect on response time (Hodgson & Waller, 2006). Interestingly, the interaction shows that set size had a larger effect on response times in the viewpoint rotated condition, and it appears that they increase monotonically with set size. This observation suggests that time dependent processes related to the number of items in the array are specifically involved in the viewpoint rotated condition. These processes might include ego-allocentric conversion or comparing reconstructed and imagined scene elements from the novel viewpoint to the test stimulus. One such process would be analogous to mental rotation. In egocentric mental rotation larger viewpoint shifts are known to take longer (Shepard & Metzler, 1971), but here, of course, the degree of view rotation was constant, so the results are not compatible with a single coherent rotation of the entire array as might be expected in ego-allo transformations (Amorim & Stucchi, 1997; Bicanski & Burgess, 2018; Klatzky et al., 1999).

Another possibility is that in the current task it may take longer to resolve decision uncertainty for larger arrays. However, analysis of the slope parameters indicates that, if anything, there is greater decision uncertainty (indicated by a shallower psychometric slope) in the 3 item, 135° condition when fewer objects were present. Thus, it seems likely that the increased response times reflect some sort of serial/item-wise *comparison* or *matching* process which is specific to viewpoint-independent memory.

The online (browser-based) nature of the task meant that data were noticeably noisier than has been seen in comparable experiments in previous chapters. It was necessary to introduce strict exclusion criteria in order to retain confidence in the results. This is a useful gauge of how applicable this type of task is to a real life setting where the testing environment might not be as well controlled as the laboratory. The data clearly indicate that if one were to develop the methods for practical clinical or medical diagnostic applications, it would be important to make the procedure as efficient as possible to avoid inattention and fatigue. One way to do this could be to use alternative adaptive procedures (see e.g., Leek, 2001) which use maximum likelihood methods to select the most informative test parameters on each trial so as to estimate the logistic model as accurately and quickly as possible.

In sum, the results of Experiment 5.1 are in line with some earlier studies (Hodgson & Waller, 2006) although other studies had suggested that we should see an effect of set size on spatial memory precision (Wang et al., 2006). Perhaps this can be explained in terms of set size (for practical reasons, our study included a maximum set size of 6 which may be below the relevant capacity thresholds in our task) or task differences (specifically, our study relied on a binary decision on each trial whereas Wang et al. used tasks involving a single continuous DV). In Experiment 5.2 we use an alternative experimental approach which addresses these points and aims to improve on the efficiency of the 2AFC/staircase-threshold method.

5.3 Experiment 5.2

5.3.1 Introduction

In Experiment 5.2, again investigating effects of set size on the precision of viewpoint-independent and viewpoint-dependent spatial memory, participants were required to replace a moved object, with replacement error providing a direct measure their spatial memory precision on each trial (see Figure 5.5).

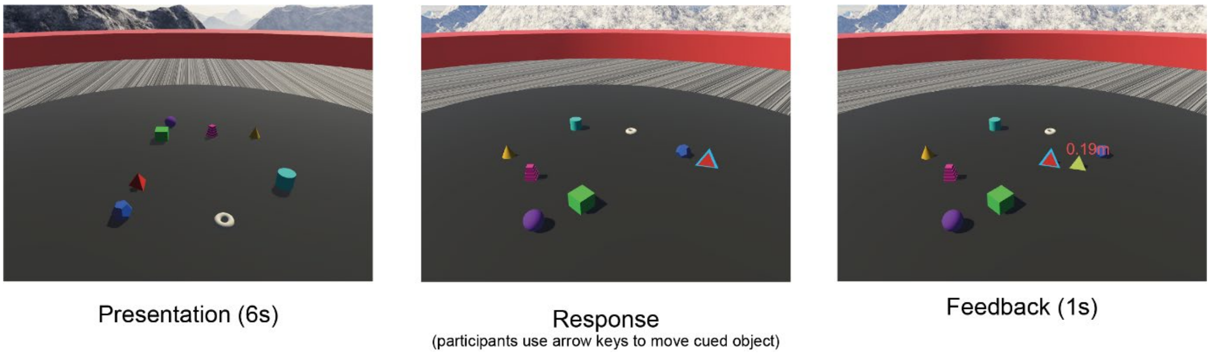

Figure 5.5 Experiment 5.2 task and materials as it appears to participants . Participants are shown a stimulus of an array of up to 8 objects. One object is moved between presentation and testing. During the response phase, the moved object is highlighted and the participant is asked to use the keyboard to move it back to its original position. Arrow keys are used to move the highlighted object respectively left, right, forward and backward relative to the current viewpoint. Participants are then briefly given feedback on their performance, indicating the original (target) location of the object (pale yellow shape) and the distance of the adjusted response as well as the distance between target and response (highlighted shape and text) from the target location.

'Continuous reproduction', also known as 'delayed estimation' involves testing participants with a single continuous response variable rather than a binary decision (2AFC, Same/Different) and has proved to be a useful technique in measuring performance in memory research. For example, in visual Working Memory, a number of studies (e.g., Ma & Wilken, 2004; Zhang & Luck, 2008) required participants to remember an array of coloured squares with a varying set size of between 2 and 8 items. At test, one item's location was highlighted, and participants were required to choose on a continuous colour-wheel the colour of the cued, target item. Accuracy on the colour wheel was used as a measure of memory precision. Similarly, Bays and Husain (2008) also used this method within a 2D spatial context in which participants responded with the cardinal direction of a moved target object.

There are several advantages to measuring precision directly as a continuous variable. As spatial locations are stored with a certain amount of error (which in previous experiments we measured in terms of change detection thresholds), continuous responses can in principle can give an indication of the *type of error* that is made. This is due to the relative richness of data compared to a binary response (Bays & Husain, 2008; Berens et al., 2020; Hartley et al., 2004; Zhang & Luck, 2008).

The benefit of a more descriptive response variable can also be applied to effects of viewpoint shift on spatial memory precision. In Experiment 5.1, and previous experiments in this thesis, we have seen a tendency for greater variation in participant responses to binary decisions after a large viewpoint rotation. This occasionally resulted in curve fitting procedures failing to converge or judged too inaccurate to be included in the results. This variability in binary responses might indicate a higher guess rate in viewpoint-independent tasks- that the participant has not encoded either cued item in spatial memory or does not recognise the array from the rotated viewpoint. By using a continuous response variable (in this case asking participants to replace the object in its original position) we may be able to see qualitative differences in responses to viewpoint-dependent and -independent tasks, and this could shed light on the nature of errors. For example, if there are 2 types of error such as an estimate of object location and a total guess, this may be expressed in the response distributions as a bimodal combination of two distinct response types.

Finally, using an object replacement task more closely replicates the approach taken by Wang et al. (2006) who did report an effect of set size on spatial memory precision, even in small arrays, contrasting with the results in Experiment 5.1. While there remain some significant differences (e.g., Wang et al.'s study took place in an immersive virtual environment within which participants could move around) the current task is similar in terms of its difficulty and complexity. Additionally, as it is laboratory-based, we were also able to increase the maximum set size from 6 to 8 objects in the current task. If the absence of set size effects in Experiment 5.1 were due to task limitations, these modifications may increase the opportunity to observe any effects of capacity limits.

5.3.1.1 Experiment 5.2 Hypothesis

Based on previous results, viewpoint change was expected to reduce the precision of object replacement resulting in a significantly larger median replacement error. We also investigated the effect of set size on median replacement error; a significant main effect would indicate that set size has an impact on spatial memory precision while an interaction between viewpoint and set size would be consistent with the operation of distinct resource limits for viewpoint-dependent and -independent spatial memory. Finally, an exploratory analysis also investigated qualitative characteristics of the response distributions that might reveal (1) different types of error and (2) differences in distributions between viewpoint-dependent and -independent tasks.

5.3.2 Methods

5.3.2.1 Participants

All participants were recruited via the University of York's Department of Psychology participant recruitment system, 11 participants were recruited in total (10 female, mean age 18.8, SD 1.2).

5.3.2.2 Task and Materials

The experiment employed a 2x3 factorial design with two viewpoint rotation conditions (0° and 135°) and 3 set size conditions (3, 4, and 8 items). Trials from each of the 6 conditions were randomly interleaved.

All testing took place in a quiet testing room in the Department of Psychology, University of York. The task and stimuli were rendered and controlled by custom software programmed in the Unity3d game engine in C# running on an HP Z400 computer and displayed on a 20” ‘Elite Display 232’ monitor.

The virtual environment and view rotations were identical to that of Experiment 5.1 apart from the addition of a high contrast dark grey ground ‘mat’ in response to participant feedback during the pilot (Figure 5.5). One key difference in the materials is that in this experiment used set sizes of 3,4 and 8. This range was chosen for same reasons as in Experiment 5.1 and to offer comparison between studies. The 8 item sets were included as in the lab we did not have the same practical considerations as with the online task and we were able to push the set size limit further than was possible in Experiment 5.1.

As outlined above, Experiment 5.2 used an object replacement method in which in the presentation phase they viewed an array of *n* objects (sampled without replacement from a pool of 10 as in earlier experiments) for 6 seconds (Figure 5.6). The screen then dimmed over 0.5 seconds, remained black for 1 seconds and then faded back into the testing phase. They then saw the same array rotated by 0° or 135° in a random rotation direction. One of the objects had been moved an amount specified by a gaussian distribution from the original (target) location with a mean of 0.5m standard deviation of 0.2, it was also bounded at 0m and 1.5m. After 1 second the participant adjusted the position of the highlighted using the arrow keys to move it left, right, forward and backward with respect to the current viewpoint. These adjustments were animated, with the object moving continuously at a slow and controlled speed so long as a key was held down. When they judged that the object was in its original position the participant recorded their response by pressing the space bar, at which point they received feedback on where the original placement was by showing them a faded image of the object in the original target location and red text displayed above the response position to indicate the distance between target and response in meters (Figure 5.5).

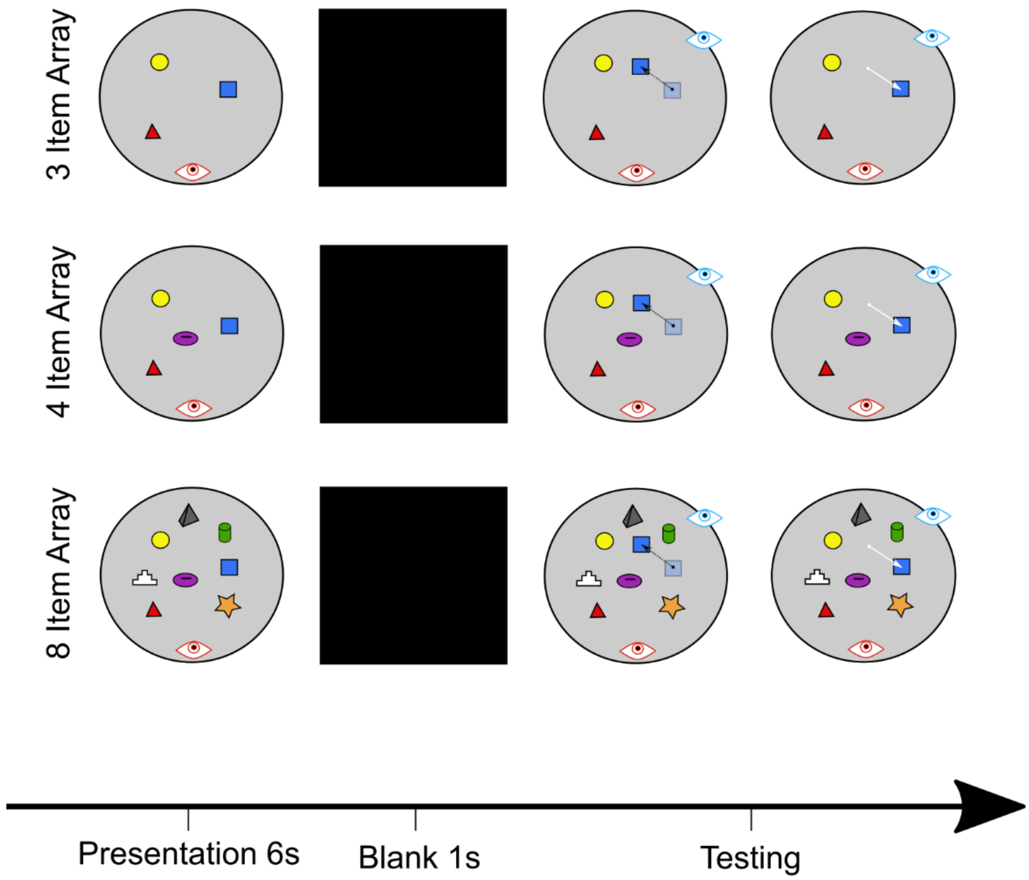

Figure 5.6 Experiment 5.2: Task and Materials . The above shows the trial procedure for the Object movement task used in Experiment 5.2. In the 'Presentation' phase, participants were shown an array of either 3, 4 or 8 Items depending on the condition of that trial. On each trial, one object is moved while the screen is blank and participants are moved around the array by either 0 or 135 degrees (denoted by the blue and red eye symbol). In the 'Testing' phase, the target object was highlighted and the participant was prompted to move the object back to its original position using the arrow keys to adjust its location. They were given an unlimited time to do this and are asked to confirm and move on by pressing the space bar.

5.3.2.3 Data Analysis

For each trial the replacement error was calculated as the absolute distance from the original target location to the response location. Median replacement error was calculated for each participant and each condition and analysed using repeated measures ANOVA.

To better understand replacement errors in terms of the array distortion, responses were additionally analysed in terms of accuracy relative to the array distortion i.e., reflecting the degree to which each response moves the object further away from or back toward the target (see section 5.3.3.2 below).

Distributions of response error (the distance from target to response) can be used in an exploratory context. Because these data reflect errors in 2D space we would expect the distribution of error to resemble a Rayleigh distribution, a curve which is bounded at 0, then peaks and is skewed so that is tails off to the right, with the peak of the distribution capturing the participants typical response accuracy.

5.3.2.4 Procedure

Before testing began, participants undertook a short, 10-minute practice session to familiarise themselves with the task. They then completed the experimental protocol. On each trial they were first presented with an array of 3, 4, or 8 objects for 6 seconds, arranged in a random configuration, generated using the same procedures as used in previous experiments (described in Chapter 2, section 2.3).

After every 30 trials the participant could take a break of an indeterminate length. The experiment ended when 50 trials had been completed in each condition (300 in total) or when 1 hour had passed. One participant did not complete all trials due to time constraints.

5.3.3 Results

5.3.3.1 Spatial Estimation: Effect of Viewpoint and Set Size

We predicted that as viewpoint rotation increased, median replacement error would increase. A factorial ANOVA showed that there was a significant main effect of viewpoint on median accuracy, with viewpoint rotation resulting in greater error (reduced precision) ($F(1,12) = 117.47$, $p<0.001$).The second hypothesis concerned whether larger set sizes would be associated with greater error in responses. Here, a significant main effect associated with set size was observed ($F(2, 24) = 30.18$, $p<0.001$). We also found a significant interaction between viewpoint shift and set size ($F(2,24) = 8.95$, $p <0.01$).

Conducting post-hoc pairwise t-tests to investigate the nature of this interaction we can see for both view conditions there was no significant differences in replacement error between set sizes of 3 and 4 ($p = 1.00$ in both cases). In the 0° viewpoint rotation conditions there is no significant difference in median replacement error between set sizes of 3 and 8 ($p = 0.08$) but there was a significant difference between replacement error in set sizes of 4 and 8 ($p = 0.01$). In the 135° viewpoint rotation conditions we find a significant difference between replacement errors for set sizes of 3 and 8 ($p < 0.01$) and also between 4 and 8 ($p < 0.01$).

In sum, across both viewpoints, participants were slightly more accurate with set sizes of 4 than 3 items but significantly less accurate when set sizes were increased to 8 items, with this effect being larger in the viewpoint shifted condition.

