THE NEURAL MICROCIRCUITRY SUPPORTING EVENT PROCESSING

Alice Laura Hickling

November 2024

Submitted for a PhD in Cognitive Neuroscience

University College London

Supervised by:

Professor Eleanor Maguire (Principal)

Professor Martina Callaghan (Subsidiary)

UCL

I, Alice Laura Hickling confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:

Date: 8th November 2024

Abstract

In my PhD, I studied how the layers of the medial prefrontal cortex (mPFC) and subfields of the hippocampus form mircrocircuits to support the perception and memory of everyday experiences. Addressing this question required the high resolution offered by 7 Tesla (7T) MRI, but also brought significant challenges. I assisted the Department's Physics Team in developing a new fMRI sequence which sought to minimise the high levels of distortion and signal drop out in mPFC. I then developed a preprocessing and analysis pipeline to detect layer- and subfield-specific multivoxel event representations. This included a tool for automating segmentation of hippocampal subfields. I applied these acquisition and analysis methods in two 7T fMRI experiments. The first examined participants while they recalled experiences from their recent and remote past. I detected representations of specific remote memories in the mPFC deep layers. This deep layer involvement suggests a role of feedback signalling from the mPFC during the remembering of remote events. In the second study, I sought to probe further the role of the mPFC in event processing. During scanning, participants watched short animated movies comprising sets of scenes viewed consecutively, that were either linked together to form meaningful events or were sets of unrelated scenes (unlinked). I found that all of the mPFC layers were engaged specifically when the events were meaningful, as were a number of hippocampal subfields, namely the dentate gyrus, Cornu Ammonis (CA)2/3, CA1 and the pre/parasubiculum. The involvement of all mPFC layers suggests that both feedforward and feedback signalling are at play during event perception. Overall, this work is starting to expose the neural microcircuitry that might ultimately lead to a full mechanistic understanding of how our experiences are initially processed, and how they often remain accessible to us in the weeks, months and years that follow.

Impact statement

Our waking hours are primarily spent experiencing present events, recalling past events, or imagining scenarios that may happen in the future. Even in sleep, we conjure up a range of, sometimes bizarre, events while we dream. Therefore, the ability to directly or mentally experience events is a fundamental aspect of human life. Understanding the neural mechanisms that support this capacity is a key issue in neuroscience.

Two brain regions, the hippocampus and the medial prefrontal cortex, are critical for event processing. Until recently, investigations were limited to the engagement of these broad brain areas, lacking the resolution to examine the finer microstructures, such as the cortical layers and hippocampal subfields. Advances in high-resolution 7 Tesla magnetic resonance imaging now allow for the exploration of event processing within these microstructures. The distinct microcircuits formed from cortical layers and hippocampal subfields support specific types of information processing. For instance, projections travelling between the deep layers of distant cortical areas support top-down information processing, such as the transmission of high-level beliefs to facilitate the rapid understanding of sensory information. Understanding the involvement of these microcircuits provides a deeper and more mechanistic understanding of event processing than ever before in human research.

Many neurological and psychiatric conditions compromise the hippocampus and medial prefrontal cortex. Furthermore, these pathologies are often highly localised within these regions during their early stages. Alzheimer's disease, for example, targets specific microstructures within the hippocampal region. By understanding how the microcircuitry supports normative event processing, we can better grasp how these circuits malfunction in disease states. This insight may eventually inform the development of targeted interventions and biomarkers for precise diagnoses.

This thesis contributes foundational knowledge about the brain's capacity to mentally represent and process events, focusing on the neural microcircuitry. This research area is still in its early stages, with few studies examining cognition in the hippocampus and medial prefrontal cortex with such a high degree of spatial resolution. Given the infancy of the research area, methods development was required, including the development of an open access tool for automatically delineating subfields of the hippocampus in MRI scans. I hope

that this thesis will provide both methodological resources and neuroscientific inspiration for future discoveries about the neural underpinnings of human cognition.

Research Paper Declaration Form

UCL Research Paper Declaration Form: referencing the doctoral candidate's own published work(s)

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2):

   (a) **What is the title of the manuscript?** Automated protocols for delineating human hippocampal subfields from 3 Tesla and 7 Tesla magnetic resonance imaging data

   (b) **Please include a link to or doi for the work:** https://doi.org/10.1002/hipo.23606

   (c) **Where was the work published?** Hippocampus online website and Hippocampus Volume34, Issue6, June 2024, Pages 302-308

   (d) **Who published the work?** Hippocampus

   (e) **When was the work published?** 09 April 2024

   (f) **List the manuscript's authors in the order they appear on the publication:** Alice L. Hickling, Ian A. Clark, Yan I. Wu, Eleanor A. Maguire

   (g) **Was the work peer reviewd?** Yes

   (h) **Have you retained the copyright?** No

   (i) **Was an earlier form of the manuscript uploaded to a preprint server (e.g. medRxiv)? If 'Yes', please give a link or doi** No

   If 'No', please seek permission from the relevant publisher and check the box next to the below statement:

   ☑ *I acknowledge permission of the publisher named under 1d to include in this thesis portions of the publication named as included in 1c.*

2. **For a research manuscript prepared for publication but that has not yet been published** (if already published, please skip to section 3):

   (a) **What is the current title of the manuscript?**

   (b) **Has the manuscript been uploaded to a preprint server 'e.g. medRxiv'?**

   If 'Yes', please please give a link or doi:

   (c) **Where is the work intended to be published?**

   (d) **List the manuscript's authors in the intended authorship order:**

(e) **Stage of publication:**

3. **For multi-authored work, please give a statement of contribution covering all authors** (if single-author, please skip to section 4): See below

4. **In which chapter(s) of your thesis can this material be found?** Chapter 4

e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

Candidate: Alice Hickling

Date: 30/10/2024

Supervisor/Senior Author signature (where appropriate): Eleanor Maguire

Date: 30/10/2024

Author contributions:
Conceptualization (AH, EM)
Data curation (AH, IC)
Formal analysis (AH, IC)
Funding acquisition (EM)
Investigation (AH, IC, YW)
Supervision (EM)
Visualization (AH)
Writing – original draft (AH, EM)
Writing – review & editing (AH, IC, YW, EM)

Acknowledgements

First and foremost, I wish to thank my principal supervisor, Eleanor Maguire. I am immensely grateful for the time and effort you have invested in me and for the support and encouragement that you have given me throughout my PhD journey. I have learned a huge amount under your guidance. I am also sincerely thankful to my subsidiary supervisor, Martina Callaghan, who, despite her many commitments, has always made time to patiently explain concepts in MR physics or discuss any other aspect of my work. I would also like to thank my Thesis Committee members, Peter Kok and Peter Zeidman, for their advice and insightful comments both during and outside of the thesis committee meetings. I am very grateful to Peter Kok for providing feedback on several of my thesis chapters and for welcoming me into his lab group meetings, where I learned a great deal about perception, memory, the hippocampus and a broad range of topics beyond my immediate research focus.

I owe thanks to numerous collaborators. I am especially grateful to Ian Clark, who provided guidance on many aspects of my PhD, including teaching me to manually segment hippocampal subfields. I also wish to thank my other friends from the Maguire Group, Yan Wu, Nic Alexander, Rob Seymour, Tom Miller and Anna Monk, for their insightful feedback on my work and for their camaraderie. Also, to the many attendees of the Kok Group lab meetings, whose stimulating discussions have enriched my perspective on various scientific topics. Thank you to the physics team, including Nadine Graedel, Martina Callaghan, Vahid Malekian and Oliver Josephs, for their hard work and commitment to the development of my 7T MRI sequences. Thanks to my fellow 7T colleagues at the FIL, notably Joost Haarsma and Oliver Warrington, have helped me to navigate various 7T fMRI analysis queries. Thank you to David Bradbury, Megan Creasy, Daniel Bates, and everyone involved in imaging support; to Ric Davis and the IT team; and to Kamlyn Ramkissoon, Maddy Scott, Sophie Berry, and Alison Sin, who have been helped me with various administrative tasks.

Finally, I feel very fortunate to have received such warm and loving support from my friends and family. Yan, I am so glad our paths crossed during our PhDs—your friendship means a great deal to me and I look forward to many more years of “Yalice”. To my partner, William, and my parents, Helen and Alastair, thank you for cheering me on through every high and helping me through every low. Last, but by no means least, a special thanks to my cat, Jasper,

who has been a great source of comfort and an exceptionally dedicated lap-warmer during the long thesis-writing days.

Table of contents

List of abbreviations

2D	two-dimensional
3D	three-dimensional
7T	7 Tesla
ASHS	Automatic Segmentation of Hippocampal Subfields
ASL	arterial spin labelling
AT	anterior temporal
B0	static magnetic field
BBR	boundary-based registration
BOLD	blood oxygen level dependent
CA	Cornu Ammonis
CBF	cerebral blood flow
CBV	cerebral blood volume
CSF	cerebrospinal fluid
DCM	dynamic causal modelling
DG	dentate gyrus
EEG	electroencephalography
EPI	echo planar imaging
fMRI	functional magnetic resonance imaging
FWHM	full-width at half maximum
GE	gradient echo
GLM	general linear model
GRAPPA	GeneRalized Autocalibrating Partial Parallel Acquisition
JLF	joint label fusion
HSG	Hippocampal Subfield Group
MEG	magnetoencephalography
mLTC	middle gyrus of the lateral temporal cortex
MNI	Montreal Neurological Institute
MP2RAGE	Magnetization Prepared 2 Rapid Acquisition Gradient Echo
mPFC	medial prefrontal cortex
MRI	magnetic resonance imaging
MT	magnetisation transfer
MVPA	multivoxel pattern analysis
NMI	normalised mutual information
PM	posterior medial
RF	radio frequency
ROI	region of interest
RSA	representational similarity analysis
sLTC	superior gyrus of the lateral temporal cortex
SNR	signal-to-noise ratio
SPACE	spin echo sequence with variable flip angles
SVM	support vector machine

TE	echo time
TR	repetition time
tSNR	temporal signal-to-noise ratio
VASO	vascular-space-occupancy
vmPFC	ventromedial prefrontal cortex

1. General Introduction

Our lives consist of a continuous series of events, ranging from the everyday brushing of one's teeth to the highly significant experience of submitting a PhD thesis. We weave these life events into the story of our lives. They shape our self-identities and world-views.

Rather than always experiencing events in the present moment, much of our time is also spent reflecting on, replaying and reimagining events from our past and imagining possible future scenarios. We can also imagine entirely fictitious events, from the fantastical stories of literature to the vivid and often bizarre narratives of dreams. This faculty to mentally visualise events that are not happening at that moment is important for survival. It allows us to learn from our past experiences and anticipate and prepare for future challenges. Given that much of our existence is spent visualising events, understanding the neural mechanisms that support this function is a fundamental question in neuroscience.

In Section 1.1, of this introduction I will first examine the psychological features of events. Subsequently, I will discuss neuroimaging studies that have identified a network of brain regions supporting event processing. Among these regions, the hippocampus and medial prefrontal cortex (mPFC) play important, cooperative roles in event processing. I will examine evidence from behavioural studies in patients with focal bilateral damage to the hippocampus or mPFC as well as neuroimaging studies in healthy people that reveal the roles of these two brain regions in event processing. Finally, I will discuss how, based on their respective roles, the hippocampus and mPFC can be organised into a functional hierarchy of event processing along three cognitive dimensions: time, specificity and executive control.

In Section 1.2, I will change gears slightly to examine the microcircuitry of the hippocampus and neocortex. I will explain the microcircuitry of the neocortex and how it is thought to support hierarchical information processing. This microcircuitry is formed from connections between the layers of hierarchically organised cortical areas as well as the hippocampal subfields.

In Section 1.3, I will present the aims and outline of my PhD thesis. The broad aim was to investigate the mapping of the hippocampal and neocortical microcircuitry onto functional hierarchies of event processing. To this end, I used ultra-high field 7 Tesla (7T) magnetic

resonance imaging (MRI) to examine brain activity in the cortical layers and hippocampal subfields during event processing.

1.1. Event processing

1.1.1. What is an event?

An important feature of events is that they consist of scene imagery. In this thesis, and in previous research, a scene is defined as a naturalistic, three-dimensional (3D) spatially-coherent representation of the world, typically including people, animals, objects and landmarks, and viewed from an egocentric (i.e., first person) perspective (Dalton et al., 2018; Maguire & Mullally, 2013). Scenes have been further defined as contexts that you could physically step into (e.g., a forest) or operate within (e.g., a desk area; Monk, Dalton, et al., 2021). They can be viewed in the real world or as two-dimensional (2D) representations of the real world, like photographs (Monk, Dalton, et al., 2021).

The relevance of real-world scene imagery during the perception of life events is intuitive. However, the significance of scenes in the retrieval and imagination of events requires further scrutiny. To investigate this, a set of studies asked a large sample of 217 participants to explicitly self-report their strategies for recalling past and imagining future events (Clark et al., 2020). They performed an autobiographical memory retrieval task, in which they recalled and described autobiographical memories of events from various time periods of their life (Levine et al., 2002). During the future imagination task, participants imagined and described possible future events (e.g., an event next weekend; Hassabis, Kumaran, Vann, et al., 2007). In each of these tasks, the use of visual scene imagery was the most commonly reported strategy. Moreover, the ability to mentally construct scene imagery was found to fully mediate performance in each of these tasks (Clark et al., 2019). By contrast, tasks requiring the perception and recall of abstract word pairs, which did not evoke mental images, relied more heavily on verbal strategies than visual imagery (Clark et al., 2020). This research highlights the fundamental role of scene imagery in internally generated mental events.

A second essential characteristic of events is their temporal dimension. Events have been defined as having beginning, middle, and end such that they can be described in a story-like narrative (Zacks & Tversky, 2001). To become an event, scene imagery must evolve, such that

the world it depicts moves. The people and animals that populate it will often perform actions, interacting with each other and elements of the environment. Like single scenes, events, can be viewed in the real world or as a 2D representations of the real world, in the form of movies.

1.1.2. Core network of brain regions

A widely distributed set of brain regions is involved in recalling life events from the past (Beaty et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Hassabis & Maguire, 2007; Maguire, 2001; McDermott et al., 2009; Schacter et al., 2007; Svoboda et al., 2006), imagining events in the future (Beaty et al., 2018; Buckner & Carroll, 2007; Hassabis, Kumaran, & Maguire, 2007; Hassabis & Maguire, 2007; Schacter, 2012; Schacter et al., 2007), and perceiving events in the present moment such as during spatial navigation (Bohbot et al., 2004; Doeller et al., 2010; Herweg & Kahana, 2018; Rodriguez, 2010; Spiers & Maguire, 2006; Spreng et al., 2009; Wolbers & Buchel, 2005; Wolbers et al., 2007). This core event processing network includes the hippocampus, mPFC, parahippocampal, entorhinal, retrosplenial, precuneus, lateral parietal, lateral prefrontal and lateral temporal cortices (Figure 1). This same set of brain areas is also active during the imagination of single scenes, providing further support for the close relationship between scene and event processing (Cukur et al., 2016; Dalton et al., 2018; Epstein, 2008; Hassabis & Maguire, 2007; Robin et al., 2018; Robin & Olsen, 2019; Zeidman, Mullally, et al., 2015).