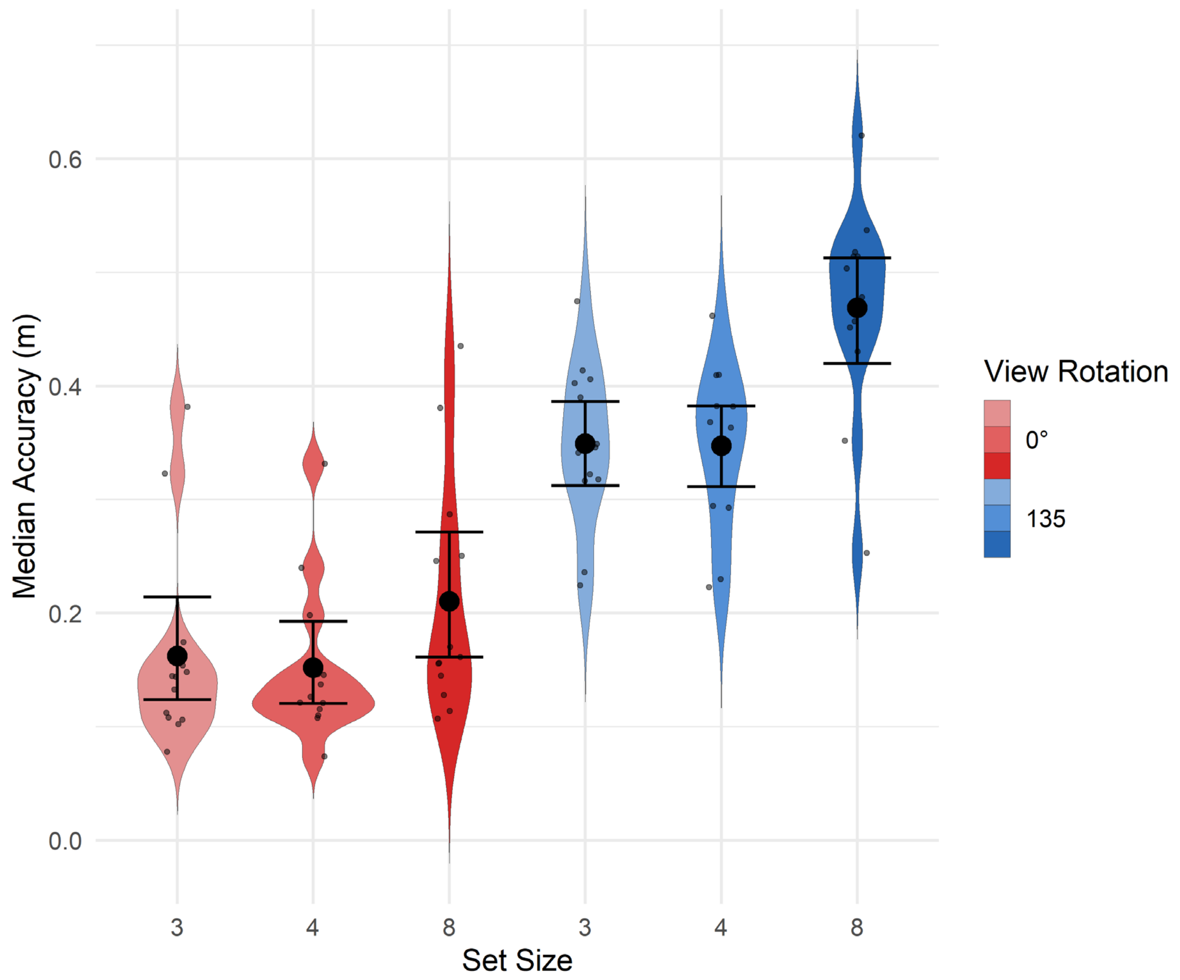

Figure 5.7 Experiment 5.2 median replacement error in meters for each condition . Set size is separated on the x axis and viewpoint shift is denoted by colour. Black points represent individual participant medians and larger black markers show group means with error bars showing standard error. Red represents viewpoint-dependent 0° viewpoint rotation conditions and blue represents viewpoint-independent, 135° viewpoint rotation conditions.

5.3.3.2 Spatial Estimation: Relative accuracy

We also calculated a metric of relative accuracy to investigate whether participants were moving objects towards or away from the target location:

Equation 5.1

$$relative\ accuracy = \frac{(d - a)}{d}$$

Where *d* is the distance an object has been moved (i.e., from presentation to the initial testing display) and *a* is the absolute distance the participant's estimate was from the target location after adjustment (replacement error). A relative accuracy of 1 is a perfect estimation in which the object is placed exactly on the original target location, a relative accuracy of 0 is when the object is not moved any nearer to the target location and a negative relative accuracy is when the object is moved further from the original target location than it was initially placed at the beginning of the response phase.

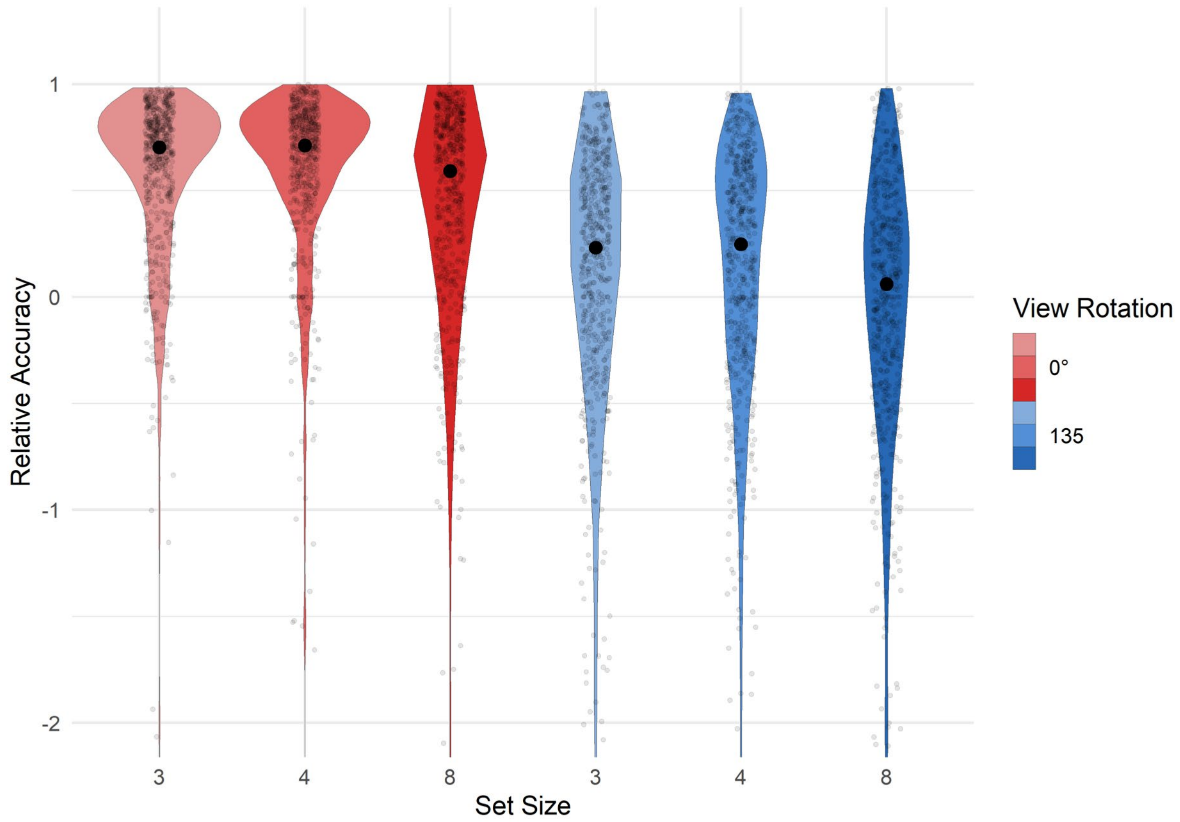

Figure 5.8 Experiment 5.2Participant accuracy relative to the amount an object was moved . A relative accuracy of 1 is equivalent to a perfectly accurate estimate, a relative accuracy of 0 is equivalent to the object being no nearer the target than when it was distorted and a negative score indicates participants moved the object further away from the target than it was initially moved. Faint points indicate raw participant data and large points represent group medians. Each condition is represented with a separate violin plot. The red colours indicate the 0° view rotation condition and the blue colours indicate the 135°. The graph is bounded at approximately -2 for visual clarity although there are raw data included within the mean- marker, that go beyond this limit.

Median relative accuracy in the view rotated conditions (especially the 8 item condition) is very close to 0 indicating that participants were moving objects farther from the target almost as often as they were moving objects towards the target.

5.3.3.3 Response Distributions

The distributions of participant error show similarities between all conditions and all resemble a unimodal Rayleigh distribution. There does not seem to be any binomial distributions indicating no clear differences in response strategy although we can see clearly the increased variance and poorer accuracy in the viewpoint-independent conditions; the 8 item, 135° condition is especially broad. All conditions seem well represented by the median value used in the main analysis (Figure 5.9).

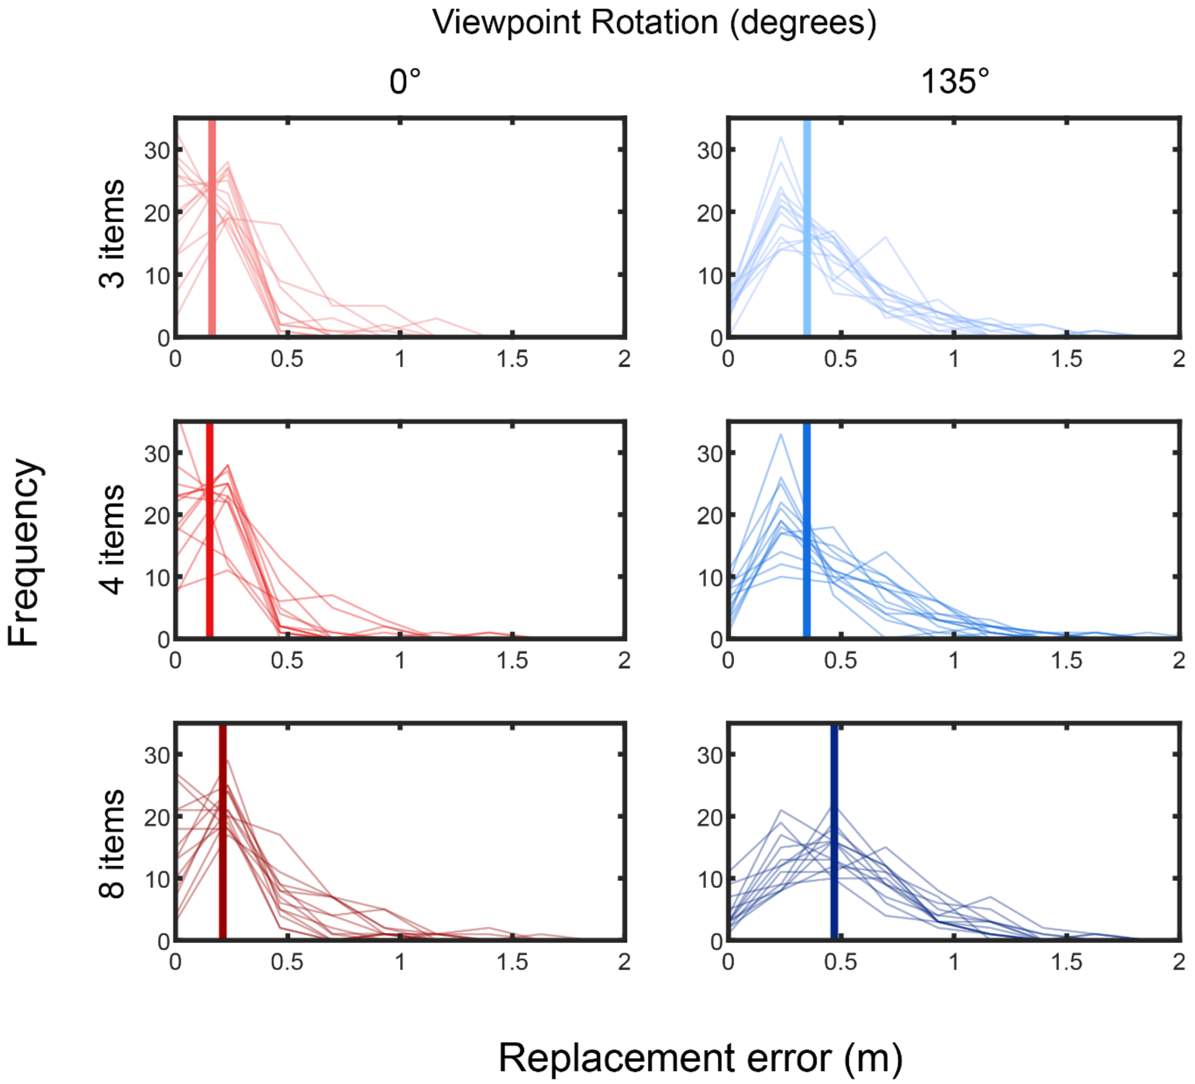

Figure 5.9 Experiment 5.2 distribution replacement error . Each plot shows a histogram displaying data from a given experimental condition divided across 12 evenly spaced bins between 0 and the maximum distance error. The thick vertical line represents the mean median response accuracy across all participants within that condition. The thin lines represent individual participant response distributions.

5.3.4 Discussion

As predicted, we found a significant effect of viewpoint on the replacement error. This is consistent with results from previous experiments investigating spatial memory precision using change detection thresholds, suggesting that the task provides a viable alternative to the use of a staircase procedure and binary decision. In this experiment, unlike Experiment 5.1, we found a significant effect of set size on spatial memory precision. Participants were significantly less precise in both viewpoint-dependent and -independent tasks when tested at a set size of 8. We also found a significant interaction in which set size had a larger impact on precision in the viewpoint rotated condition.

Generally, the results of the present study show similarity to visual Working Memory studies in which there is significant drop in performance after approximately 4 items (Cowan, 2001) suggesting that this task relies on or is constrained by limiting properties shared with visual Working Memory. In this task a set size of 8 markedly reduced the precision of participants' responses.

Based on median performance alone, it is unclear whether we are seeing a hard limit of spatial memory in terms of information or number of items. In terms of relative accuracy, the data showed that participants were, on average, moving objects towards the target. Median relative accuracy was greater than 0 in all conditions. However, the viewpoint-independent- 8 item condition mean was very close to 0 indicating that participants were moving the target object further from the target almost as often as towards the target. This result on its own this might indicate that participants were guessing on a large proportion of the 8 item, view independent trials. However, when we take into account distributions of absolute accuracy, the 8 item, 135° response patterns follow a very similar but slightly flatter, Rayleigh distribution to those seen in other conditions which may indicate a single response type. If participants were truly guessing we might expect even flatter distributions or perhaps a second peak (a bimodal distribution reflecting guesses as well as distinct imprecise responses). The poor performance in the 8 item, 135° condition described by the median, relative responses and response distribution may not be a hard capacity limit resulting in a guess, but just a much lower spatial memory precision. This would then indicate that error did not show qualitatively distinct 'type' (as might be expected in a discrete slot model), but degraded gradually and continuously. A steady decrease in precision as the number of items increase is in line with the informational trade off hypothesis (i.e., Luck & Vogel, 1997) in which a higher number of items results in a lower precision.

Our results reflect the conflicting nature of the spatial memory capacity literature. Unlike Hodgson and Waller (2006), in the current task we did observe an effect of set size on precision implying that this task relies on an active and limited process as opposed to a long term memory representation. This replicates the findings of Wang et al (2006), however unlike that study we saw extremely little effect of set size until the array contained a much larger number of items.

These disparities with the literature may be associated with task differences. The similarity between these tasks is that they each rotate the viewpoint between presentation and test, measure spatial memory precision and manipulate set size. However, there are also important differences in terms of (1) spatial cues available, (2) different measurements of spatial memory precision and (3) how stimuli were presented.

Firstly, the task used in Hodgson and Waller (2006) involved no physical movement through the environment and only rotation in place, from the centre of the array whereas in Wang et al (2006) participants moved around the environment and so richer idiothetic and spatial updating/movement cues were available. There is evidence of difference processes involved for spatial memory from within and outside of environments (Motes et al., 2006) and it may be considerably easier to spatial update the location of objects while rotating in place, resulting in the participants in Waller and Hodgson performing at ceiling. The present study was virtual and so idiothetic cues were unavailable however, participants did change location in the environment as well as rotate which bears similarity to the methods of Wang et al. (2006).

Secondly, regarding measurement of performance in each task; precision in the Hodgson and Waller (2006) study was measured in terms of the angle between estimate and target from the centre the array. This makes the task easier as participants only had to retain direction and not distance or placement in 3D space. In the present study, and also Wang et al. (2006), participant had to reconstruct arrays, replacing objects freely in the available space. Not only is this more difficult in principle, it is also potentially a more sensitive measure of accuracy, taking into account distance as well as direction error.

Finally, in Hodgson & Waller (2006), the locations of objects did not change between trials and testing did not take place until participants had adequately learned object locations and so, unlike Wang et al. (2006) and the present study in which object were reconfigured on each trial, participants had significantly more opportunities for rehearsal or other strategies to be employed which may indeed have activated long term memories which were not formed in the other tasks.

One unique finding to the present experiment was an interaction between viewpoint rotation and set size where set size exerted a larger effect in viewpoint-independent memory conditions. This could be interpreted as being due to distinct resources and capacity limits. While the present experiment does not resolve the question, it is consistent with the idea that view dependent and independent behaviour may depend on separate (ego- and allocentric) representations, each subject to its own capacity constraints. However, in the current task viewpoint-independent memory seems to have a smaller capacity, something we might not have expected at the outset of these experiments.

5.4 General Discussion

In two experiments we investigated the effects of set size on viewpoint-dependent and -independent memory precision. In Experiment 5.1, using a 2AFC task we found no significant effect of set size on spatial memory change detection thresholds in either viewpoint-dependent or independent memory tasks. However, we did see significant effects of set size on response times and an interaction showing that set size had a larger negative effect on viewpoint-independent memory response times. There was also a significant interaction between set size and response time with increased latency for viewpoint rotated conditions and larger sets. This raises the possibility that time-dependent serial/itemwise comparison or matching processes are specifically involved in the viewpoint-independent task.