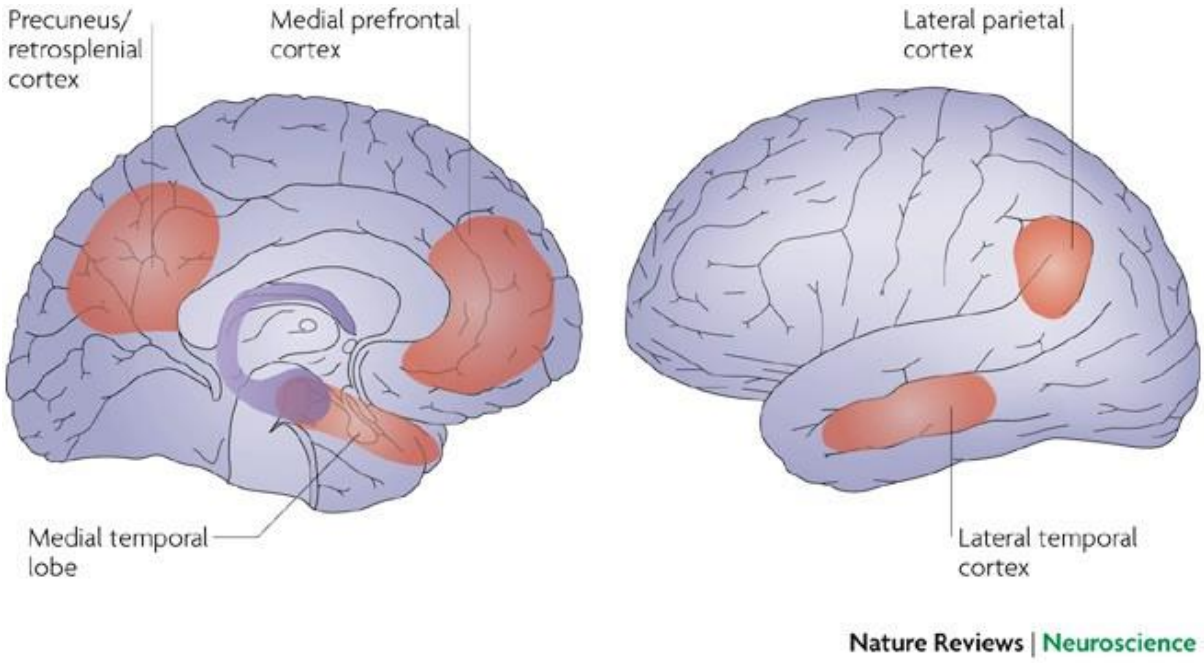

Figure 1. Core network of brain areas involved in event processing. A schematic showing the core set of brain areas that are engaged when perceiving present, retrieving past or imagining future events as well as during other forms of scene and event visualisation. These include the mPFC, the posterior medial (including the retrosplenial cortex and precuneus) and lateral parietal cortices, medial temporal cortex (including the hippocampus, entorhinal cortex and parahippocampal cortex), and lateral temporal regions. Figure from Schacter et al. (2007).

Two brain areas within the core network, the hippocampus and mPFC, have received particular attention for their distinct yet cooperative roles in event processing (McCormick, Ciaramelli, et al., 2018). In Sections 1.1.3 and 1.1.4, I will explore the existing literature that examines these roles.

1.1.3. The role of the hippocampus in event processing

1.1.3.1. The anatomy of the hippocampus

The hippocampus is a curved structure deep within the medial temporal lobe. It is composed of several distinct subfields, defined based on their cellular composition and organisation. They include the dentate gyrus (DG), Cornu Ammonis (CA)1, CA2, CA3, CA4, the subicular cortices, which include the subiculum, presubiculum and parasubiculum, and the uncus (Figure 2). Except the uncus, which is only in the most anterior portion of the hippocampus, all of the hippocampal subfields extend along the full anterior-to-posterior length of the

hippocampus. The hippocampal subfields have unique patterns of connectivity and cognitive functions (see Section 1.2.3; Coras et al., 2014; Zeidman, Lutti, et al., 2015).

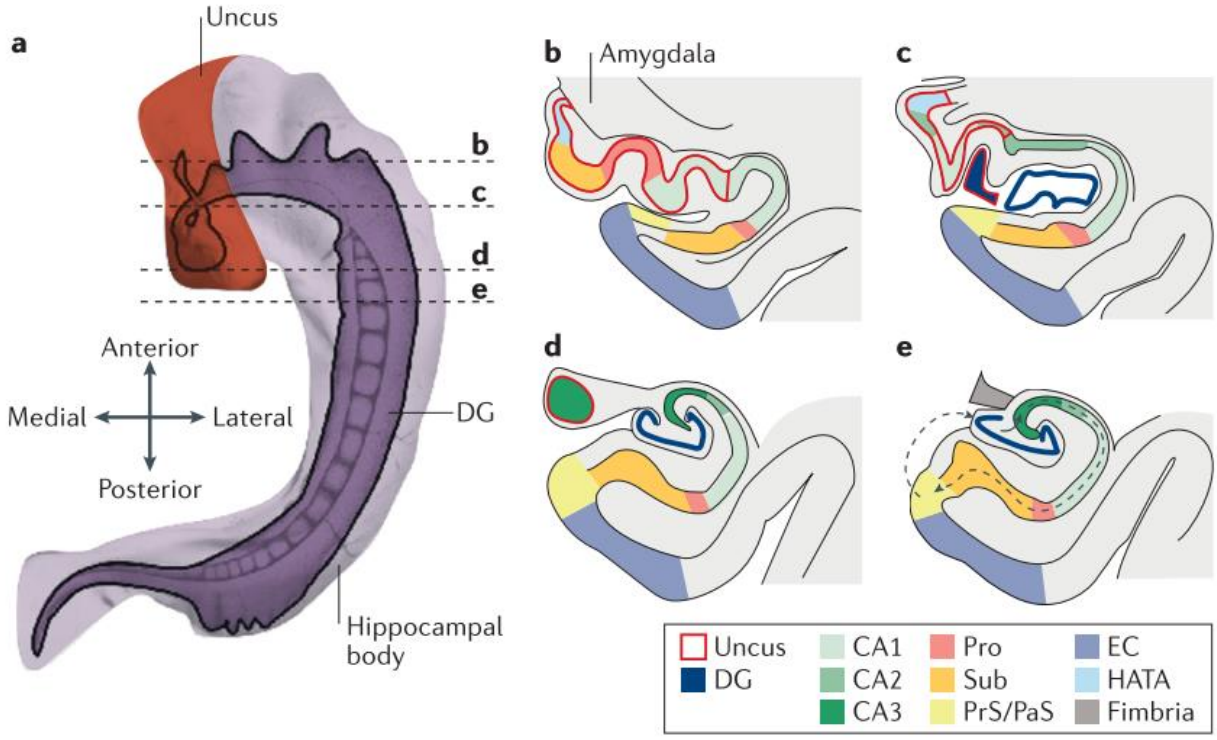

Figure 2. The anatomy of the hippocampus. A schematic showing (a) the whole hippocampus (purple) with the DG highlighted (dark purple). (b-e) Coronal slices at different points along the anterior-to-posterior axis, specified by dashed lines in (a). These slices show the hippocampal subfields and entorhinal cortex. DG = dentate gyrus, EC = entorhinal cortex; HATA = hippocampal–amygdaloid transition area; Pro = prosubiculum; PrS/PaS = presubiculum and parasubiculum. Figure from Zeidman & Maguire (2016).

1.1.3.2. The role of the hippocampus in memory over time

It has been known for decades that the hippocampus plays a critical role in the initial encoding of autobiographical events into memory (Addis et al., 2009; Addis, Wong, et al., 2007; Hainmueller & Bartos, 2020; Scoville & Milner, 1957). Furthermore, we known that, following their initial encoding, autobiographical memories are consolidated in neocortical areas in a process called systems-level consolidation (Barry et al., 2018; Frey & Morris, 1998; Marr, 1971; Redondo & Morris, 2011).

There are conflicting findings regarding the involvement of the hippocampus in recalling past autobiographical events after the memories have undergone systems level consolidation. One way to investigate this is by observing patients with bilateral hippocampal damage. There is consensus amongst patient studies that the recall of autobiographical events that occurred relatively recently before the hippocampal damage is significant impaired (Miller et al., 2020; Rosenbaum et al., 2008; Scoville & Milner, 1957; Steinvorth et al., 2005). However, while many hippocampal-damaged patients have similar impairments in retrieval of older memories (Cipolotti et al., 2001; Corkin, 2002; Miller et al., 2020; Nadel & Moscovitch, 1997; Rosenbaum et al., 2008; Steinvorth et al., 2005; Viskontas et al., 2000), others appear to have retained their ability to remember remote events (Kirwan et al., 2008; Scoville & Milner, 1957).

Another approach to investigate the involvement of the hippocampus in retrieval of remote and recent memories is by conducting neuroimaging studies in healthy participants. While some functional MRI (fMRI) studies have found higher hippocampal activation during retrieval of recent (less than 5 years ago) compared to remote (from childhood) autobiographical events (Gilmore et al., 2021; Piefke et al., 2003), many others have found similar hippocampal responses for both types of memory (Addis, Moscovitch, et al., 2007; Gilboa et al., 2004; Soderlund et al., 2012; Steinvorth et al., 2006; Viard et al., 2007) or higher activation during retrieval of remote compared to recent memories (Rekkas & Constable, 2005). Other fMRI studies have probed hippocampal representations of specific memory content, using an approach called multivoxel pattern analysis (MVPA). These studies have found stable representations of both recent and remote autobiographical memories in the hippocampus during their recall (Bonnici, Chadwick, Lutti, et al., 2012; Bonnici et al., 2013; Bonnici & Maguire, 2018), and stronger representations of individual remote memories were found in the CA3 and DG subfields (Bonnici et al., 2013), further indicating the sustained involvement of the hippocampus across time.

These inconsistent findings have given rise to many opposing theories of hippocampal function in the retrieval of memories over time (Barry & Maguire, 2019; Moscovitch et al., 2016; Nadel et al., 2007; Squire et al., 2015). Standard Consolidation Theory posits that memories are initially stored in the hippocampus and therefore the hippocampus is required for retrieval of recent memories (Figure 3A). However, once memories are sufficiently consolidated in the neocortex, they are no longer stored in the hippocampus and therefore

no longer require the hippocampus for retrieval (Squire, 1992; Squire & Alvarez, 1995). However, as discussed above, there is considerable evidence, from patient studies and neuroimaging studies in healthy individuals, for hippocampal involvement in retrieval of remote, consolidated memories.

Another theory, Multiple Trace Theory, proposes that the hippocampus continues to store memories after they have been consolidated in the neocortex, making it necessary for their retrieval across the lifespan (Figure 3B; Moscovitch et al., 2016; Nadel et al., 2007). This theory was later updated to Trace Transformation Theory, suggesting that the degree to which hippocampal memory representations are engaged during retrieval depends on the quality of the memory and the retrieval demands (Robin & Moscovitch, 2017; Sekeres et al., 2018). Memories that are recalled with rich episodic details are thought to always engage hippocampal memory traces. This theory would explain why the hippocampus appears to be involved in the retrieval of both recent and remote memories.

However, some flaws have been raised with Multiple Trace and Trace Transformation theories (Barry & Maguire, 2019). A given memory is thought to be stored in a set of neurons that contain strong synaptic connections with one another. This means that, when one cell fires, the full ensemble of neurons also fire, causing the full memory trace to activate. Rodent research suggests that there is a high turnover of the neurons and synapses that form these memory traces in the hippocampus (e.g., Attardo et al., 2015; Spalding et al., 2013). For example, it has been observed that dendritic spines, which contain synapses, only last for short periods before degenerating (Attardo et al., 2015), and full neurons are periodically lost to give way to newly born neurons, in a process known as adult neurogenesis (Spalding et al., 2013). Therefore, the prolonged existence of memory traces in the hippocampus is unlikely. Based on these observations, a third theory, Scene Construction Theory, posits that remote autobiographical memories are indeed no longer stored in the hippocampus, as suggested by Standard Consolidation Theory, but the hippocampus is nevertheless required to reconstruct them during retrieval, allowing events to be vividly reexperienced in the mind's eye (Figure 3C; Barry & Maguire, 2019). This may explain why the hippocampus is also involved in other cognitive processes that involve the vivid visualisation of events. In the next section, I will explore this idea further by reviewing the involvement of the hippocampus in other cognitive

processes, outside of autobiographical memory retrieval, that also require the (re)construction of events.

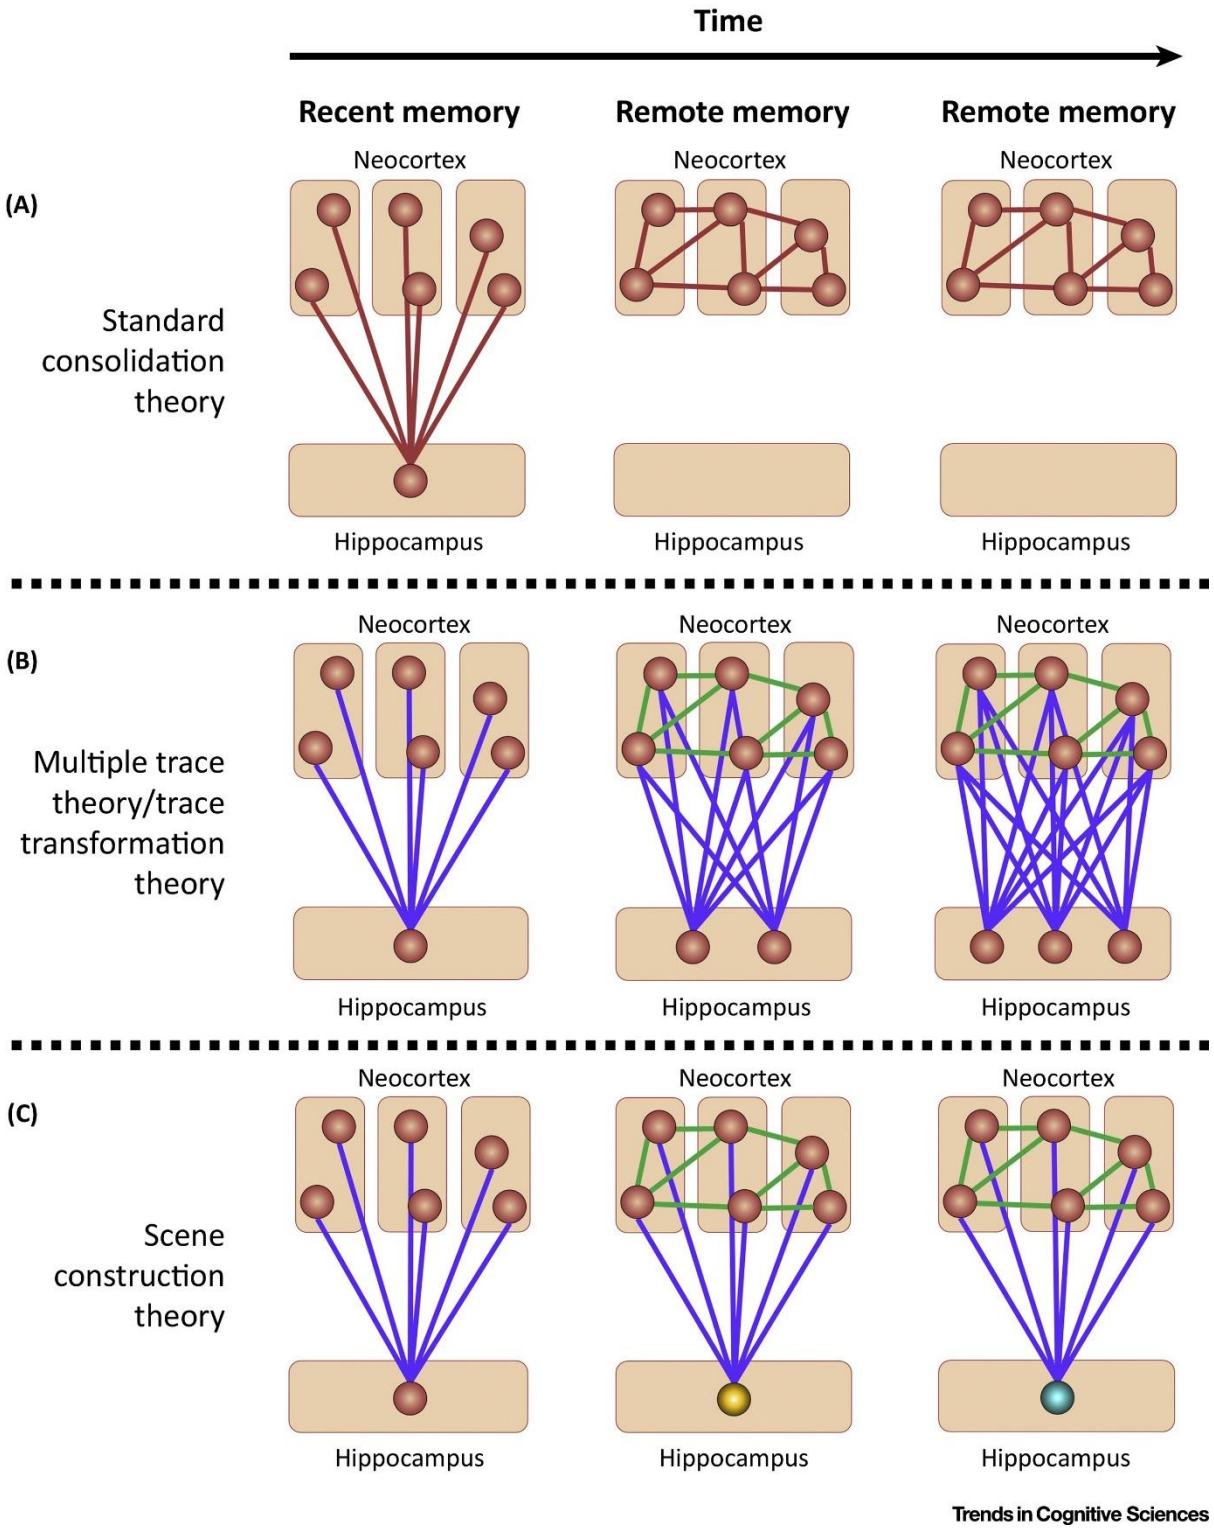

Figure 3. Theories of systems-level consolidation. (A) Standard consolidation theory posits that an episodic memory is initially stored as a joint hippocampal-neocortical trace but, over a period of

consolidation, transitions to a purely neocortical trace, losing its hippocampal element. (B) Multiple trace theory and trace transformation theory posit that an episodic memory is initially stored in a hippocampal-neocortical trace, and over time, the neocortical trace is strengthened to produce a more schematic and semantic memory. Also, with each repeated retrieval, a new hippocampal trace is additionally encoded. (C) Scene construction theory proposes that an episodic memory is initially encoded by a hippocampal-neocortical trace but with consolidation the hippocampal trace degrades rapidly and the neocortical trace is strengthened. Also, with each repeated retrieval, the elements of the neocortical memory trace are reconstructed to form a temporary hippocampal memory trace, which functions to construct the series of scenes that make up the episode. Figure from Barry & Maguire (2019).