In Experiment 5.2 we increased the maximum set size to 8 objects and changed the task to one in which participants were required to replace a moved target object to its original position. We found significant effects of viewpoint and set size on participants' accuracy, observing significantly larger errors when they were required to remember 8 item arrays. We also saw an interaction that showed set size had a larger effect on error in the viewpoint-independent task. By calculating a metric of relative accuracy that described whether responses moved the object towards or further away from the target location, we found that in the viewpoint-dependent condition participants were performing near ceiling, consistently moving objects towards the target with very small errors. In the viewpoint-independent, 8 item condition, participant performance was near floor and could be considered equivalent to guessing in that, objects were moved further away from the target location almost as often as they were moved towards it.

The responses from each condition show similar shaped distributions, resembling a Rayleigh distribution in all conditions with broader distributions as set size and viewpoint rotation increased. There was no evidence of a bimodal distribution which we might have expected to see if participants were using systematically incorrect strategies or in a simplistic slot model in which object locations are either stored with precision or not at all.

Together, although the results from these experiments do not provide definitive answers, they do go a long way towards exposing potential capacity limits in viewpoint-dependent and -independent memory, showing that there are significant differences in the capacity limits in spatial memory when the viewpoint is static compared with when it is moved between presentation and testing. There are also indications that there are Working Memory like informational limits on processes involved in viewpoint rotation tasks; a trade-off between the number of objects to be retained, and the precision with which they are represented. We found that participants could remember the location of up to 6 objects without any negative effect on precision but that objects from arrays of 8 were replaced with significantly less precision than those from arrays of 3 or 4.

We saw interactions between set size and viewpoint on response times (Experiment 5.1) and precision (Experiment 5.2). In each case, set size had a larger negative effect on performance in the viewpoint-independent memory task. This is an interesting result and contradicts the predictions of some authors (and our own initial intuitions) who argue based on links between allocentric and episodic memory (Hodgson & Waller, 2006), the opposite interaction should be seen (i.e., view independent memory should be insensitive or less sensitive to set size). This is a good indication that there are differences in the capacities of viewpoint-dependent and viewpoint-independent tasks although there are a number of possible explanations that could be compatible with this observation. For example, it could be due to the use of separate ego- and allocentric resource pools, or perhaps additional task requirements added by the viewpoint rotation involving itemwise processes as suggested by the response time interaction in Experiment 5.1 (e.g., matching/comparison of locations from different viewpoints), or some combination of these factors.

Although the tasks from Experiment 5.21 and 5.2 are very different, it is possible to express both the change detection thresholds and replacement errors in the same units and thus compare the findings (Figure 5.10).

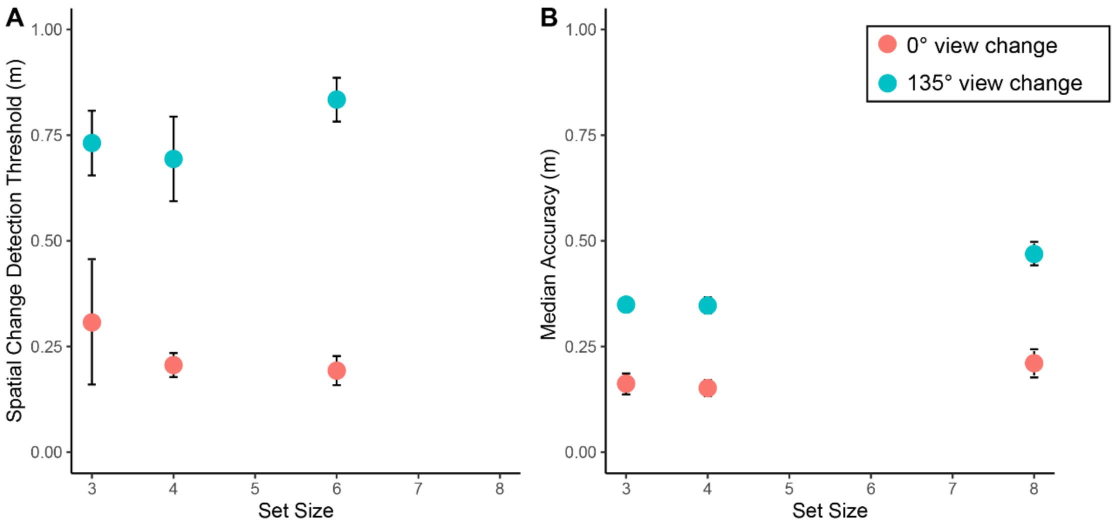

Figure 5.10 Comparing measurements of spatial precision in meters from Experiment 5.1 (A) and 2 (B) . A shows spatial change detection thresholds as recorded in Experiment 5.1 at set sizes of 3, 4 and 6 and viewpoint changes of 0 and 135°. B shows mean median accuracy in meters in each condition of Experiment 5.2- set sizes of 3,4, and 8 and viewpoint changes of 0 and 135 degrees. Error bars show standard error.

We see that replacement errors are typically smaller than spatial change detection thresholds and this is especially true in the rotated view condition. The effect of viewpoint rotation was much smaller in object movement task. Spatial precision declined much less due to view rotation (the offset of blue and red markers in Figure 5.10). This suggests that there may be something specific to the 2AFC task which is more difficult for viewpoint-independent memory. This difference may concern the number of objects that are probed in the testing phase. In order to complete the 2AFC task, the participant must both query the spatial location of two objects. In task 2, as the single target object was preselected for participants to replace, this reduced the number of items needed to be probed in memory. While no definitive conclusions can be made by comparing two quite different tasks and measures, this comparison perhaps give yet more indication of the task related nature variability of viewpoint-dependent and -independent memory

One pattern in the data which can be seen in the results from Experiment 5.2 and in previous chapters is an effect on spatial precision in a viewpoint-dependent task which is then amplified in the viewpoint-independent task (the interaction between viewpoint and set size). This brings us back to the notion of interdependence between egocentric and allocentric information. It is possible that set size affects the precision with which object-locations are initially encoded, egocentrically, and this effect is amplified during the process of conversion of retrieval to and from an allocentric reference frame.

It is also possible that there were more strategies available to participants in the static view condition, due to the relative simplicity of the task. This was suggested as a possible reason for response time differences in Hodgson and Waller (2006). Another possibility for future research would be to investigate possible set size related strategies such as spatial chunking or verbal strategies. One solution could be to investigate the spatial relationship within arrays and how that relates to threshold. For other strategies, a distractor could be employed in the task disrupt strategies that were available in only one condition (a common approach in the Working Memory literature that has been used to dissociate e.g., visual and verbal components of Working Memory). If this were the case, disrupting these strategies should result in a similar capacity/ precision balance in both viewpoint-dependent and -independent tasks and no significant interaction between set size and viewpoint manipulations.

5.5 Conclusion

This chapter focussed on spatial memory capacity with mixed results. Using the change detection paradigm as in other chapters, no significant effect of set size between 3, 4, and 6 objects on change detection thresholds was found. However, there was shown to be a significant effect of set size on response time. We used a variation on the methodology and measured participant accuracy when they were asked to replace a cued object. Here, a significant effect of set size in both viewpoint conditions was shown when the size was 8. This indicates that there are Working Memory-like limits imposed at some point in the process of viewpoint-independent spatial memory and that there are lower capacity limits for viewpoint-independent as opposed to viewpoint-dependent memories. Further research will be needed to determine whether these differences are related to separate informational resources responsible for each type of memory or whether processes specifically involved in viewpoint rotation incur additional 'resource cost' which lowers capacity limits.

Chapter 6. Discussion and summary

6.1 Introduction and contributions

Spatial memory is supported by multiple representations that are distinct in how spatial information is encoded. Previous work has suggested that both egocentric and allocentric representations contribute to spatial behaviour, but their detailed properties, contributions and interactions are not fully understood. One way to characterise a memory system is by determining its limits, but at the outset of this project, it was not clear how the capacity and precision of viewpoint-independent memory might be measured in the laboratory (or indeed online). Many studies had measured performance in a variety of relevant tasks, but few had investigated the fidelity of allocentric spatial memory in any detail. Notably, it was previously unclear how the precision of spatial behaviour would be affected in tasks where the observer's viewpoint changes unpredictably, is affected by the scale of the stimulus or environment, or by the number of items and locations to be stored.

By examining these questions this thesis has made several new contributions.

i. I developed new psychophysical methods for the assessment of spatial memory precision and its relationship to viewpoint, scale and capacity (set size).
ii. I identified a function that relates changes in viewpoint to spatial change detection thresholds.
iii. The function separately quantifies viewpoint-specific, viewpoint-dependent and viewpoint-independent components of spatial precision.
iv. I determined that the scale of (1) stimulus arrays and (2) environments affect viewpoint-dependent and -independent memory precision. Larger scales negatively affect spatial memory precision independent of viewpoint rotation.
v. I established that both viewpoint-dependent and -independent spatial memories are consistent with Weber's law for position, meaning that spatial memory precision scales with the visual scale rather than absolute physical size of the stimulus across all tasks and conditions so far investigated.
vi. I found that set size negatively impacts precision for sets of 8 items in an object replacement 3D spatial memory task indicating working memory-like constraints.

vii. Set has a larger effect on viewpoint-independent precision indicating distinct capacity limits in viewpoint-dependent and -independent tasks.

In the next section I review the most important findings in more detail before considering recurring themes that integrate the results from different chapters.

6.2 Main findings

6.2.1 Developing methods

Prior to this thesis, few methods existed that explicitly measured the spatial precision of viewpoint-dependent and -independent memory to a high degree of accuracy. By combining techniques from psychophysics (such as the adaptive staircase method and psychophysical threshold estimation) with a spatial memory task that included viewpoint manipulation to vary the dependence of responses on ego- and allocentric information, a set of methods was developed that could reliably estimate observers' spatial change detection threshold. This threshold denoted the distance beyond which a change in an object's location was noticeable after a short delay. This provided a quantitative measure of the precision or fidelity of the underlying spatial information retained by the observer between presentation and testing.

The development and subsequent use of the spatial change detection task (Chapter 2) was largely successful. The results from pre-pilot and pilot experiments were an encouraging basis for using the methods through the thesis. The most notable and important take-aways from the process relate to the practice of psychophysics and the importance of task efficiency and correct parameter setting.

The statistical power of psychophysical methods comes from the number of trials rather than the sample size (Baker et al., 2020) and the abundance of data collected from individual participants as well as the high degree of replicability between subjects provides confidence in the validity of the results. One practical point to emerge from the process of developing and piloting the methods was that the number of independent variables (each requiring a valid psychometric curve) and the number of trials (more trials resulted in better fitting curves) would have to be carefully managed to ensure that participants were able to complete the experimental sessions while cooperating with and fully attending to the tasks. Subsequently, an important part of developing the experiments in Chapters 3, 4, and 5 was in balancing and honing the scope of each experiment to reconcile the need for high-quality data (to support well-fitting psychophysical curves) with the number of independent variables and conditions required to address the research questions.

As my primary interest was with allocentric spatial representations, at the outset of the project the plan was to conduct most experiments using with a fixed viewpoint change. In this original design, all trials were rotated 60° between presentation and test, with a focus on investigating the effects of other independent variables.

The pilot studies were helpful in indicating the range over which change detection thresholds might be measured, but it was not possible to interpret these in terms of allocentric representation in the absence of a more egocentric version of the task. This suggested that a zero-degree condition should be included for comparison. On reflection, we realised that rather than limits being either allocentric or egocentric there might be a more complex and parametric relationship between precision and viewpoint. To investigate this, we set about designing an experiment incorporating multiple levels of viewpoint change.

The scope of this method meant that it was important to distinguish terms that separated spatial representations from their expression in behaviour. While our task aimed to isolate and understand *ego*- and *allocentric* contributions (referring to the encoding of space e.g., by spatial cells within the parietal cortex or hippocampal formation), it tapped into *viewpoint-dependent* and *viewpoint-independent* behaviour yielding measurements that might, in principle, reflect the combination of these representations.

6.2.2 Spatial memory precision and viewpoint change

The main finding from initial experiments, which measured spatial change detection thresholds at different systematic viewpoint rotations, showed a monotonic but non-linear increase in thresholds as viewpoint rotation increased.

Using results from experiments in this thesis, we can model how change detection thresholds vary as a function of viewpoint rotation. This descriptive model parameterises putative cognitive components involved in this process and allows us to interpret performance at different degrees of rotation in terms of the spatial representations that may be contributing to behaviour. By parameterising these components, we are also able to provide a highly specific and detailed comparison of spatial memory precision between participants or subject groups potentially revealing different profiles.

This is a descriptive model which does not attempt to address the underlying mechanisms. However, one can tentatively map performance in the zero-degree condition, as determined principally by availability of a very high-fidelity egocentric information. This is evidenced by the negative and additive effect of viewpoint rotation, whereas performance in the post-plateau region is determined principally by the allocentric information. It is less clear how these sources interact in determining viewpoint-dependent responses, i.e., it is not clear whether performance between zero-degree and plateau is driven solely by degraded egocentric or whether it involves integration of ego- and allocentric sources.

A lawful pattern of curve and plateau as viewpoint rotation increased was present in all participants in both a Same/Different and 2AFC task. However, when comparing model parameters, we find there were significant task differences in viewpoint-dependent and -independent precision. Compared to a 2AFC task, performance in a Same/Different task showed significantly higher thresholds at lower degrees of rotation and significantly lower thresholds at higher degrees. When participants were asked to pay attention to the entire array rather than two cued items, they may be more sensitive to allocentric object-object or object-environment cues, and subsequently perform better in the viewpoint-independent conditions.

The stimuli developed for this task and used throughout the thesis were designed to allow unpredictable but carefully controlled viewpoint changes while avoiding unintended confounds from uncontrolled environmental and spatial cues (e.g., a sparse environment, directional cues that did not assist in the task, instant movement). We might expect that in the real world, viewpoint-independent spatial precision is better than that observed in the lab, both because of task/context effects like those seen in the comparison of Same/Different and 2AFC tasks and because of the availability of richer spatial information including idiothetic cues and richer environments. It would be informative to see how well the patterns seen in precision of these experiments apply to other tasks.

In sum, based on a high quality measurement of spatial precision, we can separate and quantify viewpoint specific, viewpoint-dependent and viewpoint-independent components of spatial memory. By applying our knowledge of the behavioural and task characteristics of spatial representations in the brain, we can model the interaction between viewpoint change and spatial memory precision and understand more about the distinct contributions made to this task.

6.2.3 Spatial memory and scale

Substantial evidence indicates that viewpoint-independent memory relies on an allocentric spatial encoding based in the hippocampal formation. However, we also must bear in mind that spatial information is initially perceived and encoded egocentrically (Section 6.3.1). Factors affecting the precision of this initial egocentric representation ultimately constrain the precision of the allocentric encoding and subsequent viewpoint-independent task performance.

One example of this can be seen by examining how different spatial scales are encoded in a viewpoint rotation task. We might expect there to be a difference in how scale is encoded between viewpoint-dependent and -independent memory. Viewpoint-dependent change detection thresholds can be expected to adhere to Weber's law, which states that precision is relative to the magnitude of the stimulus (i.e., a larger array would be encoded with lower precision). By contrast, allocentric spatial cells encode locations in a way that is anchored to the environment. As such we might have expected viewpoint-independent precision to retain this fixed and absolute character thus being insensitive to changes to the scale of the stimulus or environment.

However, when the footprint of a stimulus array was doubled a significant effect of array scale on change detection thresholds was found, and no interaction between the effects of scale and viewpoint change. This indicates that both viewpoint-dependent and -independent thresholds obeyed Weber's law, a property which is thought to be based on egocentric perceptual properties, such as visual angles.

This finding was confirmed in a second experiment. Here, an innovative approach used immersive virtual reality to make large changes to the absolute scale of the environment (and to participants perception of scale) while controlling the participant's visual image (the relative locations of objects and environment in the field of view) such that pictorial cues remained the same. Comparing change detection thresholds between room-sized and amphitheatre-sized environments we see that thresholds scale 1:1 with the size of the environment in a viewpoint static and viewpoint rotated task.

Together, this information allows us to conclude that the effect of scale seen in a spatial change detection task was not due to a change in the absolute scale of the stimulus, but due to a change in its extent within the observer's field of view. Overall, the results indicate that the precision of ego- and allocentric contributions to spatial memory are affected in the same way. When found in a viewpoint-independent task, this effect can be interpreted as a consequence of the impact of visual scale on initial egocentric encoding which has repercussions on allocentric precision.

6.2.4 Spatial memory capacity

The limits of memory systems are classically considered in terms of capacity, the number of items that can be retained. The spatial change detection task developed in this thesis provides an excellent opportunity to investigate differences in the capacity of viewpoint-dependent and -independent spatial memory and its relationship with precision. Because of its links with allocentric representation and the hippocampal formation, viewpoint-independent memory might be expected share qualities of episodic memory, which is typically seen as having an extremely large or unlimited capacity. This characteristic contrasts with visuo-spatial working memory which, like other forms of working memory, is known to be severely limited in its capacity. However, visuo-spatial working memory tasks typically do not involve any change in viewpoint and are thus expected to rely on egocentric representations. Previous investigations with spatial memory tasks that incorporate viewpoint manipulations (Hodgson & Waller, 2006; Wang et al., 2006) have shown very mixed results, so the capacity limits of viewpoint-independent spatial memory were not clear at the outset.