1.1.3.3. Event processing beyond memory

Neuroimaging studies have reported increased activation in the hippocampus during a broad range of cognitive functions beyond autobiographical memory, including during the perception (Aly et al., 2013; Graham et al., 2010; Lee et al., 2013; Lee et al., 2005; Lee et al., 2012; McCormick et al., 2021; Read et al., 2024; Zeidman, Mullally, et al., 2015) and imagination of events (Addis, Moscovitch, et al., 2007; Addis et al., 2009; Gaesser et al., 2013; Hassabis, Kumaran, & Maguire, 2007) as well as during spatial navigation (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978).

Furthermore, patient studies have shown that individuals with hippocampal damage struggle to mentally project themselves into future events, providing fewer details of imagined future events compared to healthy controls (Addis & Schacter, 2011; Hassabis, Kumaran, Vann, et al., 2007; Klein et al., 2005; Kurczek et al., 2015; but also see Squire et al., 2010). Hippocampal-damaged patients also report fewer dreams during sleep and the few dreams that they have lack episodic content (Spano et al., 2020). They also spend less time day-dreaming about past and future events and their day-dreams are more abstract and verbal in content (McCormick, Rosenthal, et al., 2018). Together, these findings suggest that the hippocampus plays an important role in mentally visualising past, future and imagined events, whether during wakefulness or sleep.

1.1.3.4. Scene construction

Evidence has accumulated to support a fundamental role of the hippocampus in processing scene imagery (reviewed in Hassabis & Maguire, 2007; Maguire & Mullally, 2013). For instance, the hippocampus produces greater responses to scenes compared to other types of stimuli. During perception, the hippocampus activates more to the presentation of single

scenes compared to single objects (Brandman & Peelen, 2017; Graham et al., 2010; Hodgetts et al., 2016; Hodgetts et al., 2017; Read et al., 2024; Zeidman, Mullally, et al., 2015), faces (Hodgetts et al., 2017; Read et al., 2024), scrambled images (McCormick et al., 2021) and geometric shapes (Read et al., 2024). Likewise, during memory recall and imagination, the hippocampus is more engaged when visualising single scenes compared to scrambled words (Summerfield et al., 2010) and objects (Zeidman, Mullally, et al., 2015).

Further supporting a role of the hippocampus in scene processing, scene-selective impairments of hippocampal damaged patients are observed during perception, memory and imagination (reviewed in Graham et al., 2010; Lee et al., 2012; McCormick, Ciaramelli, et al., 2018). Patients are unable to imagine single scenes, but can imagine single objects (Hassabis, Kumaran, Vann, et al., 2007). They are unable to visualise anything that might exist beyond the borders of a presented scene, but can describe possible semantic, conceptual, and contextual details of an extended scene (Mullally et al., 2012). They perform poorly when tasked with recognising previously presented scenes from their memory, but can recognise faces (Bird et al., 2008; Taylor et al., 2007). Their day-dreams lack vivid scene imagery, but remain rich in conceptual and semantic thoughts (McCormick, Rosenthal, et al., 2018). Together, these findings suggest that the internal representation of scenes relies heavily on the hippocampus when compared to representations of other stimuli, such as objects, faces, words and concepts.

One study probed the specific role of the hippocampus in scene processing, revealing an important role in their spatial construction, that is, the spatial arrangement of individual elements (e.g., people, objects and landmarks) to internally visualise a scene (McCormick et al., 2017). McCormick et al. (2017) found that patients were impaired at detecting implausible scenes when the violations were constructive in nature (e.g., an endless staircase). By contrast, they were able to detect semantically impossible scenes (e.g., an elephant with butterflies for ears). The detection of constructive violations relies on the ability to internally construct and compare multiple versions of a scene. Patients have also been found to perform poorly in a perceptual discrimination task that similarly requires the internal construction and comparison of scenes (Barense et al., 2010; Lee et al., 2013; Lee et al., 2005; but also see Kim et al., 2011). In this task, an odd scene must be identified in a group of scenes shown from different perspectives. It requires the ability to mentally reconstruct a given presented scene

from various different perspectives, in order to compare it to the other presented scenes (Zeidman & Maguire, 2016). In contrast, patients are able to discriminate between different views of non-scene stimuli such as objects and faces, indicating that their ability to visualise and mentally rotate non-scene images is intact.

Further research, using this same scene discrimination paradigm, found that the subicular cortex, which included the subiculum, presubiculum and parasubiculum, was the only hippocampal region to show this scene-selective construction effect (Hodgetts et al., 2017). This finding supports other neuroimaging studies which showed that the medial-most portion of the hippocampus, an area that corresponds with the position of the pre/parasubiculum, is engaged during the perception and imagination of individual scenes (Dalton & Maguire, 2017; Dalton et al., 2018; Zeidman, Lutti, et al., 2015) and during memory retrieval of autobiographical events (Leelaarporn et al., 2024). Furthermore, the pre/parasubiculum is functionally connected to the parahippocampal cortex and entorhinal cortex during perception and memory of single scenes (Grande et al., 2022) and to the mPFC during the recollection of autobiographical memories (Leelaarporn et al., 2024). Therefore, the pre/parasubiculum has been proposed as a crucial hub for scene-based cognition (Dalton & Maguire, 2017).

There is also considerable evidence that hippocampal subfields other than the pre/parasubiculum are involved in processing scenes and spatial information more generally. For example, another study using the scene discrimination task described above, found a scene-selective effect in the inferior portion of CA1 (Read et al., 2024). Furthermore, studies in rodents and non-human primates have identified neurons that fire when an animal is in a specific location in space, place cells, in CA1, CA2, CA3 and the DG (Leutgeb et al., 2007; Oliva et al., 2016) and neurons that fire when an animal looks at a specific part of the environment, spatial view cells, in CA1 and CA3 (Robertson et al., 1998; Rolls, 2023). Moreover, a range of subfields have been identified as important for the recollection of past events including the DG (Bonnici et al., 2013; Chadwick et al., 2014), CA3 (Bonnici et al., 2013; Chadwick et al., 2014; Miller et al., 2020), CA1 (Bartsch et al., 2011) and the subicular cortices (Barry et al., 2021; Leelaarporn et al., 2024; Palombo et al., 2018).

In summary, evidence points to a role for the hippocampus in constructing mental imagery, particularly in the form of scene representations, which is fundamental to event processing.

This evident preference for constructing scenes led to the scene construction theory of the hippocampus (Hassabis & Maguire, 2007; Maguire & Mullally, 2013). Many other theories of hippocampal function exist (e.g., Behrens et al., 2018; Hassabis & Maguire, 2007; Maguire & Mullally, 2013; Mayes et al., 2007; Mayes & Roberts, 2001; Olsen et al., 2012; Schiller et al., 2015; Turk-Browne, 2019). One common theme amongst these theories is the importance of the hippocampus for associating multiple elements, and scenes may be just one example of stimuli that require this function. In this thesis, I seek to understand the role of the hippocampus in event processing. Its involvement in scene construction is highly relevant due to the integral nature of scene imagery in events.

1.1.4. The role of the mPFC in event processing

The anatomy of the mPFC

Along with the hippocampus, the mPFC also consistently engages during tasks requiring the processing of events (Benoit et al., 2014; D'Argembeau, 2013; Lieberman et al., 2019; Maguire, 2001; McDermott et al., 2009; Roy et al., 2012; Spreng et al., 2009; Svoboda et al., 2006). The mPFCis a broad region that encompasses several Brodmann areas: 10, 12, 25, as well as parts of 32, 9 and 11 (Figure 4; Catani et al., 2012). The area of activation during autobiographical memory retrieval is large, spanning most of these Brodmann areas (Figure 5; Roy et al., 2012).

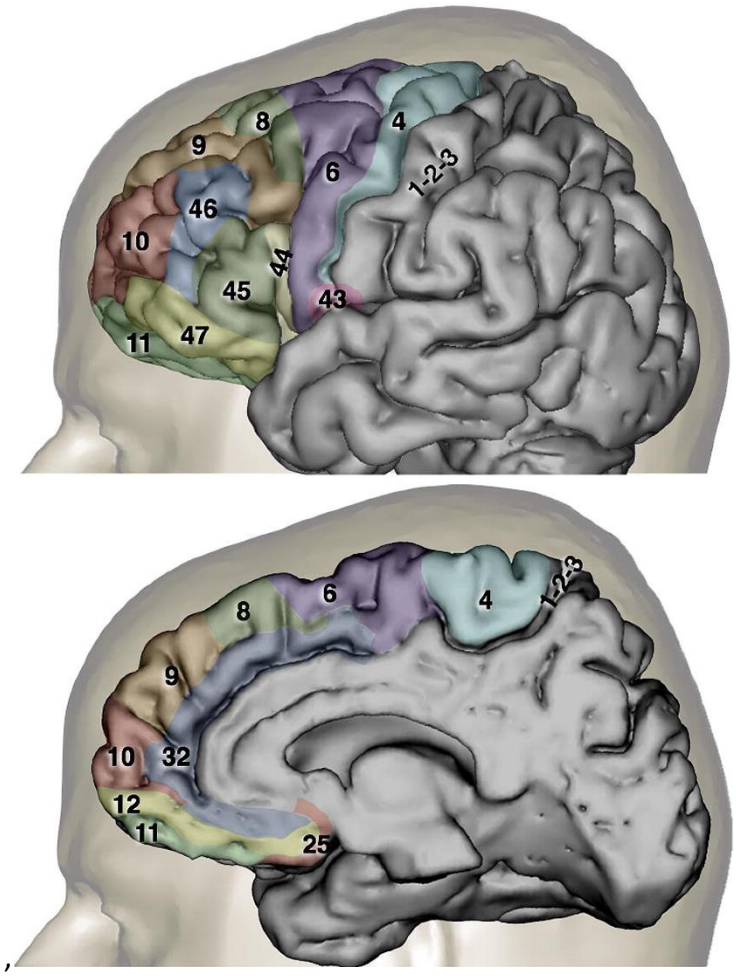

Figure 4. The anatomy of the frontal lobe. The parcellation of the frontal lobe into Brodmann areas, including the areas of the mPFC: 10, 12, 25, as well as parts of 32, 9 and 11. Figure from Catani et al. (2012).

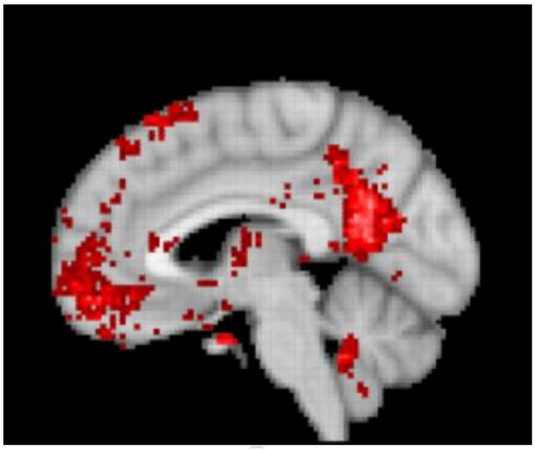

Figure 5. Functional MRI activation during autobiographical memory tasks. Based on a meta-analysis of neurosynth.org database of 84 fMRI studies whose abstracts include the term "autobiographical memory" once or more. There is strong and vast activation in the mPFC spanning several Brodmann areas. Figure produced using neurosynth.org.

The role of the mPFC in memory over time

As described previously, it is widely agreed that memories are stored in cortical regions in the longer term (Figure 3; Barry & Maguire, 2019; Moscovitch et al., 2016; Squire, 1992; Squire et al., 2015). Unlike in the hippocampus, in the neocortex, dendritic spines can last throughout a lifetime providing capacity for the long-term storage of memory traces (Yang et al., 2009). Although, many spines are pruned in order to forget information that is no longer useful (reviewed in Faust et al., 2021).

Functional MRI studies using MVPA methods have identified the mPFC as a potential storage site for remote autobiographical memories. They found stable multivoxel representations of individual remote (4 month – 10 year old), but not recent (<2 week old), autobiographical memories during retrieval (Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici et al., 2013; Bonnici & Maguire, 2018). However, although the mPFC preferentially represents older memories, it shows similar levels of univariate activation during the retrieval of autobiographical memories of all ages (Barry et al., 2018; McCormick et al., 2020). Furthermore, patients with mPFC damage are impaired in retrieval of autobiographical memories of all ages (Bertossi et al., 2016; Della Sala et al., 1993; Kopelman et al., 1999). This suggests a broader involvement of the mPFC in retrieval of autobiographical memories across time.

Interactions between the mPFC and hippocampus appear to be relevant for retrieval of both recent and remote autobiographical memory retrieval. Fuentemilla et al. (2014) found, using magnetoencephalography (MEG), strong synchronisation of theta power between the hippocampus and mPFC during recall of autobiographical memories (between 2 and 7 months old), indicating functional connectivity between the two areas during autobiographical memory retrieval (Fuentemilla et al., 2014). This synchrony did not exist during retrieval of semantic memories. Further investigations used dynamic causal modelling (DCM; Friston et al., 2003), a method that has been applied to electroencephalography (EEG), MEG and fMRI to test biologically plausible models of effective (i.e., causal) connectivity between brain regions, for example, the direction of influence that one region has on another. DCM studies found that activity in the mPFC drives activity in the hippocampus during the initial retrieval as well as the full elaborated recollection of past autobiographical events, regardless of their age (McCormick et al., 2020; Nawa & Ando, 2019, 2020).

Another approach used to infer the direction of connectivity between regions is using neuroimaging modalities such as MEG and EEG, which have high temporal resolution, to investigate the relative timing of responses in different brain regions (Figure 6; McCormick et al., 2020). McCormick et al. (2020) adopted this approach finding that, although the mPFC drove hippocampal activity during retrieval across all memories, it responded earlier than the hippocampus specifically during the retrieval of remote memories (4 months to 5 years old). In contrast, there was no difference in response time between when recent (less than 1 month old) memories were recalled. This suggests that although the mPFC is involved in retrieving memories of all ages, it may initiate retrieval only when remote memories are being retrieved.

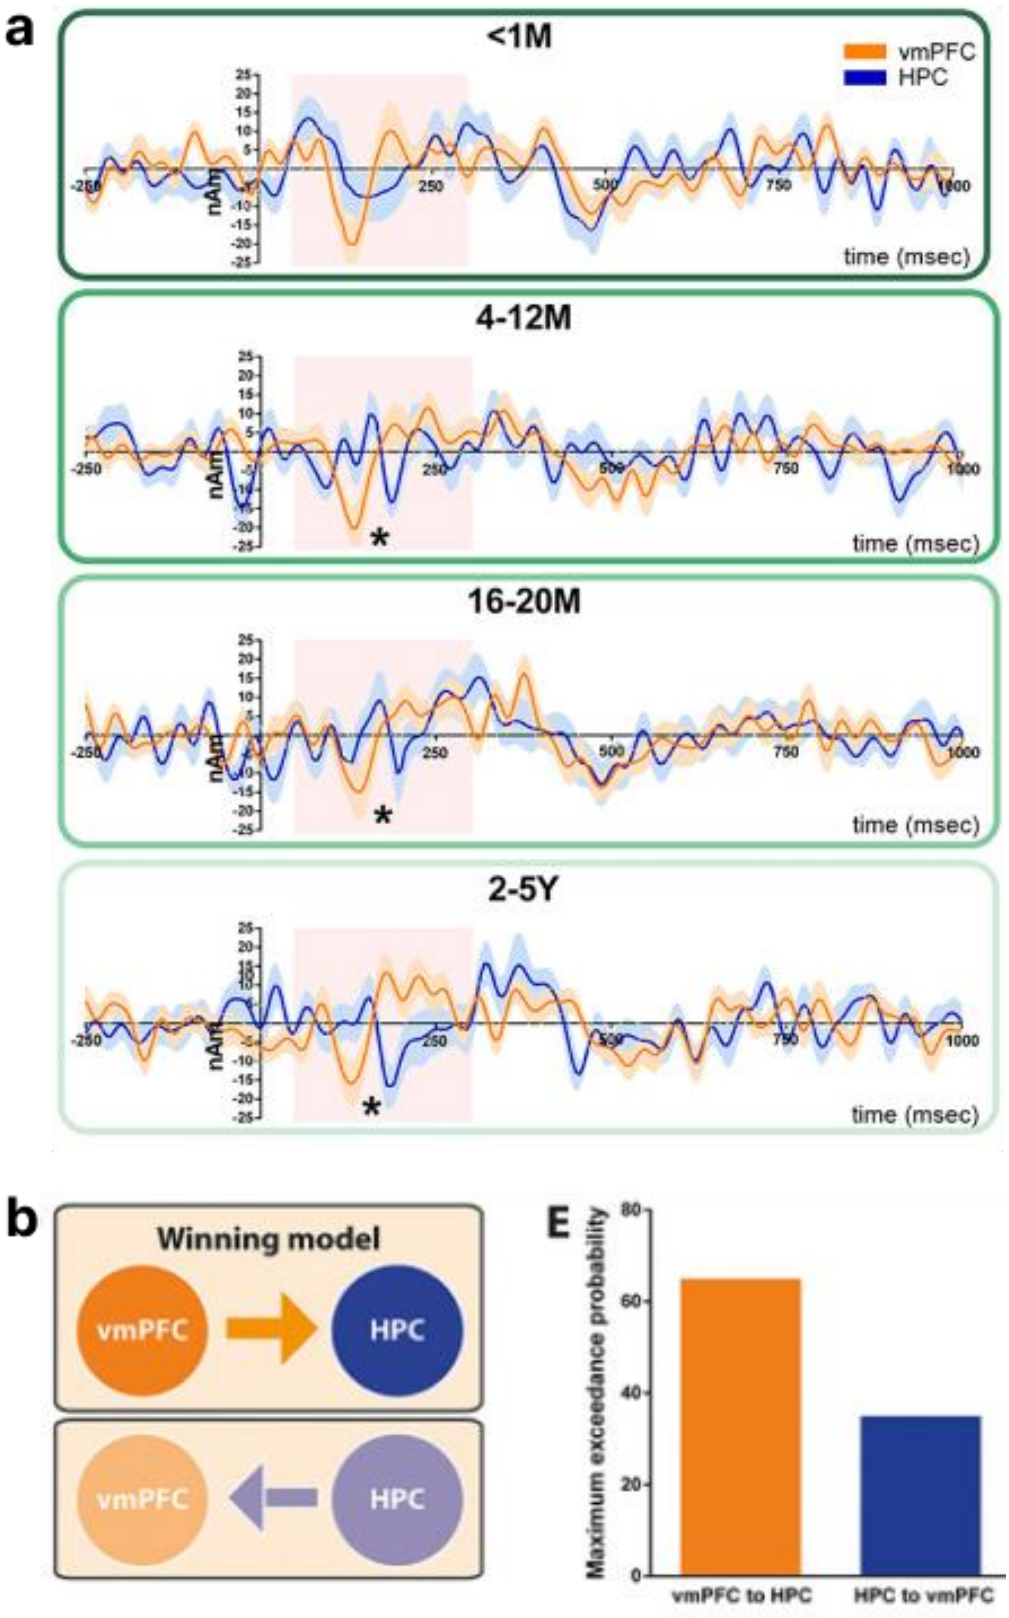

Figure 6. The mPFC initiates and drives activity in the hippocampus during autobiographical memory retrieval. (a) During autobiographical memory there was a significant difference (indicated by the *) in the timing of responses in the vmPFC (orange line) and the hippocampus (blue line) during retrieval of 4-12 month, 16-20 month, and 2-5 year old memories (lower 3 panels), with the mPFC responding prior to the hippocampus. No difference in the timing of responses was observed during the retrieval

of <1 month old memories (upper panel). Shaded areas around the solid line represent the SEM. Pink shaded boxes represent the period in which the maximum response was determined. (b) Two models of effective connectivity between the vmPFC and hippocampus were tested using DCM. The most likely model, determined by Bayesian model comparison, was the one in which the mPFC drove activity in the hippocampus. vmPFC = ventromedial prefrontal cortex, HPC = hippocampus. Figure from McCormick et al. (2020).

In summary, while autobiographical memories are stored more stably in the mPFC over time, the mPFC appears to be actively engaged in the retrieval of autobiographical memories of all ages. Connectivity between the mPFC and hippocampus appears to be important, with the mPFC playing a supervisory role in initiating and driving the retrieval process, particularly of remote autobiographical memories. In the next section, I will explore the functions of the mPFC in other areas of cognition that involve the visualisation of events.

The mPFC and event construction

As well as being active during autobiographical memory retrieval, the mPFC is consistently engaged during processing of future and imaginary events (Benoit et al., 2014; Roy et al., 2012; Schacter, 2012; Spreng et al., 2009). Furthermore, patients with mPFC damage experience difficulty spontaneously imagining future events (Bertossi et al., 2017; Bertossi et al., 2016) and report fewer day-dreams (Bertossi & Ciaramelli, 2016), indicating a more general issue with event processing, beyond memory-related issues.

Unlike damage to the hippocampus, bilateral damage to the mPFC does not prevent patients from vividly imagining in detail single scenes (Kurczek et al., 2015). It does, however, cause deficits in the temporal representation of events. For example, patients are limited in their ability to place memories accurately on a timeline (Tranel & Jones, 2006). Therefore, rather than processing single scenes, the mPFC may be required for internally representing and processing temporally extended events (McCormick, Ciaramelli, et al., 2018).

A second observation from mPFC-damaged patients is that they have a general challenge with producing scene and event content that is relevant and appropriate to reality. Patients have difficulty predicting typical objects that could exist within an imagined scene (De Luca et al., 2018). Moreover, people with mPFC damage often exhibit confabulation, a condition where they describe false, often bizarre and implausible memories, which they believe to be true (Moscovitch, 1989; Moscovitch & Melo, 1997). Therefore, the mPFC might play a role in

selecting appropriate elements for integration within a scene or event (Gilboa et al., 2006; Gilboa & Marlatte, 2017; Moscovitch & Melo, 1997; Preston & Eichenbaum, 2013). This may be due to an impairment in the ability to determine whether information is consistent with prior knowledge. Indeed, the mPFC has been found to represent schemas, abstract prior knowledge structures that have been extracted over multiple repeated events (e.g., the typical sequence of subevents when dining out at a restaurant; Ghosh & Gilboa, 2014). These kinds of event schemas may provide a template for supporting the temporal unfolding of specific events, thereby explaining the involvement of the mPFC in processing temporally extended events and timelines of multiple events.

Finally, there is evidence for a role of the mPFC in initiating scene and event processing beyond autobiographical memory retrieval. Despite being unable to *spontaneously* imagine or recall events, patients are able to do so when highly detailed external cues are provided (Bertossi et al., 2017; Bertossi et al., 2016; Della Sala et al., 1993; Kopelman et al., 1999). Furthermore, although patients report fewer instances of day-dreaming, an example of internally-driven cognition, the few day-dreams that they have are focused on the external environment (Bechara, 2004; Bertossi & Ciaramelli, 2016; Leopold et al., 2012), indicating a potential impairment in the internal generation of event processing.

I have focused here on the role of the mPFC in event processing. However, it has other functions outside of event processing. For example, damage to the mPFC can lead to deficits in moral reasoning, emotion regulation, social behaviour, and understanding others' perspectives (reviewed in McCormick, Ciaramelli, et al., 2018). These deficits may also be related to problems with internal initiation of cognition and selecting appropriate beliefs. Therefore, the mPFC may be required to initiate internally-driven cognition, including the internal visualisation of scenes and events.

1.1.5. Functional hierarchies of event processing

Based on the findings from neuroimaging studies and observations from hippocampal- and mPFC-damaged patients, several hierarchical relationships between the mPFC and hippocampus have been recognised. First, a temporal hierarchy, where the lower-order hippocampus represents single instances in time (i.e., single scenes), while the higher-order mPFC represents temporally extended events under which single scenes unfold (McCormick,

Ciaramelli, et al., 2018). Second, an abstraction hierarchy has been identified, where the higher-order mPFC represents schematic information that is generalised across multiple events, while the lower-order hippocampus processes specific events (Gilboa & Marlatte, 2017; McClelland et al., 1995; McCormick, Ciaramelli, et al., 2018; Preston & Eichenbaum, 2013; Robin & Moscovitch, 2017). Finally, there is a hierarchy of executive control, where the higher-order mPFC has been proposed to initiate and monitor the processing of scenes and events in the lower-order hippocampus (McCormick, Ciaramelli, et al., 2018). Note that under each view, the mPFC is positioned above the hippocampus in a hierarchy of information processing.

These hierarchical relationships have been formulated into a single model of scene and event construction (Figure 7; McCormick, Ciaramelli, et al., 2018). According to this model, the mPFC internally initiates event construction. Schematic representations in the mPFC act to select the appropriate scene elements (e.g., people, objects and landmarks) and inhibit irrelevant elements by activating the representations of these elements in lower-order neocortical areas (Ghosh et al., 2014; Gilboa & Marlatte, 2017; Preston & Eichenbaum, 2013; van Kesteren et al., 2010; van Kesteren et al., 2012). The scene elements are then passed on to the hippocampus where a spatially-coherent scene is constructed (Hassabis & Maguire, 2007; Maguire & Mullally, 2013). To envisage an extended event, the constructed scene is then passed on to the mPFC and the process repeats for each successive scene of the event. As such, the mPFC exerts top-down control of the hippocampus and neocortex to initiate and coordinate unfolding of temporally extended events.

This model views events somewhat like movies such that they are made up of individual image frames, each containing a slight change relative to the previous frame (Monk, Barry, et al., 2021). Because we are continuously blinking and shifting our gaze, much like a movie reel, our visual input is a discreet sequence of image snapshots. However, we are able to perceive the world as a continuously moving image, indicating that the brain is able to link image snapshots across time (Monk, Barry, et al., 2021).

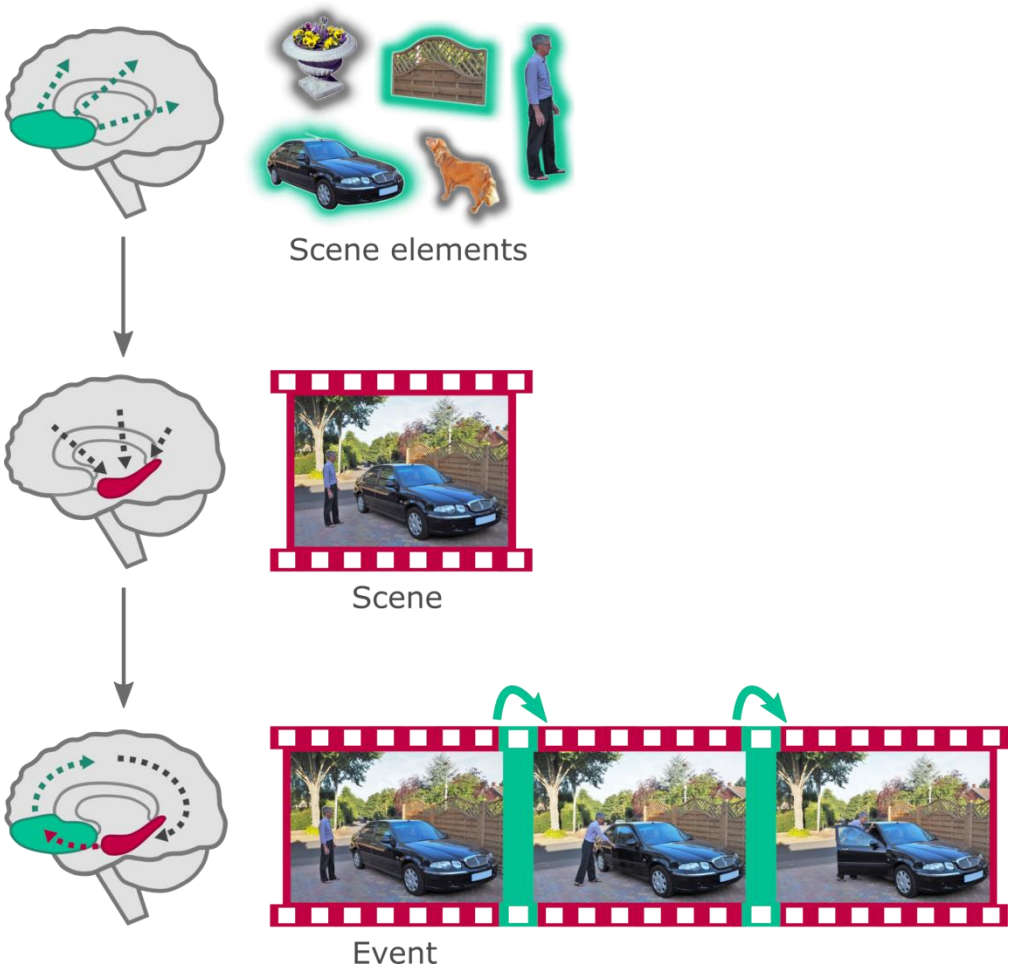

Figure 7. A proposed hierarchical model of event construction. The mPFC (in green) selects relevant scene elements (with green outlines) encoded by neocortical areas. These are then directed to the hippocampus (in red) where they are arranged to form a scene (with red outline). The scene is then passed onwards to the mPFC and the process is repeated (green bands) to allow the unfolding of an extended event. Figure from Ciaramelli et al., (2019).

Retrieval of remote, consolidated autobiographical memories have been hypothesised to require more initialising and coordinating by the mPFC to successfully facilitate the reconstruction of the scenes and events in the hippocampus (Barry & Maguire, 2019; McCormick, Ciaramelli, et al., 2018). This is because the hippocampal memory representation may be become weaker or absent over time (Barry & Maguire, 2019).

It is worth noting that the hippocampus and mPFC have been implicated in a variety of functions beyond scene and event visualisation. For example, extensive research has shown their roles in various executive and decision making functions, particularly in planning and navigating towards behavioural goals (e.g., Crivelli-Decker et al., 2023, Gauthier & Tank, 2018). Many of these functions likely involve the visualisation of events, which could explain the

hippocampus and mPFC involvement. For example, deciding on the optimal course of action necessitates the ability to envisage potential action sequences and their expected outcomes.

1.2. Anatomical hierarchies and microcircuits

So far, I have explained the role of the mPFC and hippocampus in scene and event processing that justifies their proposed hierarchical organisation. In the next section, I will explain an anatomical hierarchy that has been identified in the brain. This anatomical hierarchy has been mapped in the occipital and medial temporal cortical regions, positioning the primary visual cortex at the base of the hierarchy and the hippocampus at the top. It is an anatomical hierarchy that is based on the laminar origins and targets of neuronal connections between cortical areas. Two pathways of information flow between hierarchical levels exist: a feedforward pathway that ascends from lower levels to higher levels and a feedback pathway that descends from higher levels to lower levels. I will explain the well-known pathways of information flow through the hippocampus and how they might correspond to feedforward and feedback pathways. I will also discuss evidence to suggest that the anatomical hierarchy extends to frontal and parietal regions, including the mPFC. Finally, I will describe work that has been performed to map the laminar-based anatomical hierarchy onto a functional hierarchy of visual information processing and to understand the relative involvement of feedforward and feedback signalling on cognitive processes.