Although one experiment showed no effect of set size on change detection thresholds for sets varying between 3 and 6 items, we did see signs of differences between viewpoint-dependent and -independent memory in terms of an interaction between set size and viewpoint on response times. In another experiment, we saw a clear effect, such that set size did significantly and negatively impact precision, but only for sets of 8 items in an object replacement task. Interpreted in terms of a trade-off between capacity and precision, this is evidence that there are working memory-like capacity limits in the viewpoint rotation task which extend to viewpoint-dependent and viewpoint-independent memory. Furthermore, there was evidence in support of different capacity limits in viewpoint-dependent and -independent memory. However, this interaction showed the opposite to what we might expect: that viewpoint-independent memory in fact had a lower capacity. This distinct capacity could be explained in several ways; in terms of distinct resource pools for egocentric and allocentric spatial memory and/or increased task difficulty and additional resources required for the viewpoint rotated condition.

While this study did not explicitly investigate models of memory capacity, responses in an object replacement task favoured a resource model in which, as the number of items in a set increased, there was a trade-off between capacity and the precision of responses. This was supported by response distributions which showed progressively flatter curves as set size increased but no evidence of distinct, different responses or guesses as we might expect in a slot-based model.

The results from this thesis emphasise the task-dependent nature of spatial memory capacity in which changes in response type and spatial cues at encoding seem to quite drastically alter the conclusions drawn. While there is now clear evidence of distinct viewpoint-dependent and -independent capacity, further study is needed in understanding factors contributing to this task-sensitive phenomenon.

6.3 Overarching themes and explorations

6.3.1 Systemic constraints on viewpoint-independent behaviour

This thesis draws on the dominant account of the relationship between ego- and allocentric reference frames in spatial memory as described in a number of reviews (Burgess, Maguire, et al., 2002; Ekstrom et al., 2014; Klatzky, 1998) and implemented in mechanistic models of the relevant brain systems (Bicanski & Burgess, 2018, 2020; Byrne et al., 2007; Sheth & Shimojo, 2004). In this view, an initial egocentric spatial memory representation that codes scene elements (boundaries, landmarks, objects) is formed in the medial parietal cortex from high-level sensory processing. This is then converted into an allocentric code in the retrosplenial cortex using a gain field circuit and input from head direction cells as gain-field modulation. This process converts an egocentric representation encoded in the parietal cortex to an allocentric code within the hippocampal formation i.e., egocentric boundary-aware cells are mapped to boundary vector cells representing the structure of the environment, and egocentric object vector cells maps to object vector cells found in the medial temporal lobes representing the objects within it and their locations (Bicanski & Burgess, 2020; Hoydal et al., 2018). To bring a spatial memory into imagery, the process is reversed. Head direction cells are used to apply a viewpoint to the allocentric code so that it can be experienced in an orientation-specific manner, allowing for the control of (body-referenced) behaviour and decisions.

The implications of this account, for those that investigate spatial memory behaviourally, is that it is impossible to conduct a ‘process pure’ allocentric task, hence the distinction early in the thesis between allocentric representations and viewpoint-independent behaviour. The interdependence of ego- and allocentric systems and the consequent unavailability of process pure allocentric tasks does not nullify behavioural results when testing viewpoint-independent spatial memory, but it means that their interpretation requires an appreciation of the ‘path’ the information has taken before it is used to facilitate performance in a task.

The process and system involved in encoding allocentric spatial information could impact the precision of the spatial information in at least two ways. Firstly, as the allocentric encoding is ‘sandwiched’ between egocentric codes means that there is the possibility that system limits (in precision or capacity) are filtered by way of an ‘egocentric bottleneck’ and secondly, the process of conversion between codes may also impact precision (Figure 6.1).

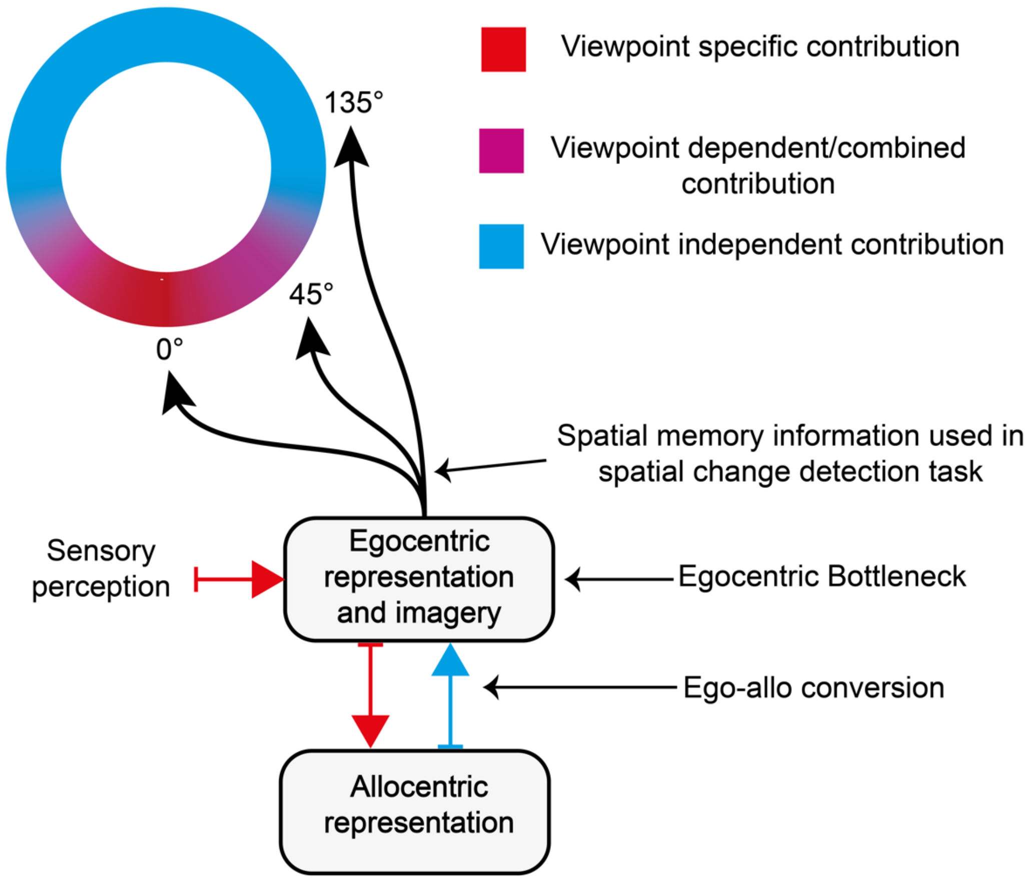

Figure 6.1 A simplified diagrammatic representation of spatial information flow for use in the spatial change detection task. Perception is encoded in an initial egocentric representation which is then converted into an allocentric encoding. To use this allocentric information it must first be converted back into an egocentric imagery. In the viewpoint rotation task we see viewpoint specific information which is susceptible to the effects of viewpoint used up until approximately 45° of viewpoint rotation. Past this point precision does not degrade as a function of view rotation indicating the use of viewpoint-independent, allocentric information. Due to this order of processing, when trying to measure allocentric memory precision through viewpoint-independent behavioural tasks, we must be aware that (1) there may be a bottleneck that constrains and shapes information going to or coming from the allocentric code and (2) that the conversion between reference frames may also impact spatial memory precision.

The egocentric bottleneck is a term coined within this thesis. It can be defined as potential systemic constraints on allocentric memory precision and capacity arising from the limits of egocentric representation during encoding and retrieval. This raises the possibility that the limits of allocentric spatial memory cannot be accessed behaviourally in viewpoint-independent memory, as they may fall outside of the limits of egocentric spatial memory. For example, in the case that the initial egocentric encoding lowers the resolution of information before it is converted or that an allocentric representation is of a higher spatial resolution than is possible for egocentric spatial representations when it is converted back to imagery.

This dependence and connectedness between ego- and allocentric spatial memory is most clear in the findings of Chapter 4 which show effects of environmental scale independent of viewpoint. These effects were shown to be based on the visual image of the stimuli rather than the absolute scale, something which is an egocentric quality and viewpoint specific. The findings from this chapter demonstrates how effects on the precision of the egocentric representation carry over and have lasting impact on the quality of viewpoint-independent memory.

A separate implication of the relationship between ego- and allocentric representations is the potential impact of a code translation between ego- and allocentric reference frames on allocentric spatial precision. Assuming a model like Bicanski and Burgess (Bicanski & Burgess, 2018), it is possible that the relative drop in precision between egocentric and allocentric codes is related to the process of translation which creates a fixed offset of additional noise and inaccuracy within the allocentric representation (Figure 6.2 A).

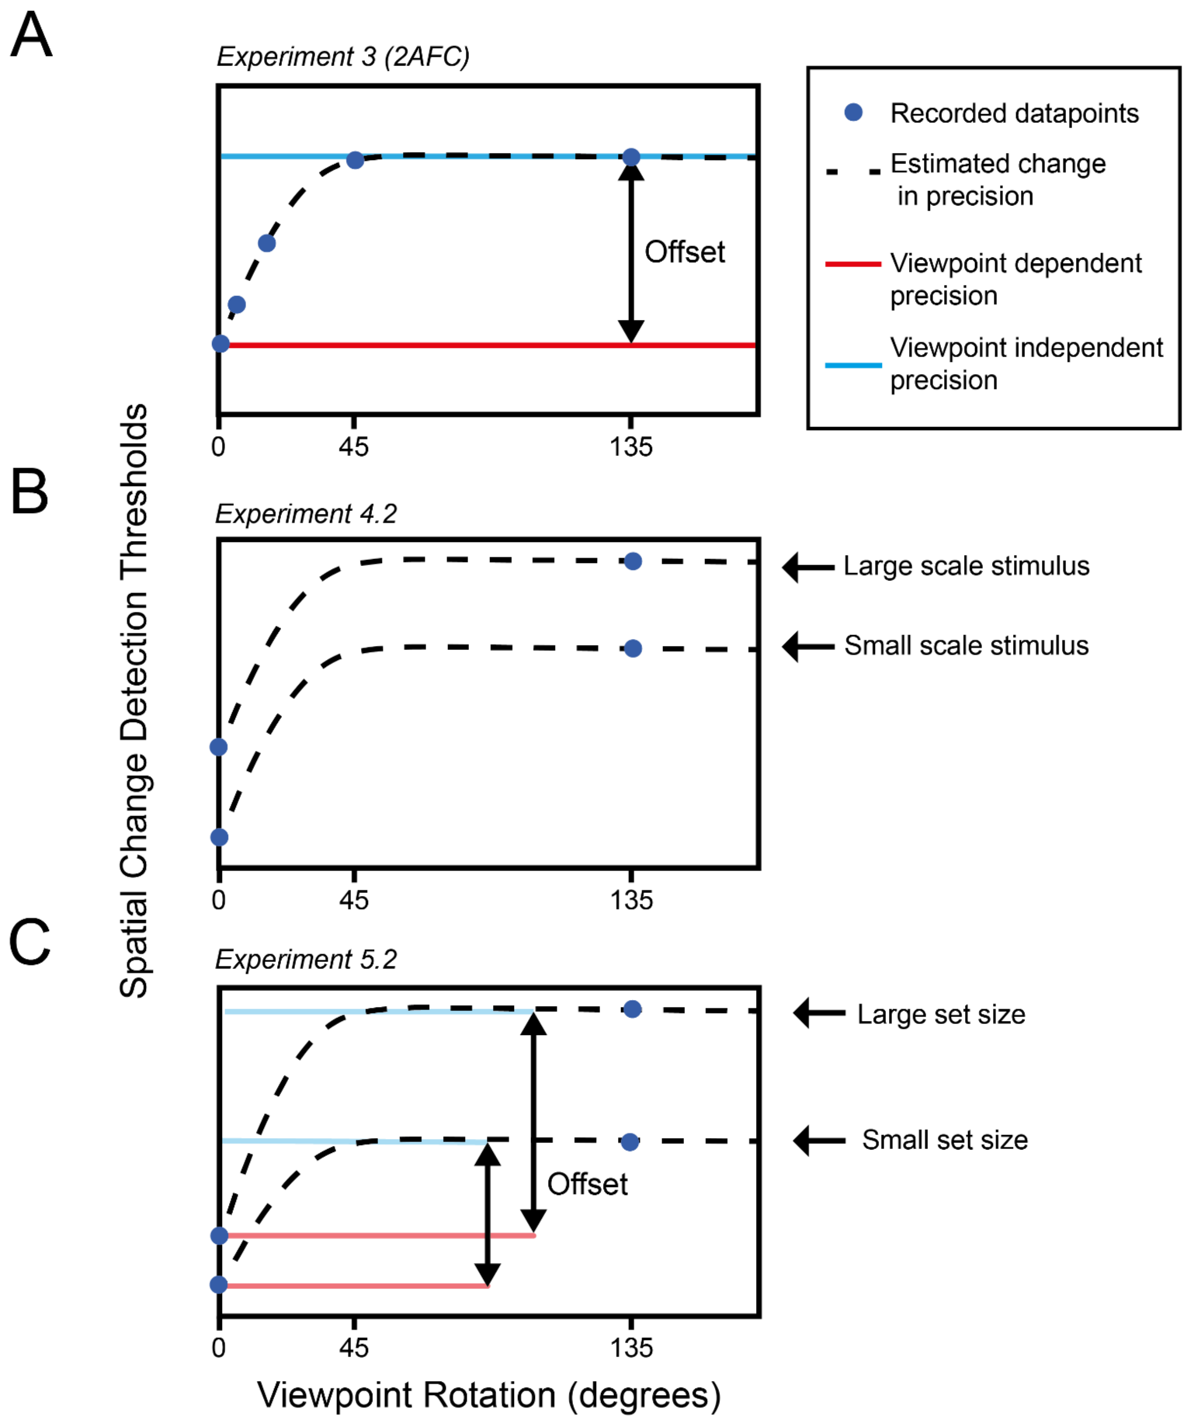

Figure 6.2 Exemplifying the model of the effect of viewpoint change on spatial memory precision theoretically applied to the results from subsequent experiments . Dark blue points represent approximations of datapoints recorded in this thesis and dotted lines show the estimated change in precision if viewpoint was systematically rotated through 180°. A shows the results from Chapter 3 and denotes a level of egocentric and allocentric precision with red and blue bars. The distance between these bars is the offset. B, shows the model theoretically applied to the data recorded in Experiment 4.1 and 4.2. Here we showed that changing the scale of the environment effected spatial precision independently i.e., in a small and large condition, the offset between egocentric and allocentric precision remained the same. C shows approximation of the results from Experiment 5.2. here set size impacted viewpoint-independent spatial precision significantly more than viewpoint-dependent spatial precision. Here the offset changed between set sizes.

This explanation can be applied to the findings of Chapter 4. The effect of scale was independent of viewpoint, and we can imagine that if we were to test more viewpoint rotation conditions we might see a similar curve and plateau as in Chapter 3, fixed and relative to the precision in the viewpoint-dependent condition (Figure 6.2A and Figure 6.2 B). Changing the environmental scale, the curve rises and falls accordingly but remains at this fixed offset. Experiment 4.2 showed that this translation process seems to relate to angular change in viewpoint rather than distance or magnitude moved, here the physical and absolute distance between viewpoints varied massively and the relative offset between viewpoint conditions was not affected by the absolute distance between viewpoint positions.

However, as shown by the interaction in Experiment 5.2, it seems that by manipulating capacity we see the viewpoint-dependent and -independent precision offset change (Figure 6.2 C). Additional items in the array seem to exaggerate and amplify the deficit in the viewpoint-independent task. It is possible that this can be attributed to the demanding process of ego-allo translation requiring additional attentional resources. Effects of set size on reaction time in Experiment 5.1 may support this theory indicating that in the viewpoint-independent task there was evidence of a serial or itemwise process that affected the response time in only the rotated view condition. This is an unexpected finding as updates to viewpoint-independent might be expected to be coherent and based on a single estimate of head-direction (Bicanski & Burgess, 2018). However, this process may not reflect spatial updating, but instead something more specific to the 2AFC task. One possibility is that constructing imagery with which to compare current locations perhaps requires a piecemeal, itemwise process whereas viewpoint specific memory may benefit from the comparison with snapshot like visual representation against which changes jump out (somewhat analogous to the difference between serial and parallel visual search, Treisman & Gelade, 1980).

In this thesis (Chapter 3) an allocentric precision limit does become clear when viewpoint rotation exceeds 45°, for a stimulus at fixed scale and capacity. As scale and capacity vary, the allocentric precision limits are determined principally by characteristic of the egocentric representation in the tasks. There are some exceptions to this, and we do see some evidence that allocentric representation (or translation to/from allocentric representation) is linked to precision for large set sizes (>6 items). The pattern of response times in this task suggest that itemwise decision processes specific to viewpoint-independent memory play a part in this.

6.4 Future directions

A major application of the findings of this collected study would be to inform current mechanistic models of spatial memory to further our understanding of how these systems interact. Beyond this, and the other specific suggestions made, it would be very revealing to apply the findings of this thesis in a neuroscientific context.

A neuropsychological study would be useful in this regard. By replicating the viewpoint rotation study in Chapter 3 with a hippocampal patient we could model performance with purely viewpoint-dependent, egocentric information. If our interpretation of the data is correct and the curve and plateau is due to two systems, in purely egocentric performance we would see no plateau and change detection thresholds would continue to rise after 45°.

Replicating the studies on spatial memory capacity with a hippocampal patient would also be highly informative. Presumably all viewpoint rotated performance in this case would be extremely low or at chance although it seems very small set sizes may be an exception to this (King et al., 2002; Wang et al., 2006). It would be interesting to experiment with smaller systematic view rotations, to find the minimum angular rotation at which performance dropped for a hippocampal patient and whether an effect of set size interacts with this.

Neuroimaging could also offer insights and help test the models developed in this thesis, potentially linking precision and capacity limits to specific brain structures in healthy people. Schmidt et al (Schmidt et al., 2007) has shown a correlation between successful completion of a viewpoint rotation task and fMRI BOLD activation of areas within the hippocampal formation. If the spatial change detection task with parametric viewpoint rotation (c.f., Chapter 3) were to be conducted within an fMRI scanner, we might expect contributions from ego- and allocentric brain system to track components of the descriptive model developed here (i.e., code change at around 45°). Although the exact nature of the relationship between model and activation in the different system is not obvious, a neuroimaging study along these lines would potentially link the model and its parameters to activation in specific brain regions and rotation conditions. This would provide an integrated account of behavioural and neuroscientific underpinnings of interaction between viewpoint-dependent and -independent memory. Such a study would help to segregate behavioural components to parts of the brain and relate brain activity to specific levels of precision.

Finally, the spatial change detection task has potential applications in a clinical context. The pathology of Alzheimer’s disease shows early atrophy of the hippocampus and associated deficits in viewpoint rotation tasks (Bird et al., 2010; Chan et al., 2016; Wood et al., 2016). If the methods developed within this thesis offer a precise reliable way to assess allocentric memory function, this task could be developed into a useful diagnostic tool, able to differentiate Alzheimer's disease from other dementias or to monitor the progression of the disease in an individual. This task offers important advantages over other viewpoint rotation tasks in that it measures spatial precision explicitly and provides a highly detailed profile of performance across a testing range. Obviously, there are several steps required before validation of such a tool could take place. Primarily, this would involve a thorough investigation of the internal validity of the test, for example specific analysis of individual differences, large sample correlations and longitudinal study. Secondly, before use within a clinical context it would be beneficial to make the test more appropriate for the patient group. This would largely involve simplifying instructions and reducing the time it takes to estimate a participant's viewpoint-independent threshold level of ability. There are several ways that this could be done including various alterative adaptive approaches (Leek, 2001) which could expedite the threshold estimation process. A final step in this process would be to conduct a longitudinal study with MCI patients to test the sensitivity of the current test in detecting spatial memory deficits associated with Alzheimer's Disease, the appropriateness of the test with this patient group and to compare performance with validated cognitive tests currently used such as The Four Mountains Test, ACE-III, MOCA (De Roeck et al., 2019).

6.5 General summary and conclusions

I began this project with the intent to investigate allocentric representations through their limits of precision and capacity. In attempting to isolate distinct contributions to memory we have uncovered an interconnected system which integrates information from egocentric and allocentric representations in a way that is sensitive to task demands. Some of the limits of allocentric memory are unavoidably inherited or influenced by egocentric processes, but in carefully designed psychophysical experiments we can begin to distinguish and measure their separate contributions to spatial behaviour, disentangling the underlying mechanisms and learning more about their nature.

Appendices

Appendix A: Calculation of decibel values

The following function is used to convert distances in meters into equivalent decibel values:

$$\mathrm{C_{dB}} = 20 \times log_{10}(\frac{C}{C_0})$$

$\mathrm{C_{dB}}$ = spatial change (distortion) in decibels.

C = spatial change (distortion) in this trial in meters.

C_0 = 0.5 This is the decibel reference value (ie 0dB = 0.5m). We chose to use 0.5m as during the pilot testing this corresponds to an approximately average level of performance in the task.

Appendix B: Chapter 3 Individual participant Bayes Factors

Experiment 1: Bayes Factors

Participant	0°	5°	15°	45°	135°
1	1438.60	2.60	2.80	1.24	2.44
3	34.97	39.17	398.24	4.05	4.86
4	17.59	1132.60	187.60	8.30	1.58
5	44.63	8.69	21.47	3.15	20.94
6	321.09	16788.18	113.54	4992.62	283.73
7	206.45	37.22	65.57	1.43	6.00
8	40.06	11.85	1.26	6.08	3.62
10	52.44	524.06	9.67	2.78	3.55
12	379077.36	20222.15	5807.40	22.44	54.66
13	5.86	224.65	35.37	4.38	1.72

Experiment 2: Bayes Factors

Participant	0°	5°	15°	45°	135°
14	7372.55	1035.91	53.85	320.58	31.37
16	387.10	1071.41	685.07	11.52	2.43
17	251.09	140646.86	434.50	152.98	75.40
18	58090.18	9367.31	7399.87	41.55	123.90
20	4.02	245.11	1092.20	36.98	17.90
22	6404.80	208.10	28.47	15.04	21.42
25	3684.98	2929.09	718.48	7.21	276.54
26	89774.44	124531.04	226.03	120.42	60.58
27	5748.44	374.87	53420.57	3.68	12.73
29	55.77	36668.24	3002.64	30.06	110.69

Table A2 showing Bayes factors for each participant/condition in Experiment 1 and Experiment 2 respectively . These describe a likelihood ratio of two different models: a model described by the null hypothesis (model 0) - a constant level of performance unaffected by the amount of spatial distortion, and the experimental hypothesis (model 1) i.e., that the data are described by a psychophysical ogive. The Bayes Factor (BF_{10}) describes the likelihood of model 1/model 0 given the data. Kass and Rafferty (1995) classify the likelihood strength based on this value as follows, BF_{10} = 1-3.2 : 'Barely worth mentioning', 3.2-10 : 'Substantial', 10-100 : 'Strong', >100 : 'Decisive'.

Appendix C: Same/Different Task D prime analysis

In order to facilitate comparisons between results from the Same/Different task and the 2AFC task, we chose to calculate thresholds using proportion correct and fitting psychometric curves in both tasks (see methods/results sections). This is a fairly typical analysis for yes/no type experiments when used in conjunction with staircases. However, it involves calculating proportion correct by binning responses according to the current staircase level. This makes the treatment of Same trials somewhat ambiguous because there is no spatial distortion in the stimulus itself to make the participant aware of the current level of the active staircase in intermixed trials.

To account for possible response bias we calculated a d' value for each participant in each condition. This involved calculating a z-score across all of the false alarm target absent (Same) trials for a given viewpoint condition (around 50% of all trials), and subtracting this from the z-scores for the hit rates of each target present (Different) condition (e.g. staircase level of spatial distortion). This showed an appropriate range of d' scores over the testing range. Participants' d' scores were higher at higher staircase levels, meaning that the stimulus was easier to differentiate from noise when there was a larger spatial distortion. Consistent with our main analysis (see Figure 2), increasing viewpoint rotation shifted the d' functions rightwards, meaning that a greater spatial distortion was required to reach a given performance level.

We also calculated each participant's response criterion across the testing range as the mean of the summed z-scores ($C = -(z(hit) + z(FA))/2$). However, given that the false alarm rate is constant across all staircase levels (because we pooled data from all target absent trials), these functions are effectively inverted d' scores and are therefore not informative (see lower row of Figure S1). This is one reason why proportion correct values are typically presented when using staircase designs, as we have done here in our main analysis.

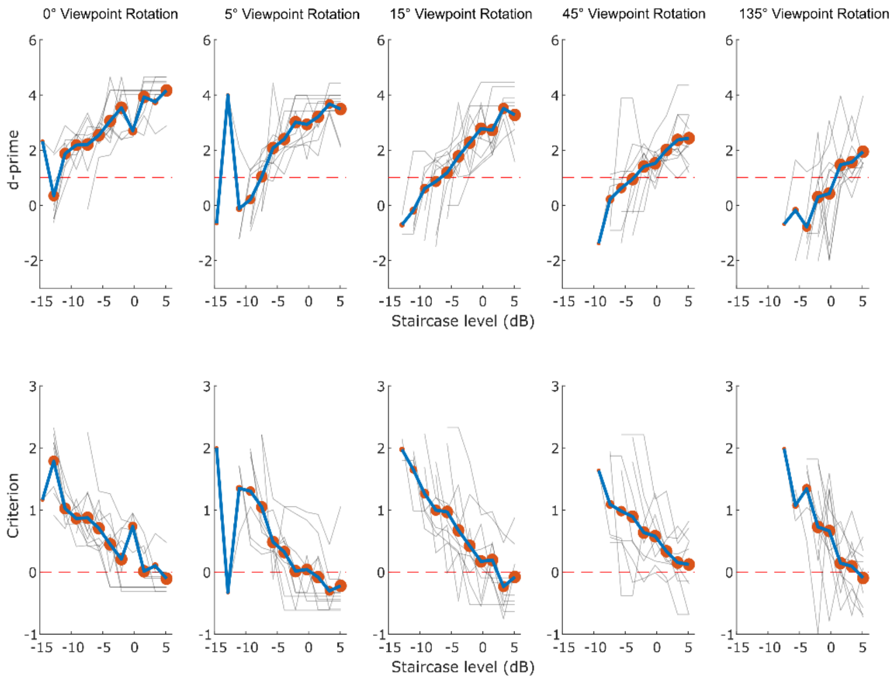

Figure A1 showing levels of d' (top) and response criterion (bottom) in each participant (faded grey line) . The circles show the mean level of d' and criterion respectively, averaged across all data points for each staircase level tested. Circle size represents number of data points. The bold line is a line plotted through these points.

Appendix D: Chapter 3 Staircase Parameters

Chapter 3 Experiment 1 Adaptive Staircase Information															
Step Criteria	3 up 1 down														
Number of Reversals to Complete	7 reversals														
Start Level	1x begin high														
	1 x begin low														
Steps sizes (change at each reversal)	6	4	2	1	1	1	1								
Levels (m)	0.05	0.06	0.08	0.09	0.11	0.14	0.17	0.21	0.26	0.32	0.39	0.48	0.6	0.73	0.9
Levels (dB)	-20	-18.2	-16.41	-14.62	-12.8	-11.03	-9.24	-7.45	-5.65	-3.86	-2.07	-0.27	1.52	3.31	5.11

Chapter 3 Experiment 2 Adaptive Staircase Information															
Step Criteria	3 up 1 down														
Number of Reversals to Complete	14 reversals														
Start Level	1x begin high														
Steps sizes (change at each reversal)	6	4	2	1	1	1	1								
Levels (m)	0.05	0.06	0.08	0.09	0.11	0.14	0.17	0.21	0.26	0.32	0.39	0.48	0.6	0.73	0.9
Levels (dB)	-20	-18.2	-16.41	-14.62	-12.8	-11.03	-9.24	-7.45	-5.65	-3.86	-2.07	-0.27	1.52	3.31	5.11

Table A3, Adaptive staircase parameters for experiments in Chapter 3 Experiment 1 (above) and Experiment 2 (below) . Step criteria refers to the sequence of correct and incorrect responses needed to change intensity. Number of reversals refers to how many direction changes required to finish the experiment. Step sizes refers to, when changing intensity, how many levels are skipped. Levels are reported in meters and corresponding decibels.

Appendix E: Chapter 4 Individual participant Bayes Factors

Experiment 4,1

Participant	Small, 0°	Large, 0°	Small, 135°	Large, 135°
1	11.13	13.32	31.99	6.79
2	457.25	488.19	4.04	3.83
3	4025.76	281.37	112.16	18.75
4	351.12	24.71	23.87	30.63
5	47.77	4054.38	4976.82	1.16
6	237.15	87.07	124.15	5.84
7	19.64	121.84	2.65	21.62
8	801.65	696.61	3.72	12.15
9	114.88	2.89	28.43	4.79
10	15.21	61.59	79.45	1.26
11	68350.55	595.43	245.58	3.07

Experiment 4,2

Participant	Small, 0°	Large, 0°	Small, 135°	Large, 135°
1	28.29	12.84	4.59	10.42
2	121.57	22.38	5.86	12.61
3	173.29	13.93	37.17	35.08
4	2.99	68.95	6.52	4.76
5	11.63	15.16	20.56	7.21
6	8.29	18.04	17.51	7.47
7	43.15	27.99	24.08	4.50
8	25.64	142.34	14.14	74.52
9	43.13	132.64	2.51	52.85
10	22.55	110.16	5.55	3.55
11	8.33	2.96	104.12	60.00
12	2939.91	7.14	19.96	17.92

Table A4 Bayes factors for each participant/condition in Experiment 4.1 and Experiment 4.2 respectively . These describe a likelihood ratio of two different models: a model described by the null hypothesis (model 0) - a constant level of performance unaffected by the amount of spatial distortion, and the experimental hypothesis (model 1) i.e., that the data are described by a psychophysical ogive. The Bayes Factor (BF_{10}) describes the likelihood of model 1/model 0 given the data. Kass and Rafferty (1995) classify the likelihood strength based on this value as follows, BF_{10} = 1-3.2 : 'Barely worth mentioning', 3.2-10 : 'Substantial', 10-100 : 'Strong', >100 : 'Decisive'.

Appendix F: Chapter 4 and 5 Staircase Parameters

Chapter 4 and 5 Adaptive Staircase Information															
Step Criteria	3 up 1 down														
Number of Reversals to Complete	14 reversals														
Start Level	1x begin high														
Steps sizes (change at each reversal)	6	4	2	1	1	1	1								
Levels (m)	0.06	0.07	0.09	0.10	0.12	0.14	0.16	0.18	0.21	0.25	0.29	0.34	0.40	0.46	…
Levels (dB)	-18.00	-16.67	-15.33	-14.00	-12.67	-11.33	-10.00	-8.67	-7.33	-6.00	-4.67	-3.33	-2.00	-0.67	…

…	0.54	0.63	0.73	0.86	1.00	1.16	1.36	1.58	1.84	2.15	2.51
…	0.67	2.00	3.33	4.67	6.00	7.33	8.67	10.00	11.33	12.67	14.00

Table A5, showing adaptive staircase parameters for experiments in Chapters 4 and 5 . Step criteria refers to the sequence of correct and incorrect responses needed to change intensity. Number of reversals refers to how many direction changes required to finish the experiment. Step sizes refers to, when changing intensity, how many levels are skipped. Levels are reported in meters and corresponding decibels.

References

Aagten-Murphy, D., & Bays, P. M. (2019). Independent working memory resources for egocentric and allocentric spatial information. *PLOS Computational Biology*, *15*(2), e1006563. https://doi.org/10.1371/journal.pcbi.1006563

Abrams, R. A., Meyer, D. E., & Kornblum, S. (1989). Speed and accuracy of saccadic eye movements: Characteristics of impulse variability in the oculomotor system. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(3), 529–543. https://doi.org/10.1037/0096-1523.15.3.529

Adamou, C., Avraamides, M. N., & Kelly, J. W. (2014). Integration of visuospatial information encoded from different viewpoints. *Psychon. Bull. Rev.*, *21*(3), 659–665. https://doi.org/10.3758/s13423-013-0538-5

Aguirre, G. K. (1999). Topographical disorientation: A synthesis and taxonomy. *Brain*, *122*(9), 1613–1628. https://doi.org/10.1093/brain/122.9.1613

Aguirre, G. K., & D'Esposito, M. (1999). Topographical disorientation: A synthesis and taxonomy. *Brain*, *122 (Pt 9)*, 1613–1628.

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychol. Sci.*, *15*(2), 106–111. https://doi.org/10.1111/j.0963-7214.2004.01502006.x

Amorim, M.-A., & Stucchi, N. (1997). Viewer- and object-centered mental explorations of an imagined environment are not equivalent. *Cognitive Brain Research*, *5*(3), 229–239. https://doi.org/10.1016/S0926-6410(96)00073-0

Andersen, R. A. (1995). Encoding of intention and spatial location in the posterior parietal cortex. *Cereb. Cortex*, *5*(5), 457–469.

Armbrüster, C., Wolter, M., Kuhlen, T., Spijkers, W., & Fimm, B. (2008). Depth Perception in Virtual Reality: Distance Estimations in Peri- and Extrapersonal Space. *CyberPsychology & Behavior*, *11*(1), 9–15. https://doi.org/10.1089/cpb.2007.9935

Astur, R. S., Taylor, L. B., Mamelak, A. N., Philpott, L., & Sutherland, R. J. (2002). Humans with hippocampus damage display severe spatial memory impairments in a virtual Morris water task. *Behav. Brain Res.*, *132*(1), 77–84.