1.2.1. A laminar-based anatomical hierarchy

The neocortical grey matter consists of 6 distinct layers, running parallel to the cortical surface, each with specific cell types (Figure 8). Brain areas have been defined based on differences in the density of the various layers between regions of the cortex (Brodmann, 1909). Particularly, the density of layer 4 differs considerably between areas, with some areas, such as the primary visual cortex, having a very high density of granule cells in layer 4 and others, like the premotor cortex, having a very sparse or non-existent layer 4.

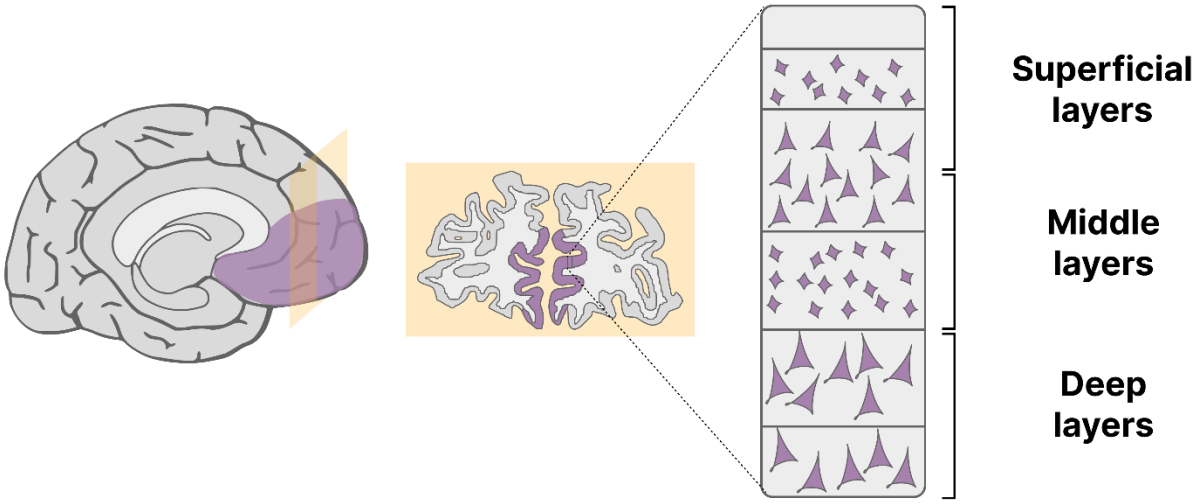

Figure 8. The neocortical layers. A schematic of the 6 distinct neocortical layers in the mPFC (purple shading), which can broadly be separated into superficial (above layer 4), middle (layer 4) and deep (below layer 4) layer depths.

Tracing studies in non-human primates have revealed distinct laminar sites of two types of projections between neocortical areas (reviewed in Markov et al., 2014). Projections from posterior to anterior areas typically originate in the layers above layer 4 (the superficial layers) of more posterior areas and terminate in layer 4 of more anterior areas (Cragg, 1969; Lund et al., 1975; Martinez-Millan & Hollander, 1975; Rockland & Pandya, 1979; Spatz et al., 1970; Wong-Riley, 1978). Conversely, projections from anterior to posterior areas originate in the layers below layer 4 (the deep layers) of more anterior areas and target the superficial and deep layers of more posterior areas (Kaas & Lin, 1977; Kennedy & Bullier, 1985; Kuypers et al., 1965; Tigges et al., 1973; Wong-Riley, 1978). These findings led to the development of a network wiring diagram that organises cortical areas into an anatomical hierarchy, with more posterior cortical areas sitting at the base of the hierarchy and more anterior brain areas sitting at the higher levels (Figure 9; Felleman & Van Essen, 1991). As such, ascending feedforward projections originate in the superficial layers of lower-level areas and terminate in layer 4 of higher-level areas (Figure 10). Whereas, descending feedback projections originate in the deep layers of higher-level areas and terminate in both the deep and superficial layers of lower-level areas.

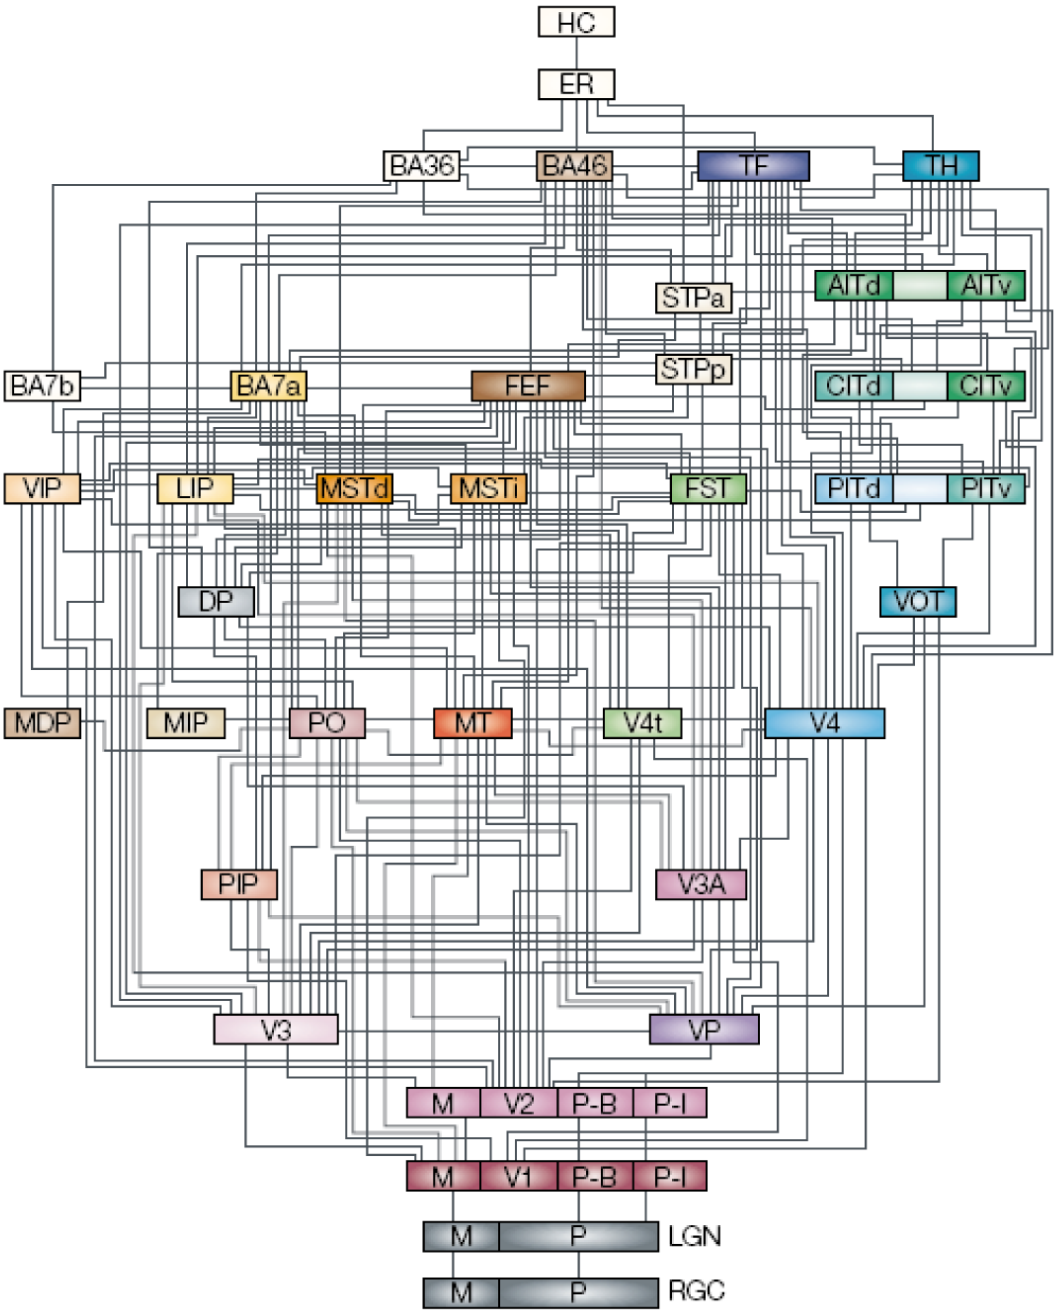

Figure 9. Network diagram of the laminar-based anatomical hierarchy. This network diagram shows the hierarchical organisation of occipital and medial temporal brain regions based on their laminar connectivity. A given area sends feedforward projections, travelling in the middle and superficial layers, to higher-level areas and receives feedback projections, travelling in the deep layers, from higher-level areas. Figure from Felleman & Van Essen (1991).

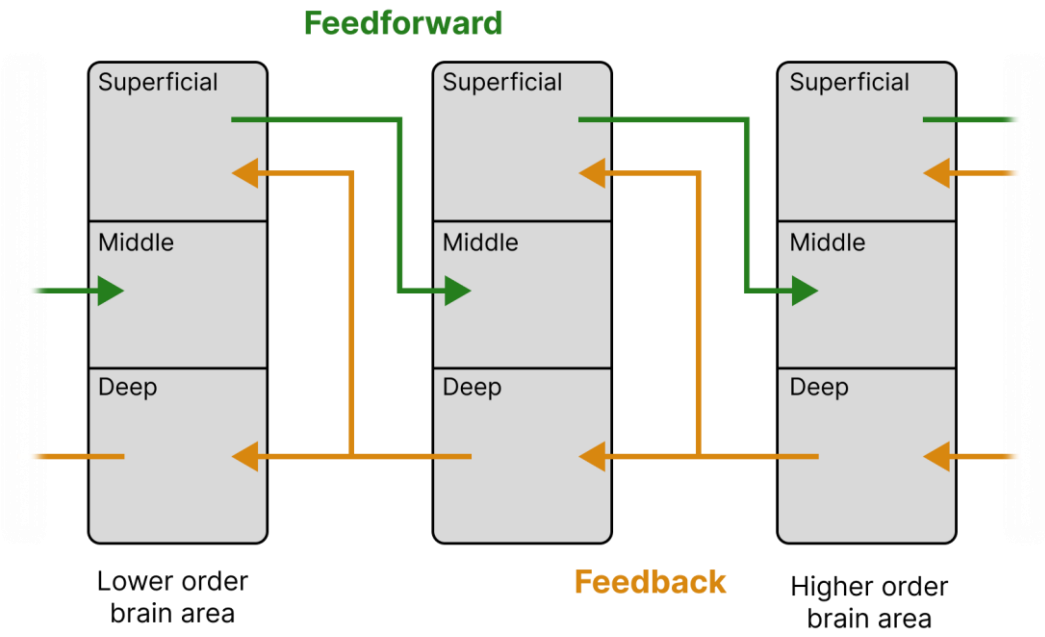

Figure 10. Feedforward and feedback signalling in the cortical layers. Feedforward signalling pathways (in green) originate in the superficial layers and target the middle layers of higher-order cortical areas. Feedback signalling pathways (in orange) originate in the deep layers and target the superficial and deep layers of lower order cortical areas. Figure based on the wiring diagram formulated by Felleman & Van Essen (1991).

Since Felleman and Van Essen's iconic wiring diagram, further anatomical tracing studies have been performed which have led to refinements and expansions to the diagram (Markov et al., 2014). Markov et al. (2014) found that feedforward projections specifically originate in layer 3B and target both layer 4 and layer 3B of higher-level cortical areas. Similarly, the feedback pathway was found to originate specifically in layer 6 and to target layer 6 and layer 1 of downstream areas. Additionally, these well characterised feedforward and feedback pathways were observed to function exclusively over longer distances. Two additional short-distance pathways were observed between adjacent cortical areas: a feedforward pathway in layers 5 and 6 and a feedback pathway in layer 2/3A.

1.2.2. Mapping the anatomical hierarchy onto a feature representation hierarchy

The laminar-based anatomical hierarchy has been successfully mapped onto a functional hierarchy of visual information processing (reviewed in Felleman & Van Essen, 1991; Ungerleider & Haxby, 1994; Van Essen & Gallant, 1994). Information pertaining to a visual stimulus arrives into the lowest level of the functional hierarchy, where neurons respond to local features of the stimulus, like changes in contrast. These representations are then passed

on to higher levels, where higher-level image features are represented, such as shapes. As the information progresses further up the hierarchy, it becomes more integrated until, eventually, the entire visual stimulus is represented.

Researchers have found that this visual processing hierarchy maps onto the laminar-based anatomical hierarchy, such that at increasingly higher levels of the anatomical hierarchy, neurons have larger receptive fields and to respond to more global, complex image features (e.g., D'Souza et al., 2022; Harris et al., 2019). For example, the primary visual cortex, located at the base of the anatomical hierarchy, responds to basic visual features like changes in image contrast, while the hippocampus, positioned at the top, responds to more complex and global features, such as entire scenes. A recent study further explored this by using a deep learning neural network to generate, from electrophysiological recordings, scenes that mice were presented with (Chen et al., 2024). They found that, when using recordings from progressively higher levels of the anatomical hierarchy, the quality of the image reconstructions decreased. This may be because the more global and meaningful features represented by higher-level areas (e.g., the hippocampus) do not require the representation of pixel-level image details. By contrast, the basic features, such as local contrast variations, represented by lower-level areas (e.g., the primary visual cortex) would contain this level of image precision.

1.2.3. Microcircuitry of the hippocampal subfields

The hippocampus has been positioned at the apex of the laminar-based anatomical hierarchy (Figure 9; Felleman & Van Essen, 1991). While the laminar-based anatomical hierarchy does not split the hippocampal subfields into hierarchical levels, we know that information flows through the hippocampal subfields in very specific ways. In this section, I will describe the well-known microcircuitry of the hippocampal subfields and how the anatomical connections may correspond to feedforward and feedback signalling pathways (Figure 11).

Most connections between the hippocampus and the neocortex pass through the entorhinal cortex, a cortical region in the medial temporal lobe, adjacent to the pre/parasubiculum. The primary feedforward inputs to the hippocampus project from the superficial layers (layers 2/3) of the entorhinal cortex, reminiscent of feedforward projections. Whereas outputs project back to the deep layers of the entorhinal cortex, reminiscent of feedback projections.

There are two main pathways of information flow from the entorhinal cortex into the hippocampus: the trisynaptic pathway and the monosynaptic pathway (Amaral & Witter, 1989; Lorente de Nò, 1934; Witter & Amaral, 1991, 2020). The trisynaptic pathway is a series of 3 projections. The first projection, the perforant path, projects from layer 2 of the entorhinal cortex to the DG and CA3. The next, mossy fiber, path projects from the DG to CA3. Finally, the Schaffer collateral path projects from CA3 to CA1 and then onwards to the subiculum. CA3 also contains dense self-projections, that is projections to itself. The second pathway of information flow into the hippocampus, the monosynaptic pathway, sometimes referred to as the temporo-ammonic pathway, comprises direct connections from layer 3 of the entorhinal cortex to CA1 and the subiculum.