Avillac, M., Denève, S., Olivier, E., Pouget, A., & Duhamel, J.-R. (2005). Reference frames for representing visual and tactile locations in parietal cortex. *Nat. Neurosci.*, *8*(7), 941–949. https://doi.org/10.1038/nn1480

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychol. Sci.*, *18*(7), 622–628. https://doi.org/10.1111/j.1467-9280.2007.01949.x

Baddeley, A. (1992). Working memory. *Science*, *255*(5044), 556–559.

Baddeley, A. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, *4*(11), 417–423. https://doi.org/10.1016/S1364-6613(00)01538-2

Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nat. Rev. Neurosci.*, *4*(10), 829–839. https://doi.org/10.1038/nrn1201

Baddeley, A. (2007). *Working Memory, Thought, and Action*. OUP Oxford.

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 8, pp. 47–89). Academic Press. https://doi.org/10.1016/S0079-7421(08)60452-1

Baker, D. H., Vilidaite, G., Lygo, F. A., Smith, A. K., Flack, T. R., Gouws, A. D., & Andrews, T. J. (2020). Power contours: Optimising sample size and precision in experimental psychology and human neuroscience. *Psychological Methods*, *26*(3), 295. https://doi.org/10.1037/met0000337

Baker, J. M., Bruno, J. L., Gundran, A., Hosseini, S. M. H., & Reiss, A. L. (2018). FNIRS measurement of cortical activation and functional connectivity during a visuospatial working memory task. *PLOS ONE*, *13*(8), e0201486. https://doi.org/10.1371/journal.pone.0201486

Barry, C., Ginzberg, L. L., O'Keefe, J., & Burgess, N. (2012). Grid cell firing patterns signal environmental novelty by expansion. *Proceedings of the National Academy of Sciences*, *109*(43), 17687–17692. https://doi.org/10.1073/pnas.1209918109

Bartsch, T., Döhring, J., Rohr, A., Jansen, O., & Deuschl, G. (2011). CA1 neurons in the human hippocampus are critical for autobiographical memory, mental time travel, and autonoetic consciousness. *Proc. Natl. Acad. Sci. U. S. A.*, *108*(42), 17562–17567. https://doi.org/10.1073/pnas.1110266108

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7.1-11. https://doi.org/10.1167/9.10.7

Bays, P. M., G., C., Raquel F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10). https://jov.arvojournals.org/article.aspx?articleid=2122354

Bays, P. M., & Husain, M. (2008). Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science*, *321*(5890), 851–854. https://doi.org/10.1126/science.1158023

Benhamou, S., Sauvé, J.-P., & Bovet, P. (1990). Spatial memory in large scale movements: Efficiency and limitation of the egocentric coding process. *Journal of Theoretical Biology*, *145*(1), 1–12. https://doi.org/10.1016/S0022-5193(05)80531-4

Berens, S. C., Richards, B. A., & Horner, A. J. (2020). Dissociating memory accessibility and precision in forgetting. *Nature Human Behaviour*, 1–12. https://doi.org/10.1038/s41562-020-0888-8

Bicanski, A., & Burgess, N. (2018). A neural-level model of spatial memory and imagery. *Elife*, *7*. https://doi.org/10.7554/eLife.33752

Bicanski, A., & Burgess, N. (2020). Neuronal vector coding in spatial cognition. *Nature Reviews Neuroscience*, *21*(9), 453–470. https://doi.org/10.1038/s41583-020-0336-9

Bird, C. M., Chan, D., Hartley, T., Pijnenburg, Y. A., Rossor, M. N., & Burgess, N. (2010). Topographical short-term memory differentiates Alzheimer's disease from frontotemporal lobar degeneration. *Hippocampus*, *20*(10), 1154–1169. https://doi.org/10.1002/hipo.20715

Bisiach, E., & Luzzatti, C. (1978). Unilateral neglect of representational space. *Cortex*, *14*(1), 129–133.

Bosch, L., & Sebastián-Gallés, N. (1997). Native-language recognition abilities in 4-month-old infants from monolingual and bilingual environments. *Cognition*, *65*(1), 33–69. https://doi.org/10.1016/S0010-0277(97)00040-1

Brady, T. F., & Alvarez, G. A. (2015). No evidence for a fixed object limit in working memory: Spatial ensemble representations inflate estimates of working memory capacity for complex objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 921–929. https://doi.org/10.1037/xlm0000075

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *J. Vis.*, *11*(5), 4. https://doi.org/10.1167/11.5.4

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proc. Natl. Acad. Sci. U. S. A.*, *105*(38), 14325–14329. https://doi.org/10.1073/pnas.0803390105

Brady, T. F., & Oliva, A. (2008). Statistical Learning Using Real-World Scenes: Extracting Categorical Regularities Without Conscious Intent. *Psychological Science*, *19*(7), 678–685. https://doi.org/10.1111/j.1467-9280.2008.02142.x

Brenner, E., & Van Damme, W. J. M. (1998). Judging distance from ocular convergence. *Vision Research*, *38*(4), 493–498. https://doi.org/10.1016/S0042-6989(97)00236-8

Brown, L. G. (1996). Additional rules for the transformed up-down method in psychophysics. *Perception & Psychophysics*, *58*(6), 959–962. https://doi.org/10.3758/BF03205497

Brun, V. H., Solstad, T., Kjelstrup, K. B., Fyhn, M., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, *18*(12), 1200–1212. https://doi.org/10.1002/hipo.20504

Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends Cogn. Sci.*, *11*(2), 49–57. https://doi.org/10.1016/j.tics.2006.11.004

Burgess, N. (2006). Spatial memory: How egocentric and allocentric combine. *Trends Cogn. Sci.*, *10*(12), 551–557. https://doi.org/10.1016/j.tics.2006.10.005

Burgess, N. (2008). *Spatial Cognition and the Brain*. *Annals of the New York Academy of Sciences*, *1124*(1), 77–97. https://doi.org/10.1196/annals.1440.002

Burgess, N., Becker, S., King, J. A., & O'Keefe, J. (2002). Memory for events and their spatial context: Models and experiments. In *Episodic Memory: New Directions in Research*.

Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198508809.003.0014

Burgess, N., Jeffery, K. J., & O'Keefe, J. (1999). *The hippocampal and parietal foundations of spatial cognition*. Oxford University Press.

Burgess, N., Jeffery, K., & O'Keefe, J. (1998). *The Hippocampal and Parietal Foundations of Spatial Cognition* (N. Burgess, K. Jeffery, & J. O'Keefe, Eds.). Oxford University Press. http://discovery.ucl.ac.uk/4521/

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The Human Hippocampus and Spatial and Episodic Memory. *Neuron*, *35*(4), 625–641. https://doi.org/10.1016/S0896-6273(02)00830-9

Burgess, N., Spiers, H. J., & Paleologou, E. (2004). Orientational manoeuvres in the dark: Dissociating allocentric and egocentric influences on spatial memory. *Cognition*, *94*(2), 149–166. https://doi.org/10.1016/j.cognition.2004.01.001

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychol. Rev.*, *114*(2), 340–375. https://doi.org/10.1037/0033-295X.114.2.340

Chafee, M. V., & Goldman-Rakic, P. S. (1998). Matching Patterns of Activity in Primate Prefrontal Area 8a and Parietal Area 7ip Neurons During a Spatial Working MemoryTask. *J. Neurophysiol.*, *79*(6), 2919–2940. https://doi.org/10.1152/jn.1998.79.6.2919

Chan, D., Gallaher, L. M., Moodley, K., Minati, L., Burgess, N., & Hartley, T. (2016). The 4 Mountains Test: A Short Test of Spatial Memory with High Sensitivity for the Diagnosis of Pre-dementia Alzheimer's Disease. *JoVE (Journal of Visualized Experiments)*, *116*, e54454. https://doi.org/10.3791/54454

Cheung, A., Stürzl, W., Zeil, J., & Cheng, K. (2008). The information content of panoramic images II: View-based navigation in nonrectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(1), 15–30. https://doi.org/10.1037/0097-7403.34.1.15

Clark, N., Binsted, G., Heath, M., & Krigolson, O. (2007). The proximity of visual landmarks impacts reaching performance. *Spatial Vision*, *20*(4), 317–336. https://doi.org/10.1163/156856807780919028

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, *386*(6625), 604–608. https://doi.org/10.1038/386604a0

Colby, C. L. (1998). Action-oriented spatial reference frames in cortex. *Neuron*, *20*(1), 15–24.

Committeri, G., Galati, G., Paradis, A.-L., Pizzamiglio, L., Berthoz, A., & LeBihan, D. (2004). Reference Frames for Spatial Cognition: Different Brain Areas are Involved in Viewer-, Object-, and Landmark-Centered Judgments About Object Location. *Journal of Cognitive Neuroscience*, *16*(9), 1517–1535. https://doi.org/10.1162/0898929042568550

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex*, *10*(9), 910–923.

Cornsweet, T. N. (1962). The Staircase-Method in Psychophysics. *The American Journal of Psychology*, *75*(3), 485. https://doi.org/10.2307/1419876

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*(1), 87–114. https://doi.org/10.1017/S0140525X01003922

Cowan, N., Elliott, E. M., Scott Saults, J., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. A. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, *51*(1), 42–100. https://doi.org/10.1016/j.cogpsych.2004.12.001

De Roeck, E. E., De Deyn, P. P., Dierckx, E., & Engelborghs, S. (2019). Brief cognitive screening instruments for early detection of Alzheimer’s disease: A systematic review. *Alzheimer’s Research & Therapy*, *11*(1), 21. https://doi.org/10.1186/s13195-019-0474-3

Deneve, S., & Pouget, A. (2004). Bayesian multisensory integration and cross-modal spatial links. *Journal of Physiology-Paris*, *98*(1), 249–258. https://doi.org/10.1016/j.jphysparis.2004.03.011

D’Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annu. Rev. Psychol.*, *66*, 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Dijkerman, H. C., Milner, A. D., & Carey, D. P. (1998). Grasping spatial relationships: Failure to demonstrate allocentric visual coding in a patient with visual form agnosia. *Conscious. Cogn.*, *7*(3), 424–437. https://doi.org/10.1006/ccog.1998.0365

Diwadkar, V. A., & McNamara, T. P. (1997). Viewpoint Dependence in Scene Recognition. *Psychol. Sci.*, *8*(4), 302–307. https://doi.org/10.1111/j.1467-9280.1997.tb00442.x

Doeller, C. F., King, J. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proc. Natl. Acad. Sci. U. S. A.*, *105*(15), 5915–5920. https://doi.org/10.1073/pnas.0801489105

Donkin, C., Kary, A., Tahir, F., & Taylor, R. (2016). Resources masquerading as slots: Flexible allocation of visual working memory. *Cogn. Psychol.*, *85*, 30–42. https://doi.org/10.1016/j.cogpsych.2016.01.002

Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychol. Rev.*, *120*(4), 873–902. https://doi.org/10.1037/a0034247

Donkin, C., Tran, S. C., & Nosofsky, R. (2014). Landscaping analyses of the ROC predictions of discrete-slots and signal-detection models of visual working memory. *Atten. Percept. Psychophys.*, *76*(7), 2103–2116. https://doi.org/10.3758/s13414-013-0561-7

Easton, R. D., & Sholl, M. J. (1995). Object-array structure, frames of reference, and retrieval of spatial knowledge. *J. Exp. Psychol. Learn. Mem. Cogn.*, *21*(2), 483–500. https://doi.org/10.1037//0278-7393.21.2.483

Eggleston, R. G., Janson, W. P., & Aldrich, K. A. (1996). Virtual reality system effects on size-distance judgements in a virtual environment. *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium*, 139–146. https://doi.org/10.1109/VRAIS.1996.490521

Ekstrom, A. D., Arnold, A. E. G. F., & Iaria, G. (2014). A critical review of the allocentric spatial representation and its neural underpinnings: Toward a network-based perspective. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00803

Ekstrom, A. D., Caplan, J. B., Ho, E., Shattuck, K., Fried, I., & Kahana, M. J. (2005). Human hippocampal theta activity during virtual navigation. *Hippocampus*, *15*(7), 881–889. https://doi.org/10.1002/hipo.20109

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184–188. https://doi.org/10.1038/nature01964

Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends Cogn. Sci.*, *12*(10), 388–396. https://doi.org/10.1016/j.tics.2008.07.004

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601. https://doi.org/10.1038/33402

Etienne, A. S., & Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus*, *14*(2), 180–192. https://doi.org/10.1002/hipo.10173

Evans, G. W., & Pezdek, K. (1980). Cognitive mapping: Knowledge of real-world distance and location information. *J. Exp. Psychol. Hum. Learn.*, *6*(1), 13–24.

Farrell, M. J., & Robertson, I. H. (1998). Mental rotation and automatic updating of body-centered spatial relationships. *J. Exp. Psychol. Learn. Mem. Cogn.*, *24*(1), 227.

Fechner, G. T. (1860). *Elemente der Psychophysik*. Breitkopf u. Härtel.

Fiehler, K., Wolf, C., Klinghammer, M., & Blohm, G. (2014). Integration of egocentric and allocentric information during memory-guided reaching to images of a natural environment. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.00636

Frankenstein, J., Mohler, B. J., Bülthoff, H. H., & Meilinger, T. (2012). Is the map in our head oriented north? *Psychol. Sci.*, *23*(2), 120–125. https://doi.org/10.1177/0956797611429467

Galati, G., Lobel, E., Vallar, G., Berthoz, A., Pizzamiglio, L., & Le Bihan, D. (2000). The neural basis of egocentric and allocentric coding of space in humans: A functional magnetic resonance study. *Exp. Brain Res.*, *133*(2), 156–164.

Galati, G., Pelle, G., Berthoz, A., & Committeri, G. (2010). Multiple reference frames used by the human brain for spatial perception and memory. *Experimental Brain Research*, *206*(2), 109–120. https://doi.org/10.1007/s00221-010-2168-8

Gallistel, C. R. (1990). *Chapter 14, The Organisation of Learning*.

Gallotta, M. C., Emerenziani, G. P., Monteiro, M. D., Iasevoli, L., Iazzoni, S., Baldari, C., & Guidetti, L. (2015). Psychophysical Benefits of Rock-Climbing Activity. *Perceptual and Motor Skills*, *121*(3), 675–689. https://doi.org/10.2466/30.PMS.121c26x9

Gao, Y., & Theeuwes, J. (2020). Learning to suppress a distractor is not affected by working memory load. *Psychonomic Bulletin & Review*, *27*(1), 96–104. https://doi.org/10.3758/s13423-019-01679-6

Gardner, J. L., Merriam, E. P., Movshon, J. A., & Heeger, D. J. (2008). Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. *J. Neurosci.*, *28*(15), 3988–3999. https://doi.org/10.1523/JNEUROSCI.5476-07.2008

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends Neurosci.*, *15*(1), 20–25.

Graziano, M. S. (1999). Where is my arm? The relative role of vision and proprioception in the neuronal representation of limb position. *Proc. Natl. Acad. Sci. U. S. A.*, *96*(18), 10418–10421.

Greenauer, N., & Waller, D. (2008). Intrinsic array structure is neither necessary nor sufficient for nonegocentric coding of spatial layouts. *Psychonomic Bulletin & Review*, *15*(5), 1015–1021. https://doi.org/10.3758/PBR.15.5.1015

Guariglia, C., Palermo, L., Piccardi, L., Iaria, G., & Incoccia, C. (2013). Neglecting the left side of a city square but not the left side of its clock: Prevalence and characteristics of representational neglect. *PLoS One*, *8*(7), e67390. https://doi.org/10.1371/journal.pone.0067390

Guo, K., & Li, C.-Y. (1997). Eye position-dependent activation of neurones in striate cortex of macaque. *Neuroreport*, *8*(6), 1405–1409.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. https://doi.org/10.1038/nature03721

Hägglund, M., Mørreaunet, M., Moser, M.-B., & Moser, E. I. (2019). Grid-Cell Distortion along Geometric Borders. *Current Biology*, *29*(6), 1047-1054.e3. https://doi.org/10.1016/j.cub.2019.01.074

Hartley, T., Bird, C. M., Chan, D., Cipolotti, L., Husain, M., Vargha-Khadem, F., & Burgess, N. (2005). The 4 mountains test: Hippocampal damage selectively impairs topographical processing of spatial scenes over very short delays. *Society for Neuroscience*, *191.14*. http://discovery.ucl.ac.uk/190507/

Hartley, T., Bird, C. M., Chan, D., Cipolotti, L., Husain, M., Vargha-Khadem, F., & Burgess, N. (2007). The hippocampus is required for short-term topographical memory in humans. *Hippocampus*, *17*(1), 34–48. https://doi.org/10.1002/hipo.20240

Hartley, T., & Harlow, R. (2012). An association between human hippocampal volume and topographical memory in healthy young adults. *Front. Hum. Neurosci.*, *6*, 338. https://doi.org/10.3389/fnhum.2012.00338

Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2014). Space in the brain: How the hippocampal formation supports spatial cognition. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, *369*(1635), 20120510. https://doi.org/10.1098/rstb.2012.0510

Hartley, T., Maguire, E. A., Spiers, H. J., & Burgess, N. (2003). The well-worn route and the path less traveled: Distinct neural bases of route following and wayfinding in humans. *Neuron*, *37*(5), 877–888.

Hartley, T., Trinkler, I., & Burgess, N. (2004). Geometric determinants of human spatial memory. *Cognition*, *94*(1), 39–75. https://doi.org/10.1016/j.cognition.2003.12.001

Hassabis, D., Kumaran, D., & Maguire, E. A. (2007). Using imagination to understand the neural basis of episodic memory. *J. Neurosci.*, *27*(52), 14365–14374. https://doi.org/10.1523/JNEUROSCI.4549-07.2007

Hassabis, D., & Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends Cogn. Sci.*, *11*(7), 299–306. https://doi.org/10.1016/j.tics.2007.05.001

Hassabis, D., & Maguire, E. A. (2009). The construction system of the brain. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, *364*(1521), 1263–1271. https://doi.org/10.1098/rstb.2008.0296

Hodgson, E., & Waller, D. (2006). Lack of set size effects in spatial updating: Evidence for offline updating. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(4), 854–866. https://doi.org/10.1037/0278-7393.32.4.854

Holdstock, J. S., Mayes, A. R., Cezayirli, E., Isaac, C. L., Aggleton, J. P., & Roberts, N. (2000). A comparison of egocentric and allocentric spatial memory in a patient with selective hippocampal damage. *Neuropsychologia*, *38*(4), 410–425. https://doi.org/10.1016/S0028-3932(99)00099-8

Hollingworth, A. (2006). Visual memory for natural scenes: Evidence from change detection and visual search. *Vis. Cogn.*, *14*(4–8), 781–807. https://doi.org/10.1080/13506280500193818

Holmes, C. A., Newcombe, N. S., & Shipley, T. F. (2018a). Move to learn: Integrating spatial information from multiple viewpoints. *Cognition*, *178*, 7–25. https://doi.org/10.1016/j.cognition.2018.05.003

Holmes, C. A., Newcombe, N. S., & Shipley, T. F. (2018b). Move to learn: Integrating spatial information from multiple viewpoints. *Cognition*, *178*, 7–25. https://doi.org/10.1016/j.cognition.2018.05.003

Holmes, M. C., & Sholl, M. J. (2005). Allocentric Coding of Object-to-Object Relations in Overlearned and Novel Environments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(5), 1069–1087. https://doi.org/10.1037/0278-7393.31.5.1069

Hölscher, C., Tenbrink, T., & Wiener, J. M. (2011). Would you follow your own route description? Cognitive strategies in urban route planning. *Cognition*, *121*(2), 228–247. https://doi.org/10.1016/j.cognition.2011.06.005

Hoydal, O. A., Skytoen, E. R., Moser, M.-B., & Moser, E. I. (2018). *Object-vector coding in the medial entorhinal cortex* [Preprint]. Neuroscience. https://doi.org/10.1101/286286

Interrante, V., Ries, B., & Anderson, L. (2006). Distance Perception in Immersive Virtual Environments, Revisited. *IEEE Virtual Reality Conference (VR 2006)*, 3–10. https://doi.org/10.1109/VR.2006.52

James, T. W., Humphrey, G. K., Gati, J. S., Menon, R. S., & Goodale, M. A. (2002). Differential effects of viewpoint on object-driven activation in dorsal and ventral streams. *Neuron*, *35*(4), 793–801. https://doi.org/10.1016/S0896-6273(02)00803-6

Jeffery, K. J., Page, H. J. I., & Stringer, S. M. (2016). Optimal cue combination and landmark-stability learning in the head direction system. *The Journal of Physiology*, *594*(22), 6527–6534. https://doi.org/10.1113/JP272945

Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of visual short-term memory. *J. Exp. Psychol. Learn. Mem. Cogn.*, *26*(3), 683–702.

Jung, M. W., Wiener, S. I., & McNaughton, B. L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *Journal of Neuroscience*, *14*(12), 7347–7356. https://doi.org/10.1523/JNEUROSCI.14-12-07347.1994

Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, *90*(430), 773–795. JSTOR. https://doi.org/10.2307/2291091

Kelly, J. W., & McNamara, T. P. (2010). Reference frames during the acquisition and development of spatial memories. *Cognition*, *116*(3), 409–420. https://doi.org/10.1016/j.cognition.2010.06.002

Kessels, R. P. C., Jaap Kappelle, L., de Haan, E. H. F., & Postma, A. (2002). Lateralization of spatial-memory processes: Evidence on spatial span, maze learning, and memory for object locations. *Neuropsychologia*, *40*(8), 1465–1473.

King, J. A., Burgess, N., Hartley, T., Vargha-Khadem, F., & O'Keefe, J. (2002). Human hippocampus and viewpoint dependence in spatial memory. *Hippocampus*, *12*(6), 811–820. https://doi.org/10.1002/hipo.10070

Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Finite scale of spatial representation in the hippocampus. *Science*, *321*(5885), 140–143. https://doi.org/10.1126/science.1157086

Klatzky, R. L. (1998). Allocentric and Egocentric Spatial Representations: Definitions, Distinctions, and Interconnections. In C. Freksa, C. Habel, & K. F. Wender (Eds.), *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing*

Spatial Knowledge (pp. 1–17). Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-69342-4_1

Klatzky, R. L., Beall, A. C., Loomis, J. M., Golledge, R. G., & Philbeck, J. W. (1999). *Human navigation ability: Tests of the encoding-error model of path integration*. 37.

Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010a). Conceptual Distinctiveness Supports Detailed Visual Long-Term Memory for Real-World Objects. *Journal of Experimental Psychology. General*, *139*(3), 558–578. https://doi.org/10.1037/a0019165

Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010b). Scene Memory Is More Detailed Than You Think: The Role of Categories in Visual Long-Term Memory. *Psychological Science*, *21*(11), 1551–1556. https://doi.org/10.1177/0956797610385359

Kowler, E., & Blaser, E. (1995). The accuracy and precision of saccades to small and large targets. *Vision Research*, *35*(12), 1741–1754. https://doi.org/10.1016/0042-6989(94)00255-K

Kravitz, D. J., Peng, C. S., & Baker, C. I. (2011). Real-world scene representations in high-level visual cortex: It's the spaces more than the places. *J. Neurosci.*, *31*(20), 7322–7333. https://doi.org/10.1523/JNEUROSCI.4588-10.2011

Kunz, L., Brandt, A., Reinacher, P. C., Staresina, B. P., Reifenstein, E. T., Weidemann, C. T., Herweg, N. A., Patel, A., Tsitsiklis, M., Kempter, R., Kahana, M. J., Schulze-Bonhage, A., & Jacobs, J. (2021). A neural code for egocentric spatial maps in the human medial temporal lobe. *Neuron*, *0*(0). https://doi.org/10.1016/j.neuron.2021.06.019

Lambrey, S., Doeller, C., Berthoz, A., & Burgess, N. (2012). Imagining Being Somewhere Else: Neural Basis of Changing Perspective in Space. *Cerebral Cortex*, *22*(1), 166–174. https://doi.org/10.1093/cercor/bhr101

Lawrence, M. A. (2016). *ez: Easy Analysis and Visualization of Factorial Experiments*. https://CRAN.R-project.org/package=ez

Learmonth, A. E., Nadel, L., & Newcombe, N. S. (2002). Children's Use of Landmarks: Implications for Modularity Theory. *Psychological Science*, *13*(4), 337–341. https://doi.org/10.1111/j.0956-7976.2002.00461.x

Lee, S. A., Winkler-Rhoades, N., & Spelke, E. S. (2012). Spontaneous Reorientation Is Guided by Perceived Surface Distance, Not by Image Matching Or Comparison. *PLOS ONE*, *7*(12), e51373. https://doi.org/10.1371/journal.pone.0051373

Leek, M. R. (2001). Adaptive procedures in psychophysical research. *Perception & Psychophysics*, *63*(8), 1279–1292. https://doi.org/10.3758/BF03194543

Lenck-Santini, P.-P., Muller, R. U., Save, E., & Poucet, B. (2002). Relationships between Place Cell Firing Fields and Navigational Decisions by Rats. *The Journal of Neuroscience*, *22*(20), 9035–9047. https://doi.org/10.1523/JNEUROSCI.22-20-09035.2002

Levi, D. M., & Klein, S. A. (1992). "Weber's law" for position: The role of spatial frequency and contrast. *Vision Research*, *32*(12), 2235–2250. https://doi.org/10.1016/0042-6989(92)90088-Z

Levi, D. M., Klein, S. A., & Yap, Y. L. (1988). "Weber's law" for position: Unconfounding the role of separation and eccentricity. *Vision Research*, *28*(5), 597–603. https://doi.org/10.1016/0042-6989(88)90109-5

Levvit, H. (1970). Transformed Up-Down Methods in Psychophysics. *The Journal of the Acoustical Society of America*, *49*(2), 467–477.

Link, S. W. (2020). *The Wave Theory of Difference and Similarity*. Routledge.

Logie, R. H. (2011). The Functional Organization and Capacity Limits of Working Memory. *Curr. Dir. Psychol. Sci.*, *20*(4), 240–245. https://doi.org/10.1177/0963721411415340

Loomis, J. M., Da Silva, J. A., Fujita, N., & Fukusima, S. S. (1992). Visual space perception and visually directed action. *J. Exp. Psychol. Hum. Percept. Perform.*, *18*(4), 906–921.

Luck, S. J., Hillyard, S. A., Mouloua, M., Woldorff, M. G., Clark, V. P., & Hawkins, H. L. (n.d.). *Effects of Spatial Cuing on Luminance Detectability: Psychophysical and Electrophysiological Evidence for Early Selection*. 18.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. https://doi.org/10.1038/36846

Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends Cogn. Sci.*, *17*(8), 391–400. https://doi.org/10.1016/j.tics.2013.06.006

Lyttle, D., Gereke, B., Lin, K. K., & Fellous, J.-M. (2013). Spatial scale and place field stability in a grid-to-place cell model of the dorsoventral axis of the hippocampus. *Hippocampus*, *23*(8), 729–744. https://doi.org/10.1002/hipo.22132

Ma, W. J., & Wilken, P. (2004). A signal detection account of visual short-term memory for orientation and spatial frequency. *J. Vis.*, *4*(8), 389–389. https://doi.org/10.1167/4.8.389

Maguire, E. A., Vargha-Khadem, F., & Hassabis, D. (2010). Imagining fictitious and future experiences: Evidence from developmental amnesia. *Neuropsychologia*, *48*(11), 3187–3192. https://doi.org/10.1016/j.neuropsychologia.2010.06.037

Makovski, T., Watson, L. M., Koutstaal, W., & Jiang, Y. V. (2010). Method matters: Systematic effects of testing procedure on visual working memory sensitivity. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *36*(6), 1466–1479. https://doi.org/10.1037/a0020851

Maloney, L. T., & Yang, J. N. (2003). Maximum likelihood difference scaling. *J. Vis.*, *3*(8), 573–585. https://doi.org/10.1167/3.8.5

McNamara, D. S., & Scott, J. L. (2001). Working memory capacity and strategy use. *Memory & Cognition*, *29*(1), 10–17. https://doi.org/10.3758/BF03195736

McNamara, T. P., Hardy, J. K., & Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *J. Exp. Psychol. Learn. Mem. Cogn.*, *15*(2), 211–227.

McNamara, T. P., Rump, B., & Werner, S. (2003). Egocentric and geocentric frames of reference in memory of large-scale space. *Psychon. Bull. Rev.*, *10*(3), 589–595. https://doi.org/10.3758/BF03196519

Meilinger, T., Riecke, B. E., & Bulthoff, H. H. (2007). *Orientation specificity in long-term-memory for environmental spaces*. *29*. https://cloudfront.escholarship.org/dist/prd/content/qt3jr3976f/qt3jr3976f.pdf

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81–97. https://doi.org/10.1037/h0043158

Miller, J. F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., Hefft, S., Merkow, M., Polyn, S. M., Jacobs, J., Kahana, M. J., & Schulze-Bonhage, A. (2013). Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science*, *342*(6162), 1111–1114. https://doi.org/10.1126/science.1244056

Milner, B. (1971). Interhemispheric differences in the localization of psychological processes in man. *Br. Med. Bull.*, *27*(3), 272–277.

Montello, D. R. (1993). Scale and multiple psychologies of space. In A. U. Frank & I. Campari (Eds.), *Spatial Information Theory A Theoretical Basis for GIS* (pp. 312–321). Springer. https://doi.org/10.1007/3-540-57207-4_21

Moodley, K., Minati, L., Contarino, V., Prioni, S., Wood, R., Cooper, R., D'Incerti, L., Tagliavini, F., & Chan, D. (2015). Diagnostic differentiation of mild cognitive impairment due to Alzheimer's disease using a hippocampus-dependent test of spatial memory. *Hippocampus*, *25*(8), 939–951. https://doi.org/10.1002/hipo.22417

Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods*, *11*(1), 47–60. https://doi.org/10.1016/0165-0270(84)90007-4

Morris, R. G. M., Schenk, F., Tweedie, F., & Jarrard, L. E. (1990). Ibotenate Lesions of Hippocampus and/or Subiculum: Dissociating Components of Allocentric Spatial Learning. *Eur. J. Neurosci.*, *2*(12), 1016–1028.

Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A., & Rosenbaum, R. S. (2006). The cognitive neuroscience of remote episodic, semantic and spatial memory. *Curr. Opin. Neurobiol.*, *16*(2), 179–190. https://doi.org/10.1016/j.conb.2006.03.013

Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., McAndrews, M. P., Levine, B., Black, S., Winocur, G., & Nadel, L. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *J. Anat.*, *207*(1), 35–66. https://doi.org/10.1111/j.1469-7580.2005.00421.x

Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.*, *31*, 69–89. https://doi.org/10.1146/annurev.neuro.31.061307.090723

Motes, M. A., Finlay, C. A., & Kozhevnikov, M. (2006). Scene recognition following locomotion around a scene. *Perception*, *35*(11), 1507–1520. https://doi.org/10.1068/p5459

Mou, W., & McNamara, T. P. (2002). Intrinsic frames of reference in spatial memory. *J. Exp. Psychol. Learn. Mem. Cogn.*, *28*(1), 162–170.

Mou, W., McNamara, T. P., Rump, B., & Xiao, C. (2006a). Roles of egocentric and allocentric spatial representations in locomotion and reorientation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(6), 1274–1290. https://doi.org/10.1037/0278-7393.32.6.1274

Mullally, S. L., Hassabis, D., & Maguire, E. A. (2012). Scene construction in amnesia: An FMRI study. *J. Neurosci.*, *32*(16), 5646–5653. https://doi.org/10.1523/JNEUROSCI.5522-11.2012

Mullally, S. L., & Maguire, E. A. (2014). Memory, Imagination, and Predicting the Future: A Common Brain Mechanism? *Neuroscientist*, *20*(3), 220–234. https://doi.org/10.1177/1073858413495091

Muller, R. U., Bostock, E., Taube, J. S., & Kubie, J. L. (1994). On the directional firing properties of hippocampal place cells. *J. Neurosci.*, *14*(12), 7235–7251.

Nakazawa, K., McHugh, T. J., Wilson, M. A., & Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience*, *5*(5), 361–372. https://doi.org/10.1038/nrn1385

Negen, J., Heywood-Everett, E., Roome, H. E., & Nardini, M. (2018). Development of allocentric spatial recall from new viewpoints in virtual reality. *Dev. Sci.*, *21*(1). https://doi.org/10.1111/desc.12496

Negen, J., Roome, H. E., Keenaghan, S., & Nardini, M. (2018). Effects of two-dimensional versus three-dimensional landmark geometry and layout on young children's recall of locations from new viewpoints. *J. Exp. Child Psychol.*, *170*, 1–29. https://doi.org/10.1016/j.jecp.2017.12.009

Neggers, S. F. W., Van der Lubbe, R. H. J., Ramsey, N. F., & Postma, A. (2006). Interactions between ego- and allocentric neuronal representations of space. *NeuroImage*, *31*(1), 320–331. https://doi.org/10.1016/j.neuroimage.2005.12.028

Newcombe, N., & Huttenlocher, J. (1992). Children's early ability to solve perspective-taking problems. *Developmental Psychology*, *28*(4), 635–643. https://doi.org/10.1037/0012-1649.28.4.635

O'Keefe, J., & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: Why they fire where they fire. *Exp. Brain Res.*, *31*(4), 573–590.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, *34*(1), 171–175.