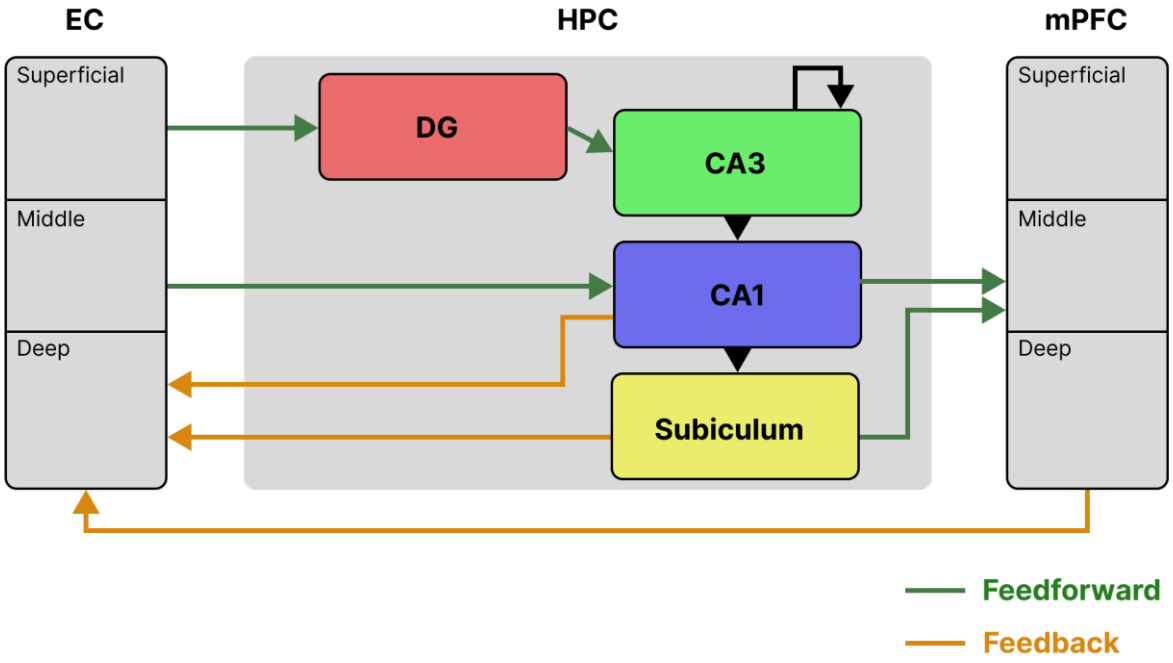

Figure 11. Hippocampal microcircuits. Information from the neocortex enters the hippocampus via the entorhinal cortex. The trisynaptic loop pathway projects from the entorhinal cortex to the DG, on to CA3, and finally on to CA1. Recurrent connections exist within CA3. The monosynaptic pathway projects directly from the entorhinal cortex to CA1. Outputs from the hippocampus project from CA1 and the subiculum back to the entorhinal cortex as well as onwards to the mPFC. The mPFC projects back to the entorhinal cortex, from where it can re-enter the hippocampus. Based on the laminar sources and/or targets of projections, some projections are coded as either feedforward (green) or feedback (orange). DG = dentate gyrus, CA = Cornu Ammonis, EC = entorhinal cortex, mPFC = medial prefrontal cortex.

Based on their circuitry, and supported by findings from rodent and human studies, unique computations are thought to occur in the different subfields of the hippocampus (Baker et al., 2016; Berron et al., 2016; Guzman et al., 2016; Leutgeb et al., 2007; Lisman, 1999; McNaughton & Morris, 1987; Rebola et al., 2017; Rolls & Treves, 1994; Treves & Rolls, 1994). Due to its large cell population, the DG is thought to create sparse neural representations. This process, known as pattern separation, allows for the existence of distinct representations of overlapping events (i.e., two events that occurred at the same time in the same location). Due to its dense recurrent collateral connections, CA3, is thought to form complex associative representations. This unique recurrent architecture allows for the activation of a full associative representation based on information that constitutes just a part of the full representation. This process, known as pattern completion, is a mechanism by which complete memories can be retrieved based on single features of a memory such as a place or person.

Outputs from the hippocampus project back to the entorhinal cortex deep layers from CA1 and the subiculum (Witter & Amaral, 2020). From here, these outputs are sent back to the entorhinal cortex superficial layers, where they are sent back into the hippocampus, creating "big loop" recurrence (Koster et al., 2018). As well as sending feedback projections to the entorhinal cortex, CA1 and the subiculum also send outputs to layer 3 of the mPFC, reminiscent of feedforward projections (Aggleton et al., 2015; Carmichael & Price, 1995; Koster et al., 2018). While there are no feedback-type pathways projecting directly from the mPFC to the hippocampus, there are dense projections from the mPFC deep layers to the entorhinal cortex deep layers, from where information may pass into the hippocampus. There are also less direct routes of feedback signalling from the mPFC to the hippocampus, synapsing in the parahippocampal cortex, retrosplenial cortex and lateral temporal cortices prior to the entorhinal cortex. Therefore, clearly information processing does not stop at the hippocampus. In the next section, I will examine evidence for the laminar-based anatomical hierarchy extending beyond the hippocampus to cortical areas such as the mPFC.

1.2.4. Extending the anatomical hierarchy to the whole brain

Much of the existing work on the laminar-based anatomical hierarchy has used relatively simple visual stimuli to map a functional hierarchy of brain areas, including the occipital and medial temporal cortices. However, as explored in Section 1.1, the complex processing of

scenes and events involves many areas outside of the occipital and medial temporal lobes, including the mPFC.

Whole-brain anatomical networks have been mapped, based on tracing studies in non-human primates (Figure 12; Averbeck & Seo, 2008). Projections exist from occipital and medial temporal cortical areas, such as the visual cortex and hippocampus, to frontal areas. Therefore, the laminar-based anatomical hierarchy is clearly a single section of a wider whole-brain anatomical network.

Unlike the visual system, the laminar sources and targets of neurons in these whole-brain networks have not been systematically mapped. However, there is evidence to suggest that the laminar connectivity patterns between cortical areas is consistent throughout the brain. Like the occipital and medial temporal lobe areas that make up the laminar-based anatomical hierarchy, areas in the frontal and parietal lobes also contain 6 distinct cortical layers. As with areas in the occipital and medial temporal lobes, some fronto-parietal areas contain a denser layer 4 while other contain a layer 4 that is very sparse. In a recent primate study using electrolytic markers and histology, a specific laminar pattern of neural oscillations was observed: gamma oscillations peaked in layers 2/3, alpha-beta oscillations in layers 5/6, and the crossover point between them was in layer 4 (Mendoza-Halliday et al., 2024). This electrophysiological laminar profile was present in all cortical areas, whether they contained a highly dense layer 4 (e.g., the primary visual cortex) or a very sparse layer 4 (e.g., the premotor cortex). As neural oscillations result from synchronised firing in ensembles of connected neurons, this is strong evidence to suggest that the same laminar connectivity patterns exist across the whole cerebral cortex, irrespective of the density of layer 4.

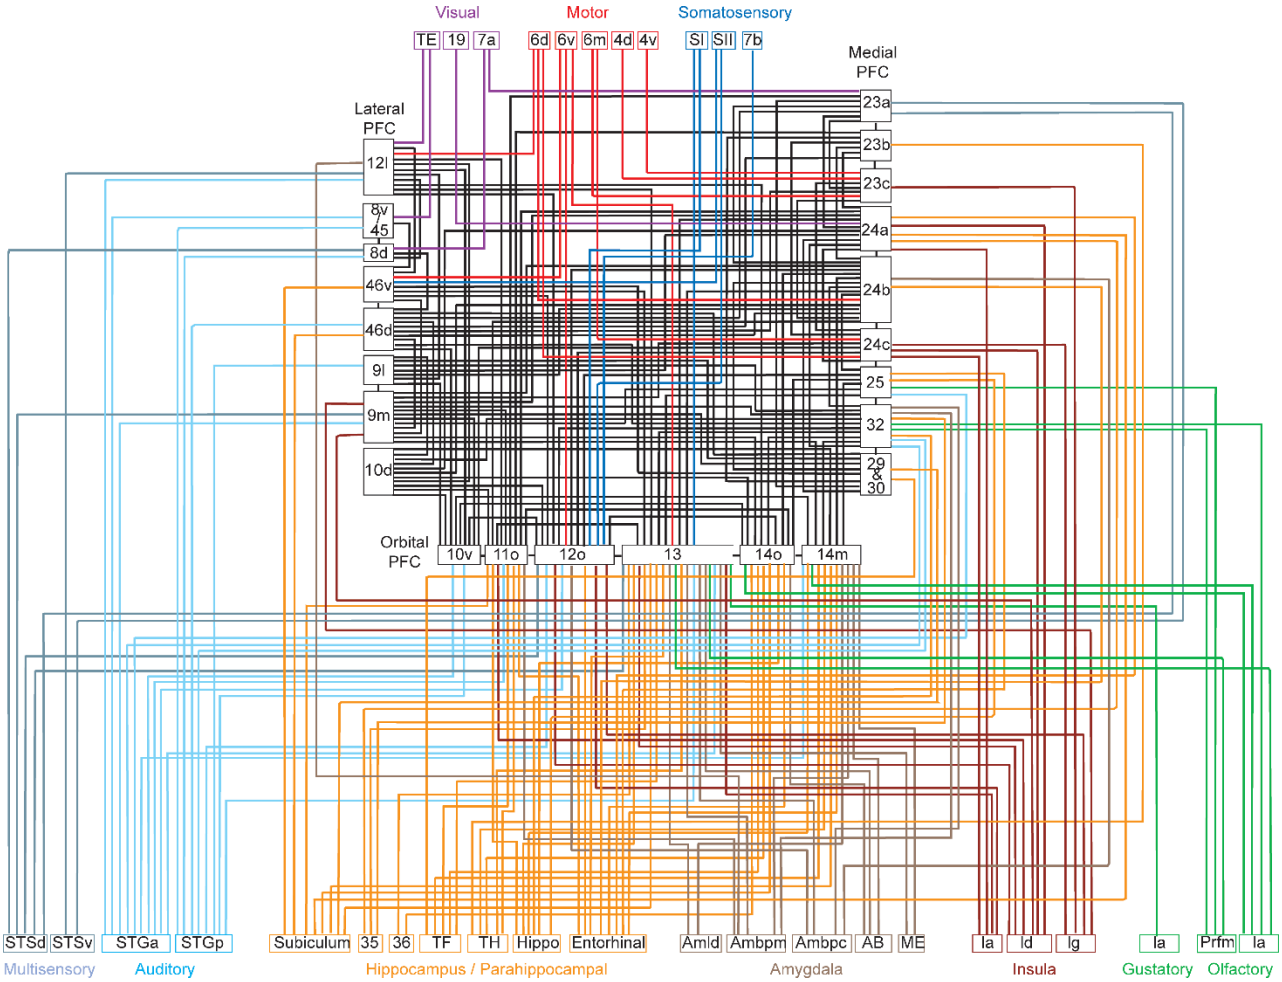

Figure 12. Whole brain network diagram. Network diagram showing the dense anatomical connections of the frontal cortex with other brain regions based on anatomical tracing studies in non-human primates. Note: this diagram only includes the intermediate and strong projections to the frontal cortex and, therefore, not the full extent of connectivity. Figure from Averbeck & Seo (2008).

1.2.5. Mapping feedforward and feedback signalling pathways onto cognitive functions

Considerable work has been carried out to understand the function of the two major information processing pathways through the laminar-based anatomical hierarchy: the feedforward and feedback signalling pathways. Much of this work has used a method called laminar fMRI, in which neuronal activity is imaged across 3 cortical depths (superficial, middle and deep) using high resolution 7T MRI scanning. By observing which depth-specific activations or representations during a given task it is possible to infer the contribution of feedforward and feedback signalling pathways.

Laminar fMRI studies have revealed that feedforward pathways transmit externally-generated sensory information coming from the sensory organs (e.g., Bergmann et al., 2024; Carricarte et al., 2024). Feedback signalling pathways, on the other hand, have been associated with modulatory cognitive processes such as attention (Liu et al., 2021) and working memory (Lawrence et al., 2018). As well as modulating external signals, feedback signalling pathways have also been found to carry informational content, sometimes even in the absence of an external stimulus. For example, activations caused by specific expected but not presented stimuli (Aitken et al., 2020; Bastos et al., 2012; Kok et al., 2016; Warrington et al., 2024; Yu et al., 2019) and imagined visual content (Bergmann et al., 2024; Carricarte et al., 2024) have been localised in the deep cortical layers. Here, information is thought to flow through the hiearchy in a top-down fashion, where higher-level visual information such as whole scenes are activated and passed downwards to lower-levels where lower-level features such as objects, shapes or contrasts are activated.

In summary, further work is required to map an anatomical hierarchy that extends to include higher-order brain areas. We know that information is sent onwards from the visual cortical hierarchy to frontal areas. We know that, during scene and event construction, frontal areas, specifically the mPFC, process highly abstracted event information. Therefore, it seems likely that they would be positioned high up in a representational hierarchy of event processing. We also know that the layers of frontal cortical areas have the same electrophysiological properties as the layers of lower order cortical areas. This indicates that they may exhibit the same inter-area laminar connectivity patterns as identified in the occipital and medial temporal lobes.

1.3. Aims and outline of the thesis

The broad aim of this thesis was to advance our understanding of how the mPFC, hippocampus and a wider system of brain regions are functionally organised to support memory and perception of events. Specifically, I was interested in their organisation at the level of their neural microcircuitry, that is, the circuits formed from neuronal projections between the cortical layers and hippocampal subfields. I asked two main questions: (1) what are the functions of the mPFC and hippocampus in event processing? (2) how does their laminar and subfield microcircuitry support these functions? An examination of mPFC and

hippocampal functions with laminar resolution provided insights into involvement of feedforward and feedback signalling pathways. Furthermore, the hippocampal subfield resolution provided insights into the involvement of pathways of information flow through the hippocampus and the computational functions associated with these pathways. Ultimately, this work aimed towards mapping a cognitive hierarchy of event processing onto a laminar and subfield-based anatomical hierarchy.

To address these aims, I used 7T MRI scanning. This ultra-high 7T magnetic field strength provides significant improvements to the spatial resolution compared to the more typical 3T MRI scanning, allowing neuronal activity to be measured within 0.8 x 0.8 x 0.8 mm^3 voxels (3D pixels) of brain tissue. Given that the neocortical grey matter is, on average, 3 mm thick, this submillimeter voxel size enables the dissociation of neuronal activity coming from three cortical depths: deep, middle and superficial (Fischl & Dale, 2000). This approach, known as laminar fMRI or depth-dependent fMRI, has been applied to determine the involvement of the different cortical layers in various cognitive processes (e.g., Aitken et al., 2020; Bastos et al., 2012; Kok et al., 2016; Warrington et al., 2024; Yu et al., 2019). Depending on cytoarchitecture of the cortical area of interest, the deep, middle and superficial cortical depths will map onto slightly different of the 6 cytoarchitectonic layers. For example, within the mPFC, the relative thickness of the cortical layers differs between Brodmann areas. Therefore, currently, laminar fMRI is reasonably coarse and the interpretation of results should be carefully considered for each area based on the specific laminar anatomy of the cortical area.

The submillimeter voxel resolution achieved by 7T fMRI also allows for the distinction of activity in the hippocampal subfields. In fact, lower resolution 3T fMRI scanning with a millimeter voxel size can successfully distinguish activity in the hippocampal subfields (Bonnici, Chadwick, Kumaran, et al., 2012; Bonnici et al., 2013; Chadwick et al., 2010). However, the smaller voxel sizes afforded by 7T fMRI help to reduce partial volume effects, which are present when a voxel crosses subfield boundaries and lead to the blurring of signal between neighbouring subfields.

There are relatively few 7T fMRI studies of event processing. Two recent 7T fMRI studies have investigated the retrieval of autobiographical memories in the hippocampal subfields (Leelaarporn et al., 2024; Pfaffenrot et al., 2024). Leelaarporn et al. (2024) found peak

activation in the pre/parasubiculum when comparing retrieval of (>1 year old) autobiographical memories to a counting baseline condition, supporting its role in scene-related processing. Pfaffenrot et al. (2024) compared activation in the inner and outer layers of the hippocampal subfields during retrieval of (<2 year old) autobiographical memories. They found higher activation in the inner compared to the outer layers, indicating the involvement of the trisynaptic loop pathway (and perhaps its pattern separation and completion functions), which synapses in the inner layers of CA1. One 3T fMRI study, has also investigated autobiographical memory retrieval in the hippocampal subfields, finding that remote (10 years old) memories are represented more strongly in CA3 and the posterior DG than recent (2 weeks old) memories, perhaps indicating higher pattern completion demands as memories age (Bonnici et al., 2013). These fMRI studies have advanced our understanding of the hippocampal microcircuitry supporting autobiographical memory retrieval. However, retrieval of recent and remote autobiographical memories has not been compared in the hippocampal subfields using 7T fMRI.

One 7T fMRI study, investigated episodic encoding and retrieval in the hippocampus and mPFC layers, along with other cortical areas (Chang et al., 2022). During the encoding task, participants were presented with the name of a famous person and an object and were tasked with imagining the person interacting with the object. During retrieval, which was immediately following encoding, they were presented with a name-object pair and responded as to whether they had previously imagined the corresponding event during the encoding task. Higher functional connectivity between the mPFC deep/middle layer and the posterior hippocampus was observed during retrieval compared to encoding, supporting previous work showing that interactions between the mPFC and hippocampus are involved in the retrieval of past episodes (Fuentemilla et al., 2014; McCormick et al., 2020; Nawa & Ando, 2019, 2020). The combination of deep layer and middle layer involvement implies the increased involvement of both feedforward and feedback signalling pathways during retrieval compared to encoding. However, unlike naturalistic event encoding, the encoding in this task relied purely on imagination, were participants were internally-generating the event from basic cues rather than externally perceiving an event. This distinction is important when investigating cognitive processes in terms of laminar-specific activations as externally-driven and internally-driven cognitive processes may rely on feedforward and feedback processing to different

degrees organs (Bergmann et al., 2024; Carricarte et al., 2024). Naturalistic autobiographical memory retrieval has not yet been investigated in the mPFC layers nor in the layers of other cortical areas in the core network and no studies have, to my knowledge, investigated the connectivity between the hippocampal subfields and the cortical layers during autobiographical memory retrieval. Furthermore, to my knowledge, no 7T fMRI studies have investigated the hippocampal and cortical microcircuitry underpinning event processing beyond episodic memory, for example during the online perception of events (although some 7T fMRI studies have examined the perception of individual scenes: Berron et al., 2016; Carricarte et al., 2024; Grande et al., 2019; Maass et al., 2014).

Measuring neural activations in the mPFC layers comes with significant challenges. While the high field strength used in 7T fMRI allows for exceptional spatial resolution, it concurrently increases the severity of certain image artefacts, such as geometric distortion, which typically manifests as stretching and compressing of the brain tissue. Due to its proximity to the sinuses, the mPFC experiences particularly high levels of geometric distortion. This can cause a blurring of the measured signal between neighbouring voxels, thereby reducing our ability to distinguish neural activity in the cortical layers.

Only three laminar fMRI study have attempted to image the frontal cortices (Chang et al., 2022; Degutis et al., 2024; Finn et al., 2019) and only one of these has imaged the highly challenging mPFC (Chang et al., 2022). At the time of conception of my PhD, there were no published works that used 7T fMRI to image the mPFC. Therefore, I spent a significant portion of my PhD working with the Physics Team to test various 7T fMRI acquisition methods, with the aim of minimising geometric distortion, while still achieving a reasonable temporal resolution of less than 4 seconds and coverage of the hippocampus and mPFC. By combining several advanced imaging methods, we were able to significantly improve the fMRI image quality. However, it was not possible to entirely remove geometric distortion in the mPFC during acquisition. Therefore, I also spent a period of my PhD testing image preprocessing and analysis methods with the aim of further correcting any remaining distortion to enable the examination of layer-specific neural activity.

A second image artefact that interferes with the detection of laminar-specific activity is the so-called superficial bias effect or venous draining effect (Chang et al., 2022; Degutis et al., 2024; Finn et al., 2019; Kay et al., 2019; Norris & Polimeni, 2019; Ugurbil, 2016). This effect

results from the fact that, in gradient echo (GE) blood oxygen level dependent (BOLD) fMRI, most of the measured signal arises from large veins (Boxerman et al., 1995). Because there are more large veins in the superficial layers, and blood drains from the deep to superficial layers (Duvernoy et al., 1981), the measured signal is highest in the superficial layers. To address this bias, I tested various preprocessing and analysis methods that generally aimed to remove signal coming from the large veins from the analysis. This work, combined with the work on correcting geometric distortion, resulted in a laminar fMRI preprocessing and analysis pipeline that was optimised towards measuring layer-specific effects in highly distorted brain areas such as the mPFC.

A final methodological challenge that I encountered during my PhD was with respect to the manual definition of hippocampal subfields in MRI scans. This manual process, while the gold standard approach, requires significant training on the anatomy of the hippocampus and, even following this training, still takes approximately 8 hours per scan. To help to remove this potential barrier to investigations into the hippocampal subfields, I developed a protocol to automatically segmenting the hippocampal subfields in both 7T and 3T MRI scans. This tool is now publicly available, with full open access, for future researchers to use (Hickling et al., 2024).

Following a general overview of the methods used in this thesis (Chapter 2), I will present 4 experimental chapters.

Chapter 3: In the first experimental chapter, I report the 7T MRI sequence pilots that I performed, with the Physics Team, with the goal of minimising geometric distortion in the mPFC. Subsequently, I report the preprocessing and analysis tests that were conducted to yield a complete pipeline for laminar analyses in the mPFC.

Chapter 4: In the second experimental chapter, I report some work that I did to develop a tool for automatically segmenting the hippocampal subfields using a detailed segmentation protocol.

Chapter 5: In the third experimental chapter, the first 7T laminar fMRI investigation of event processing, I sought to identify the laminar and subfield microcircuitry supporting the retrieval of recent and remote autobiographical memories. By investigating recent and remote autobiographical memory retrieval at the level of the cortical layers, I was able to characterise

these processes in terms of feedforward and feedback signalling. Furthermore, a subfield-level analysis of the hippocampus provided insights into the specific hippocampal microcircuitry underpinning the retrieval of autobiographical memories over time.

Chapter 6: In the second 7T laminar fMRI study on event processing, I examined the microcircuitry supporting event perception. I probed the specific roles of the mPFC and hippocampus in the processing of individual scenes and unfolding events, testing the hypothesis that the mPFC is more involved in the processing of dynamic extended events, while the hippocampus is necessary for any task requiring the visualisation of scene imagery. By investigating event perception at the level of the cortical layers, I was able to characterise it in terms of feedforward and feedback signalling. Furthermore, a subfield-level analysis of the hippocampus provided insights into the specific hippocampal microcircuitry underpinning scene and event perception.

Finally, I draw together the results from all of these experimental chapters for a discussion in Chapter 7, considering in particular what can be learned about the involvement of the hippocampal subfields, mPFC layers and the layers of other cortical areas, as well as their relationships, in the retrieval and perception of events.

Given that the areas in the core network are widespread, and much of the extant studies have investigated the mPFC and other cortical areas in terms of broad regions, I did not have specific questions around the short-range inter-area connectivity (e.g., between local Brodmann areas within the mPFC brain region). Therefore, I chose to investigate reasonably broad cortical areas, and focus on the long-distant laminar connectivity patterns between these areas and the hippocampal subfields.

My work provides new insights into the laminar and subfield microcircuitry supporting event processing, which, until the relatively recent advancements in 7T fMRI, was only possible with invasive studies typically involving animals. This precluded a microcircuit level understanding of cognitive processes that are not possible to research in animals, such as the perception and memory of autobiographical events. Understanding the how the human brain's microcircuitry functions in these aspects of cognition is important as it provides a baseline for investigations into its malfunction in disease states.

Publications

During my PhD, I produced the following publications.

Hickling, A. L., Clark, I. A., Wu, Y. I., & Maguire, E. A. (2024). Automated protocols for delineating human hippocampal subfields from 3 Tesla and 7 Tesla magnetic resonance imaging data. *Hippocampus*, *34*(6), 302–308. https://doi.org/10.1002/hipo.23606

Hickling, A. L., Graedel, N. N., Clark, I. A., Josephs, O., Malekian, V., Kok, P., Callaghan, M. F., Maguire, E. A. Feedforward and feedback signalling during recent and remote autobiographical memory retrieval. *In Preparation.*

Hickling, A. L., Monk, A. M., Clark, I. A., Maguire, E. A. The laminar and subfield microcircuitry supporting scene and event perception. *In Preparation.*

I also contributed towards the following publications. These publications are not covered in detail in this thesis.

Malekian, V., Graedel, N. N., **Hickling, A.**, Aghaeifar, A., Dymerska, B., Corbin, N., Josephs, O., Maguire, E. A., & Callaghan, M. F. (2023). Mitigating susceptibility-induced distortions in high-resolution 3DEPI fMRI at 7T. *NeuroImage*, *279*, 120294. https://doi.org/10.1016/j.neuroimage.2023.120294

Miller, T. D., **Hickling, A. L.**, Wu, Y. I., Zhou, J., Handel, A. E., Pollak, T. A., Coutinho, E., Zandi, M. S., Rosenthal, C. R., Maguire, E. A. Decoding memory dysfunction in hippocampal amnesia. *In preparation*.

Wu, Y. I., **Hickling, A. L.**, Alexander, N., Graedel, N. N., Seymour, R. A., Josephs, O., Malekian, V., Callaghan, M. F., Maguire, E. A. 7T laminar fMRI responses during encoding and retrieval of naturalistic virtual experiences. *In preparation.*

Wu, Y. I., **Hickling, A. L.**, Alexander, N., Seymour, R. A., Maguire, E. A. 7T laminar fMRI responses during retrieval of naturalistic virtual experiences after sleep consolidation. *In preparation.*

Wu, Y. I., **Hickling, A. L.**, Alexander, N., Seymour, R. A., Maguire, E. A. 7T laminar fMRI responses during retrieval of naturalistic virtual experiences after short-term versus long term consolidation. *In Preparation.*

2. General Methods

2.1. Introduction

In this chapter, I will describe the various methods that I used for data collection and analysis. I will first explain how I selected and recruited the participants (Section 2.2), as well as the procedures I followed when a participant visited the Centre for behavioural and neuroimaging testing sessions (Section 2.3). Next, I describe how structural and functional images were acquired during MRI scanning (Section 2.4) and why high field strength (i.e., 7T) MRI scanning provides exciting opportunities for cognitive neuroscience investigations, but also comes with challenges (Section 2.5). Following this, I will outline the specific 7T MRI sequences that I used (Section 2.6). I will then explain the methods I deployed for preprocessing (Section 2.7) and analysing the fMRI data. This includes using both whole-brain fMRI analyses (Section 2.8), which examine activity across the entire brain volume (in my case it was actually a partial-brain volume rather than a whole-brain volume), and region-of-interest (ROI) fMRI analyses, which focus on predefined brain regions (Section 2.9).

Note that all data were acquired under challenging circumstances, with very stringent COVID-19 safety requirements, as described in Section 2.3. Furthermore, Experiment 3, while reported here first (Chapter 5), was actually performed second. This is because it required participants to have 6 distinct memories (involving different people, activities and locations) from the past month, which was simply not possible during government-enforced lockdowns.

2.2. Participants

All 7T MRI participants were healthy young adults (exact demographic details are provided in the specific experimental chapters) with no self-reported history of neurological or psychiatric conditions, and with normal or corrected-to-normal vision. The experiments were approved by the University College London Research Ethics Committee (project ID: 18721/001) before recruitment began. They were recruited through two different sources: (1) the UCL Division of Psychology and Language Sciences SONA participant database; (2) the Call for Participants recruitment service (https://www.callforparticipants.com). From these platforms, potential participants were then directed to the study website that contained key information about

the study, eligibility requirements and reimbursement (https://ahickling1.wixsite.com/brainimagingstudy).

Potential participants were contacted by telephone to answer a number of 7T MRI safety screening and study eligibility questions. Key safety requirements included: no history of surgeries involving the insertion of metal, no tattoos in the upper half of the body, and no claustrophobia. These potential participants also read the participant information sheet and the COVID-19 procedures information sheet and had the opportunity to ask questions about the experiment.

For Experiment 2 (Chapter 4), I also availed of 3T MRI data that was previously acquired between March 2015 and June 2017 (Clark & Maguire, 2023). The experiment was approved by the University College London Research Ethics Committee (project ID: 6743/001) before recruitment began. These participants were also healthy young adults, with no self-reported neurological or psychiatric conditions, who passed the safety screening for 3T MRI scanning.

The final samples across experiments totalled 52 participants scanned using 7T MRI and 140 participants using 3T MRI. Several participants took part in multiple experiments (Table 1).

Table 1. Number of participants that took part in each experiment.

Experiment	Participants (n)
1	2
2 (7T)	15
2 (3T)	140
3	2
4	8
2 (7T) and 3	10
2 (7T) and 4	7
1, 2 (7T) and 3	4
2 (7T), 3 and 4	3
1, 2 (7T), 3 and 4	1

2.3. Experimental set up

The research took place at the Wellcome Centre for Human Neuroimaging, Department of Imaging Neuroscience, UCL Queen Square Institute of Neurology at University College London, UK. In all experiments, participants underwent a combination of neuroimaging and behavioural testing sessions. In Experiment 3 (Chapter 5), the sessions were split across two separate visits. All other experiments involved a single visit to the Centre. The specific details and order of the neuroimaging and behavioural sessions varied by experiment and are described in full in each of the experimental chapters. Examples of behavioural sessions were interviews, pre-scan task training, post-scan testing and post-scan debriefs. The neuroimaging sessions all involved 7T MRI scanning, which is described in more detail in Section 2.5. The 3T data had already been collected and were openly available (Clark & Maguire, 2023; and see Experiment 2, Chapter 4).

Several measures were taken to reduce the risk of COVID-19 transmission during a visit. Within 24-hours of their visit, the participant was telephoned and asked a series of questions to screen them for COVID-19 symptoms. If they were experiencing any symptoms, their appointment was cancelled and rescheduled once they felt well. All equipment was thoroughly cleaned immediately before and after the participant's visit. Prior to their arrival, I changed into medical scrubs and put on a face covering. Upon their arrival, the participant was provided with a face covering and their temperature was recorded. If it was higher than 38 degrees Celsius, the appointment was terminated and rescheduled for a later date. They were asked to maintain a 2-meter distance between themselves and anyone else at all times. They were also encouraged to regularly sanitise their hands with hand sanitiser that was provided throughout the building.

The participant was escorted to a designated quiet room, which was equipped with two workstations, one for the participant and one for myself, separated by a Perspex screen. All behavioural testing sessions were conducted in this room. They reviewed the participant information sheet and the COVID-19 information sheet once more. They had the opportunity to ask any remaining questions before providing written informed consent. Once they had consented to the experiment, they changed into medical scrubs. If the visit included task-

based fMRI as part of the neuroimaging session, the pre-scan task training was conducted next.

In preparation for the neuroimaging session, they removed their shoes, along with any metal items or make-up that could cause signal artefacts in the MRI data. They underwent a thorough screening process, answering the same 7T MRI safety questions they had previously answered during the telephone screening. If they were deemed suitable for MRI scanning, they then signed the MRI Safety Questionnaire.

The participant was allowed to remove the face covering when being MRI scanned. They were positioned in the scanner by a radiographer who wore full PPE (face covering, gloves and apron). They lay on their back on the scanner bed with their head positioned inside the MRI head coil. Cushions and foam pads were placed around their head for comfort and to minimise head movement. Foam earplugs were used to protect their ears from the sound of the scanner. An emergency alarm bell was taped to their abdomen and was easily accessible to their left hand. They were instructed to use this bell at any time to signal that they wanted to leave the scanner. While they were in the scanner, an intercom system was used to communicate with them. If the neuroimaging session involved task-based fMRI, the visual stimuli and task instructions were projected onto a screen via a mirror. Auditory stimuli were presented using Etymotic Ear-Tone stereo sound system (Etymotic Research Inc, Illinois) headphones. These were used in place of the protective foam earplugs as they also provide ear protection from the noise produced by the scanner. The sounds were transmitted via plastic tubes to foam ear buds. The plastic tubes were taped to the collar bone of the participant to avoid them from tugging on the ear buds and coming loose. Task responses were made using an MRI-compatible 4-way button box, which was placed in the right hand of the participant. The task stimuli were delivered and responses stored using scripts written with the Psychtoolbox for MATLAB. This setup was standardised across all participants and experiments. After scanning, the participant was escorted back to the quiet room, where they underwent additional behavioural debriefing session(s).

They were paid £10 per hour at the end of their final visit to cover their travel and meal expenses.

2.4. The biophysics of MRI

I will now provide some general background to MRI, the neuroimaging technique that I used throughout this PhD.

MRI is a powerful, non-invasive brain imaging modality that can produce detailed structural and functional images of the human brain. In this section I will explain the biophysical basis of the MRI signal, how this signal is harnessed to distinguish between different brain structures and functional states, and the spatial coding of the signal to produce 3D images of the brain. These explanations are based on the "Introducing MRI" course by the Albert Einstein College of Medicine (Albert Einstein College of Medicine, 2014b).

2.4.1. The MRI signal

The MRI signal originates from hydrogen protons within the water molecules of biological tissues. A hydrogen proton is a single positively charged proton. Its positive charge gives it a property called nuclear magnetism. Essentially, it will act like a small magnet.

Under normal background conditions, a population of hydrogen protons will be randomly oriented, resulting a net magnetisation of zero. However, when they are placed into the static magnetic field (B0) of an MRI scanner, they will either orient themselves in the same direction as the externally applied magnetic field or in the exact opposite direction. In a large population, an excess will orient in the opposite direction, resulting in a net magnetisation that is antiparallel to the externally applied magnetic field, known as the longitudinal magnetisation (Figure 13a).