O'Keefe, J., & Nadel, L. (1978a). *The hippocampus as a cognitive map*. Oxford: Clarendon Press. http://arizona.openrepository.com/arizona/handle/10150/620894

O'Keefe, J., & Nadel, L. (1978b). *The hippocampus as a cognitive map*. Oxford: Clarendon Press. http://arizona.openrepository.com/arizona/handle/10150/620894

O'Keefe, J., & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, *68*(1). https://doi.org/10.1007/BF00255230

Pardo-Vazquez, J. L., Castiñeiras-de Saa, J. R., Valente, M., Damião, I., Costa, T., Vicente, M. I., Mendonça, A. G., Mainen, Z. F., & Renart, A. (2019). The mechanistic foundation of Weber's law. *Nature Neuroscience*, *22*(9), 1493–1502. https://doi.org/10.1038/s41593-019-0439-7

Parslow, D. M., Rose, D., Brooks, B., Fleminger, S., Gray, J. A., Giampietro, V., Brammer, M. J., Williams, S., Gasston, D., Andrew, C., Vythelingum, G. N., Ioannou, G., Simmons, A., & Morris, R. G. (2004). Allocentric Spatial Memory Activation of the Hippocampal

Formation Measured With fMRI. *Neuropsychology*, *18*(3), 450–461. https://doi.org/10.1037/0894-4105.18.3.450

Pengas, G., Patterson, K., Arnold, R. J., Bird, C. M., Burgess, N., & Nestor, P. J. (2010). Lost and found: Bespoke memory testing for Alzheimer's disease and semantic dementia. *J. Alzheimers. Dis.*, *21*(4), 1347–1365.

Pessoa, L., Gutierrez, E., Bandettini, P., & Ungerleider, L. (2002). Neural correlates of visual working memory: FMRI amplitude predicts task performance. *Neuron*, *35*(5), 975–987.

Philbeck, J. W., & Loomis, J. M. (1997). Comparison of two indicators of perceived egocentric distance under full-cue and reduced-cue conditions. *J. Exp. Psychol. Hum. Percept. Perform.*, *23*(1), 72–85.

Pouget, A., & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *J. Cogn. Neurosci.*, *9*(2), 222–237. https://doi.org/10.1162/jocn.1997.9.2.222

Qi, X.-L., Katsuki, F., Meyer, T., Rawley, J. B., Zhou, X., Douglas, K. L., & Constantinidis, C. (2010). Comparison of neural activity related to working memory in primate dorsolateral prefrontal and posterior parietal cortex. *Front. Syst. Neurosci.*, *4*, 12. https://doi.org/10.3389/fnsys.2010.00012

Richardson, A. E., Montello, D. R., & Hegarty, M. (1999). Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Mem. Cognit.*, *27*(4), 741–750.

Riecke, B. E., Veen, H. A. H. C. van, & Bülthoff, H. H. (2002). Visual Homing Is Possible Without Landmarks: A Path Integration Study in Virtual Reality. *Presence:*

Teleoperators and Virtual Environments, *11*(5), 443–473. https://doi.org/10.1162/105474602320935810

Rieser, J. J., Ashmead, D. H., Talor, C. R., & Youngquist, G. A. (1990). Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception*, *19*(5), 675–689. https://doi.org/10.1068/p190675

Rieser, J. J., & Rider, E. A. (1991). Young children's spatial orientation with respect to multiple targets when walking without vision. *Developmental Psychology*, *27*(1), 97–107. https://doi.org/10.1037/0012-1649.27.1.97

Robertson, R. G., Rolls, E. T., & Georges-François, P. (1998). Spatial View Cells in the Primate Hippocampus: Effects of Removal of View Details. *Journal of Neurophysiology*, *79*(3), 1145–1156. https://doi.org/10.1152/jn.1998.79.3.1145

Rosenbaum, R. S., Ziegler, M., Winocur, G., Grady, C. L., & Moscovitch, M. (2004). ?I have often walked down this street before?: FMRI Studies on the hippocampus and other structures during mental navigation of an old environment. *Hippocampus*, *14*(7), 826–835. https://doi.org/10.1002/hipo.10218

Roskos-Ewoldsen, B., McNamara, T. P., Shelton, A. L., & Carr, W. (1998). Mental representations of large and small spatial layouts are orientation dependent. *J. Exp. Psychol. Learn. Mem. Cogn.*, *24*(1), 215–226.

Rouder, J. N., Morey, R. D., Cowan, N., Zwilling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of fixed-capacity models of visual working memory. *Proc. Natl. Acad. Sci. U. S. A.*, *105*(16), 5975–5979. https://doi.org/10.1073/pnas.0711295105

Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychon. Bull. Rev.*, *18*(2), 324–330. https://doi.org/10.3758/s13423-011-0055-3

Salinas, E., & Abbott, L. F. (2001). Coordinate transformations in the visual system: How to generate gain fields and what to compute with them. In *Progress in brain research* (Vol. 130, pp. 175–190). Elsevier. https://www.sciencedirect.com/science/article/pii/S0079612301300122

Sargent, J., Dopkins, S., Philbeck, J., & Chichka, D. (2010). Chunking in spatial memory. *J. Exp. Psychol. Learn. Mem. Cogn.*, *36*(3), 576–589. https://doi.org/10.1037/a0017528

Sargent, J., Dopkins, S., Philbeck, J., & Modarres, R. (2008). Spatial memory during progressive disorientation. *J. Exp. Psychol. Learn. Mem. Cogn.*, *34*(3), 602–615. https://doi.org/10.1037/0278-7393.34.3.602

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The future of memory: Remembering, imagining, and the brain. *Neuron*, *76*(4), 677–694. https://doi.org/10.1016/j.neuron.2012.11.001

Schindler, A., & Bartels, A. (2013). Parietal Cortex Codes for Egocentric Space beyond the Field of View. *Current Biology*, *23*(2), 177–182. https://doi.org/10.1016/j.cub.2012.11.060

Schmidt, D., Krause, B. J., Weiss, P. H., Fink, G. R., Shah, N. J., Amorim, M.-A., Müller, H.-W., & Berthoz, A. (2007). Visuospatial working memory and changes of the point of view in 3D space. *NeuroImage*, *36*(3), 955–968. https://doi.org/10.1016/j.neuroimage.2007.03.050

Schmiedek, F., Hildebrandt, A., Lövdén, M., Wilhelm, O., & Lindenberger, U. (2009). Complex span versus updating tasks of working memory: The gap is not that deep. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(4), 1089–1096. https://doi.org/10.1037/a0015730

Schöning, S., Engelien, A., Kugel, H., Schäfer, S., Schiffbauer, H., Zwitserlood, P., Pletziger, E., Beizai, P., Kersting, A., Ohrmann, P., Greb, R. R., Lehmann, W., Heindel, W., Arolt, V., & Konrad, C. (2007). Functional anatomy of visuo-spatial working memory during mental rotation is influenced by sex, menstrual cycle, and sex steroid hormones. *Neuropsychologia*, *45*(14), 3203–3214. https://doi.org/10.1016/j.neuropsychologia.2007.06.011

Schulze, K., & Koelsch, S. (2012). Working memory for speech and music. *Ann. N. Y. Acad. Sci.*, *1252*, 229–236. https://doi.org/10.1111/j.1749-6632.2012.06447.x

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2018). *Psychological Scaling Reveals a Single Parameter Framework For Visual Working Memory*. https://doi.org/10.1101/325472

Segen, V., Colombo, G., Avraamides, M., Slattery, T., & Wiener, J. M. (2021). Perspective taking and systematic biases in object location memory. *Attention, Perception, & Psychophysics*. https://doi.org/10.3758/s13414-021-02243-y

Sereno, M. I., Pitzalis, S., & Martinez, A. (2001). Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science*, *294*(5545), 1350–1354. https://doi.org/10.1126/science.1063695

Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cogn. Psychol.*, *43*(4), 274–310. https://doi.org/10.1006/cogp.2001.0758

Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior*, *6*(1), 156–163. https://doi.org/10.1016/S0022-5371(67)80067-7

Shepard, R. N., & Metzler, J. (1971). Mental Rotation of Three-Dimensional Objects. *Science*, *171*(3972), 701–703. https://doi.org/10.1126/science.171.3972.701

Sheth, B. R., & Shimojo, S. (2004). Extrinsic Cues Suppress the Encoding of Intrinsic Cues. *Journal of Cognitive Neuroscience*, *16*(2), 339–350. https://doi.org/10.1162/089892904322984616

Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., & Gerstner, W. (2009). Is there a geometric module for spatial orientation? Insights from a rodent navigation model. *Psychological Review*, *116*(3), 540–566. https://doi.org/10.1037/a0016170

Sholl, M. J. (1987). Cognitive maps as orienting schemata. *J. Exp. Psychol. Learn. Mem. Cogn.*, *13*(4), 615–628. https://doi.org/10.1037//0278-7393.13.4.615

Simons, D. J., & Wang, R. F. (1998). Perceiving Real-World Viewpoint Changes. *Psychol. Sci.*, *9*(4), 315–320. https://doi.org/10.1111/1467-9280.00062

Smith, D. M., & Mizumori, S. J. Y. (2006). Hippocampal place cells, context, and episodic memory. *Hippocampus*, *16*(9), 716–729. https://doi.org/10.1002/hipo.20208

Souman, J. L., Frissen, I., Sreenivasa, M. N., & Ernst, M. O. (2009). Walking Straight into Circles. *Current Biology*, *19*(18), 1538–1542. https://doi.org/10.1016/j.cub.2009.07.053

Spiers, H. J., Burgess, N., Maguire, E. A., Baxendale, S. A., Hartley, T., Thompson, P. J., & O'Keefe, J. (2001). Unilateral temporal lobectomy patients show lateralized topographical and episodic memory deficits in a virtual town. *Brain*, *124*(Pt 12), 2476–2489.

Standing, L. (1973). Learning 10,000 pictures. *Q. J. Exp. Psychol.*, *25*(2), 207–222. https://doi.org/10.1080/14640747308400340

Stein, J. F. (1989). Representation of egocentric space in the posterior parietal cortex. *Q. J. Exp. Physiol.*, *74*(5), 583–606.

Stürzl, W., Cheung, A., Cheng, K., & Zeil, J. (2008). The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(1), 1–14. https://doi.org/10.1037/0097-7403.34.1.1

Suchan, B., Botko, R., Gizewski, E., Forsting, M., & Daum, I. (2006). Neural substrates of manipulation in visuospatial working memory. *Neuroscience*, *139*(1), 351–357. https://doi.org/10.1016/j.neuroscience.2005.08.020

Swets, J. A. (1961). Detection theory and psychophysics: A review. *Psychometrika*, *26*(1), 49–63. https://doi.org/10.1007/BF02289684

Tamber-Rosenau, B. J., Fintzi, A. R., & Marois, R. (2015). Crowding in Visual Working Memory Reveals Its Spatial Resolution and the Nature of Its Representations. *Psychological Science*, *26*(9), 1511–1521. https://doi.org/10.1177/0956797615592394

Taube, J. S., Muller, R. U., & Ranck, J. B., Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.*, *10*(2), 420–435.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Tresch, M. C., Sinnamon, H. M., & Seamon, J. G. (1993). Double dissociation of spatial and object visual memory: Evidence from selective interference in intact human subjects. *Neuropsychologia*, *31*(3), 211–219.

Tu, S., Spiers, H. J., Hodges, J. R., Piguet, O., & Hornberger, M. (2017). Egocentric versus Allocentric Spatial Memory in Behavioral Variant Frontotemporal Dementia and

Alzheimer's Disease. *Journal of Alzheimer's Disease*, *59*(3), 883–892. https://doi.org/10.3233/JAD-160592

Tulving, E., & Markowitsch, H. J. (1998). Episodic and declarative memory: Role of the hippocampus. *Hippocampus*, *8*(3), 198–204. https://doi.org/10.1002/(SICI)1098-1063(1998)8:3<198::AID-HIPO2>3.0.CO;2-G

Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nat. Rev. Neurosci.*, *10*(11), 792.

Vargha-Khadem, F. (1997). Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory. *Science*, *277*(5324), 376–380. https://doi.org/10.1126/science.277.5324.376

Vicari, S., Bellucci, S., & Carlesimo, G. A. (2003). Visual and spatial working memory dissociation: Evidence from Williams syndrome. *Dev. Med. Child Neurol.*, *45*(4), 269–273.

Voss, J. L. (2009). Long-term associative memory capacity in man. *Psychon. Bull. Rev.*, *16*(6), 1076–1081. https://doi.org/10.3758/PBR.16.6.1076

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems ofp values. *Psychonomic Bulletin & Review*, *14*(5), 779–804. https://doi.org/10.3758/BF03194105

Waller, D., & Hodgson, E. (2006). Transient and enduring spatial representations under disorientation and self-rotation. *J. Exp. Psychol. Learn. Mem. Cogn.*, *32*(4), 867.

Wang, R. F., & Brockmole, J. R. (2003). Simultaneous spatial updating in nested environments. *Psychon. Bull. Rev.*, *10*(4), 981–986.

Wang, R. F., Crowell, J. A., Simons, D. J., Irwin, D. E., Kramer, A. F., Ambinder, M. S., Thomas, L. E., Gosney, J. L., Levinthal, B. R., & Hsieh, B. B. (2006). Spatial updating relies on an egocentric representation of space: Effects of the number of objects. *Psychon. Bull. Rev.*, *13*(2), 281–286.

Wang, R. F., & Simons, D. J. (1999). Active and passive scene recognition across views. *Cognition*, *70*(2), 191–210.

Wang, R. F., & Spelke, E. S. (2000). Updating egocentric representations in human navigation. *Cognition*, *77*(3), 215–250.

Wang, R., & Spelke, E. (2002). Human spatial representation: Insights from animals. *Trends Cogn. Sci.*, *6*(9), 376.

Weber, E. H. (1834). *De Pulsu, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae ...* C.F. Koehler.

Werner, S., & Schmidt, K. (1999). Environmental reference systems for large-scale spaces. *Spat. Cogn. Comput.*, *1*(4), 447–473. https://doi.org/10.1023/A:1010095831166

Whitaker, D., & Latham, K. (1997). Disentangling the Role of Spatial Scale, Separation and Eccentricity in Weber's Law for Position. *Vision Research*, *37*(5), 515–524. https://doi.org/10.1016/S0042-6989(96)00202-7

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. https://ggplot2.tidyverse.org

Wolbers, T., & Büchel, C. (2005). Dissociable retrosplenial and hippocampal contributions to successful formation of survey representations. *J. Neurosci.*, *25*(13), 3333–3340. https://doi.org/10.1523/JNEUROSCI.4705-04.2005

Wolbers, T., & Hegarty, M. (2010). What determines our navigational abilities? *Trends in Cognitive Sciences*, *14*(3), 138–146. https://doi.org/10.1016/j.tics.2010.01.001

Wolbers, T., & Wiener, J. M. (2014). Challenges for identifying the neural mechanisms that support spatial navigation: The impact of spatial scale. *Front. Hum. Neurosci.*, *8*, 571. https://doi.org/10.3389/fnhum.2014.00571

Wood, R. A., Moodley, K. K., Lever, C., Minati, L., & Chan, D. (2016). Allocentric Spatial Memory Testing Predicts Conversion from Mild Cognitive Impairment to Dementia: An Initial Proof-of-Concept Study. *Frontiers in Neurology*, *7*, 215. https://doi.org/10.3389/fneur.2016.00215

Wraga, M., Creem, S. H., & Proffitt, D. R. (2000). Updating displays after imagined object and viewer rotations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(1), 151–168. https://doi.org/10.1037/0278-7393.26.1.151

Wraga, M., Shephard, J. M., Church, J. A., Inati, S., & Kosslyn, S. M. (2005). Imagined rotations of self versus objects: An fMRI study. *Neuropsychologia*, *43*(9), 1351–1361. https://doi.org/10.1016/j.neuropsychologia.2004.11.028

Wystrach, A., & Graham, P. (2012). View-Based Matching Can Be More than Image Matching: The Importance of considering an Animal's Perspective. *I-Perception*, *3*(8), 547–549. https://doi.org/10.1068/i0542ic

Xu, Y., Regier, T., & Newcombe, N. S. (2017). An adaptive cue combination model of human spatial reorientation. *Cognition*, *163*, 56–66. https://doi.org/10.1016/j.cognition.2017.02.016

Yonelinas, A. P. (2002). The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language*, *46*(3), 441–517. https://doi.org/10.1006/jmla.2002.2864

Zaehle, T., Jordan, K., Wüstenberg, T., Baudewig, J., Dechent, P., & Mast, F. W. (2007). The neural basis of the egocentric and allocentric spatial frame of reference. *Brain Research*, *1137*, 92–103. https://doi.org/10.1016/j.brainres.2006.12.044

Zago, L., & Tzourio-Mazoyer, N. (2002). Distinguishing visuospatial working memory and complex mental calculation areas within the parietal lobes. *Neuroscience Letters*, *331*(1), 45–49. https://doi.org/10.1016/S0304-3940(02)00833-9

Zhang, H., Mou, W., & McNamara, T. P. (2011). Spatial updating according to a fixed reference direction of a briefly viewed layout. *Cognition*, *119*(3), 419–429. https://doi.org/10.1016/j.cognition.2011.02.006

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860

Zhao, M., & Warren, W. H. (2015). How You Get There From Here: Interaction of Visual Landmarks and Path Integration in Human Navigation. *Psychological Science*, *26*(6), 915–924. https://doi.org/10.1177/0956797615574952