Hydrogen protons also possess a quality called spin angular momentum. This means that when they are placed in a magnetic field, they will rotate, or precess, around the orientation of the magnetic field in a cone-shaped trajectory. The rate at which the hydrogen protons precess is known as the lamor frequency and is determined by both the properties of the hydrogen nucleus and the B0 strength. In a population of hydrogen protons in a static magnetic field, none will be precessing in the same phase, resulting in a net magnetisation that is parallel and opposite to the magnetic field (Figure 13b).

During MR imaging, radio frequency (RF) energy is put into the system. If this energy has the same frequency as the lamor frequency of the hydrogen protons, then the energy will pass

between all the hydrogen protons and they will start precessing in the same phase as one another. This phenomenon is known as resonance. It means that, now, as well as having a longitudinal magnetisation, the population of hydrogen protons will also have a transverse magnetisation (Figure 13c). The longer the RF pulse, the more hydrogen protons precess in the same phase and the greater the magnitude of the transverse magnetisation. As the longitudinal magnetisation decreases, the transverse magnetisation increases. This results in a net magnetisation that is diagonal, or sometimes even perpendicular, to the static B0 magnetic field. The degree of rotation of the net magnetisation is known as the flip angle.

When the RF pulse is turned off, the hydrogen protons will move back to their lower energy resting state in a process known as relaxation (Figure 13d). There are two types of relaxation which are independent of one another. Longitudinal relaxation describes the process by which the longitudinal magnetisation returns to its baseline and energy is transferred into molecules in the surrounding “lattice” of the tissue. The rate of longitudinal relaxation is determined by the time constant T1. Transverse relaxation is the process by which the transverse magnetisation returns to its baseline, zero, as hydrogen protons return to random phase. The time constant that governs the transverse magnetisation rate of decay is called T2. As the transverse magnetisation decays, RF energy is released, which can be detected by RF receiver coils. This is the MRI signal.

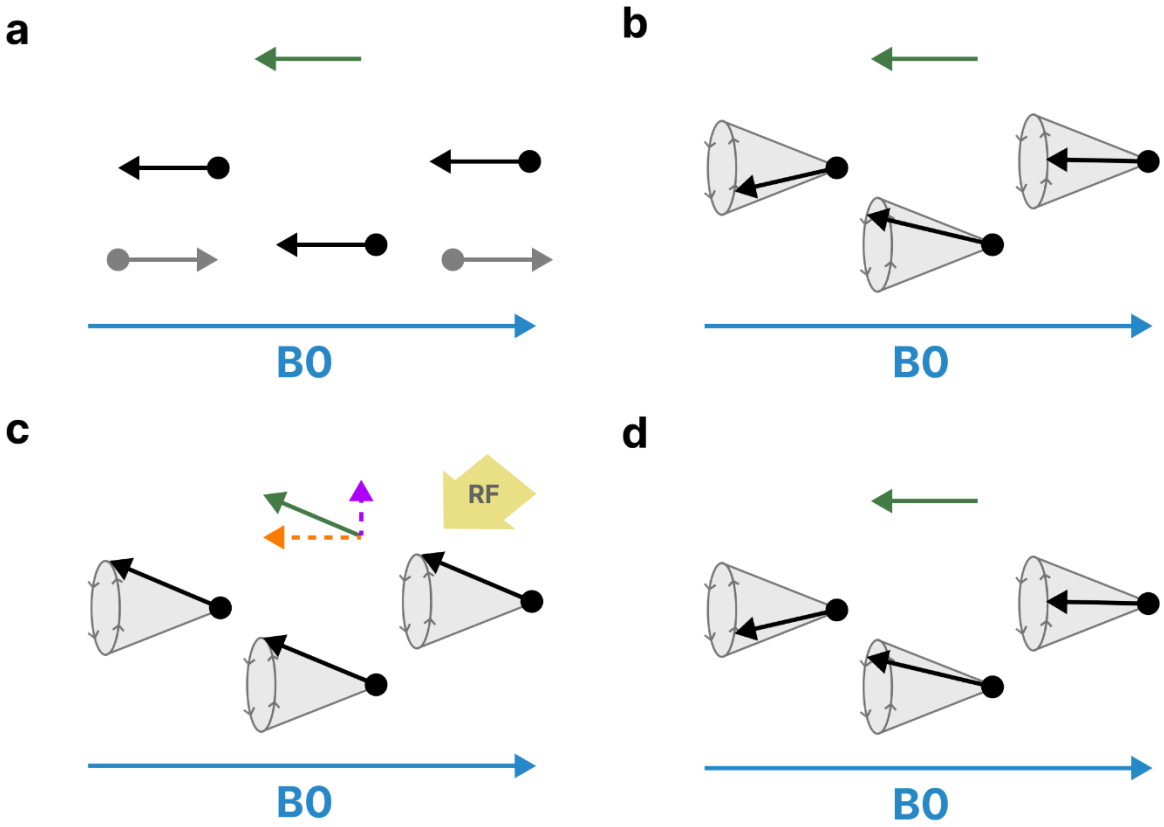

Figure 13. Generation of the MRI signal. (a) When placed in a static magnetic field (B0, blue line), hydrogen protons align with B0, with a relative excess in the opposite direction to B0 (black arrows). This results in a net magnetisation (green arrow) that is antiparallel to B0. (b) Hydrogen protons (black arrows) precess at the lamor frequency (depicted by the grey cones) at random phase. (c) When a radio frequency pulse (yellow arrow) is applied at the lamor frequency, the hydrogen protons (black arrows) precess at the same phase, resulting in a net magnetisation (green arrow) that includes both longitudinal (orange arrow) and transverse (purple arrow) components. (d) When the RF pulse is switched off, the hydrogen protons (black arrows) gradually return to random phase causing the net magnetisation to return to its baseline.

2.4.2. T1, T2, T2* and BOLD contrast

It is important to note that T1 and T2 are independent; longitudinal relaxation is much slower than transverse relaxation. The T1 of a tissue is determined by its lattice composition, whereas T2 is determined by its proton density, that is the concentration of hydrogen protons, typically from water molecules.

After a single excitation, the reading at the echo time (TE; the time between the RF pulse and the receipt of signal by the coil) will be a function of the T2. However, if a second excitation is administered before the longitudinal magnetisation has completely returned, then the next measurement at TE will be a function of both the T1 and the T2. The time between each RF

excitation is referred to as the repetition time (TR). Therefore, to achieve a better T1 contrast (i.e., to detect differences in T1 between different tissue types) it is optimal to have a short TR and TE. An MR image with good T1 contrast is referred to as a T1-weighted image. Whereas to achieve better T2 contrast (i.e., to detect differences in T2 between different tissue types), a long TR and TE are preferable. This would produce a T2-weighted image.

Importantly, the B0 magnetic field is not homogenous. This is because, the biological tissues can interact with and distort the externally applied magnetic field. The degree to which each biological tissue (e.g., brain, bone, fat) becomes magnetised is referred to as its magnetic susceptibility. When adjacent tissues have different magnetic susceptibilities, an additional magnetic field is produced, a so-called susceptibility field, which locally perturbs the B0 field, making it inhomogeneous. This means that the lamor frequency of the protons varies across the brain, as the lamor frequency of a proton depends on the magnetic field applied to it. So, after the protons are excited, they will gradually move out of phase with one another as they precess at different frequencies. This results in faster decay of signal compared to the T2 signal decay in a homogenous B0 field. The signal decay under conditions of B0 field inhomogeneity is known as T2* decay and is defined by the time constant T2*. Although it decays faster, in some cases, a reading of T2* is more desirable than a T2 reading, for example in BOLD contrast imaging, which I will explain next.

Hemoglobin, the molecule in red blood cells that carries oxygen, has a different magnetic susceptibility depending on whether it is oxygenated or deoxygenated. Oxyhemoglobin is weakly diamagnetic and does not distort the B0 field significantly. Whereas, deoxyhemoglobin is paramagnetic and does perturb the B0 field. As described above, these B0 field inhomogeneities increase the T2* decay time, causing a faster loss in MRI signal. This means that a larger ratio of oxyhemoglobin to deoxyhemoglobin results in a higher measure of MRI signal.

When neurons within a brain area become active, they metabolise oxygen from the local blood supply, reducing the ratio of oxyhemoglobin to deoxyhemoglobin and causing an initial dip in the MRI signal. Then, oxygenated blood is rerouted to the neurons, causing local blood vessels to dilate. This increase in blood flow increases the ratio of oxyhemoglobin to deoxyhemoglobin causing a large increase in the MRI signal. As the neurons metabolise this oxygen, the MRI signal decreases. The ratio dips below its baseline before gradually returning

to baseline. This entire process, known as the hemodynamic response (Figure 14), is used in fMRI analyses as an indirect measure of neural activity (Blamire et al., 1992). Being the most pronounced part of the hemodynamic response, fMRI analyses aim to detect the main peak in MRI signal that occurs around 6 seconds after the onset of neuronal activity.

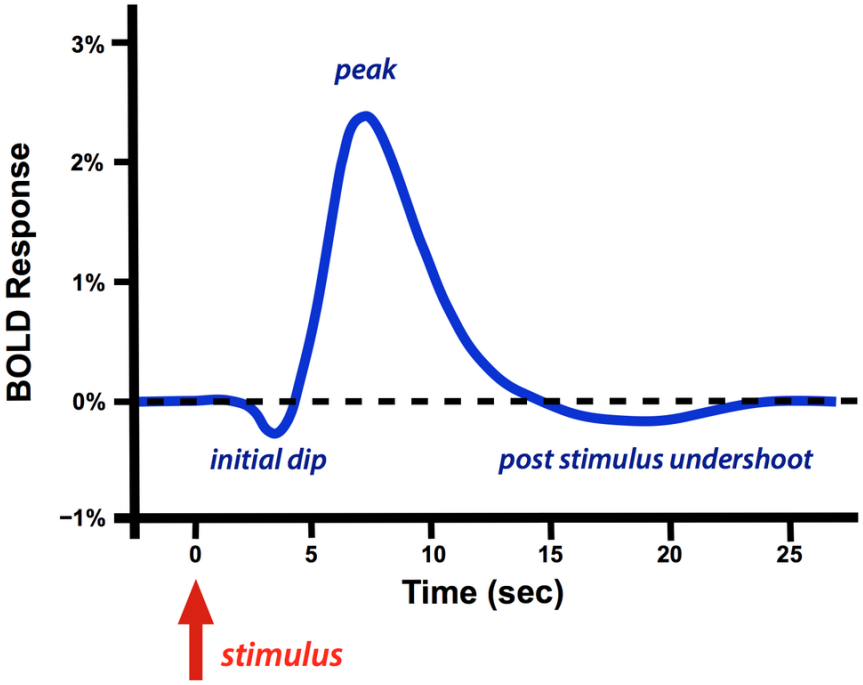

Figure 14. The haemodynamic response function. Following a single brief stimulus (red arrow), the hymodynamic response function (blue line) is characterised by an initial dip in the BOLD response followed by a large peak, then an undershoot before returning to baseline. Figure courtesy of Allen D. Elster, MRIquestions.com.

2.4.3. MR image formation

So far, I have explained how the detected MRI signal can be used to distinguish between different types of biological tissue or different metabolic states of the same tissue. Next, I will explain how the MRI signal from each unit of volumetric space, a 3D pixel known as a "voxel", is localised to form a 3D image of the brain.

The first step is slice selection. Recall that, in order to excite hydrogen protons, the RF pulse must be administered at precisely the lamor frequency of the protons, which is determined by the strength of the externally applied magnetic field. The strength of the magnetic field can be altered in small increments along a single axis of the scanner using a gradient coil. When the RF pulse is applied, only hydrogen protons located at a certain point along the axis will

become excited, the point at which the hydrogen protons are precessing at the same frequency as the RF pulse. The thickness of this slice of excited protons can be increased by increasing the RF frequency range or decreasing the gradient.

Once the slice has been excited, the slice selection gradient magnet is switched off. Then, at the moment the MRI signal is measured (i.e., TE), a second gradient magnet is switched on, this time along one of the in-plane axes of the slice. This is called the frequency encoding gradient as it causes the hydrogen protons to precess at different frequencies along the axis. A Fourier transform can be applied to the MRI signal to decompose it into the different contributing frequencies, allowing localisation of the MRI signal along the frequency encoding axis.

This leaves one final axis to locate the signal along, the second in-plane axis of the slice. After the slice gradient but before the frequency gradient, a third, phase-encoding, gradient along the final axis is switched on and then off. This causes the precessing hydrogen protons to either increase in frequency (i.e., accelerate) or decrease in frequency (i.e., decelerate) by a certain degree depending on their location along phase encoding axis, thereby knocking them out of phase from one another. The MRI signal is measured and the same slice is excited again. On this next repetition, a different strength of phase-encoding gradient is applied, resulting in a different rate of phase shift along the phase encoding axis. The process is repeated for multiple acquisitions, each using a different strength of phase encoding gradient. At a given point along the phase encoding axis, there may be a large difference in phase caused by the different gradients, whereas at another point, there will be a small difference in phase caused by the different gradients. A second Fourier transform is applied to the acquired data to decompose the signal into each magnitude of phase difference, allowing localisation of the MRI signal along the phase encoding axis.

This entire process is repeated for each slice. For each slice, the raw data is read out over time on a 2D grey-scale grid known as k-space, where the x axis is time, the y axis is the acquisition number (e.g., the acquired data with each different phase encoding gradient) and the grey scale of the data points is the MRI signal intensity value. The centre of k-space contains the highest amplitude signal intensities, while the corners contain the weakest signal. The two Fourier transformations, described above, are applied after the whole of k-space is acquired, along the x and then y axis of k-space to yield the MR image.

An issue with slice selection is that selective excitation of the desired slice is not possible. Instead, the signal will fade out at each end of the slice. To solve this issue, 3D imaging was developed. In 3D imaging, instead of exciting a single slice, the whole area that is being imaged, known as the slab, is excited. A slice gradient is then switched on and off prior to the phase encoding gradient. The same phase and frequency encoding gradients are applied as described above and the whole of the 2D k-space is acquired. This is repeated multiple times, using different slice gradient strengths and producing 2D k-space planes with different degrees of dephasing. These 2D k-space planes can be combined to form a 3D k-space, where the z axis is the slice acquisition number (i.e., the acquisition for each slice gradient). Three Fourier transformations are then used to localise the signal along the slice axis as well as the frequency and phase encoding axes, thereby producing a 3D image of the brain made up of voxels. The size of the voxels is determined by the slice thickness and/or the size of the k-space matrix.

2.5. 7T MRI opportunities and challenges

MRI scanning at higher field strengths (e.g., 7T) provides a lot of advantages, but also comes with some challenges.

2.5.1. Signal-to-noise ratio

The main advantage of increasing the magnetic field strength is the increase in MRI signal. This is because, when the B0 magnetic field is stronger, more hydrogen protons will align parallel to the magnetic field, resulting in stronger longitudinal magnetisation. This provides a larger population of precessing protons to be excited by the RF pulse, which also leads to stronger transverse magnetisation. The result is an overall increase in the MRI signal and in turn a higher signal-to-noise ratio (SNR).

The stronger B0 magnetic field also leads to a stronger susceptibility field as the magnetic susceptibilities of tissues are enhanced. This leads to an increase in T2* and BOLD signal and in turn an increase in temporal SNR (tSNR). In other words, this improves the ability to detect changes in the metabolic state of brain tissues, the proxy for neural activity.

2.5.2. Spatial resolution

A major appeal of 7T fMRI is its improved spatial resolution. It is important to recognise the difference between nominal spatial resolution and effective spatial resolution. Nominal spatial

resolution refers to the theoretical resolution based on the voxel size. Whereas, effective spatial resolution is the actual resolving power of the image, considering factors like SNR and artefacts. Effective resolution is typically lower than nominal resolution due to these influences.

As discussed previously, the voxel size depends on the size of the k-space matrix and the slice thickness (in 2D imaging). Decreasing the voxel size, decreases the SNR because smaller voxels are the accumulation of signal from a smaller number of hydrogen protons. Therefore, the increase in SNR achieved by the higher B0 field strength allows for smaller voxel sizes, while maintaining good effective resolution. The use of stronger gradient magnets in 7T MRI scanners also improves the spatial resolution as it allows the dissociation of signal between closer points along the gradient axes. Taken together, this means that 7T fMRI scanning achieves voxel sizes of less than 1 x 1 x 1 mm^3, compared to the typical 3 x 3 x 3 mm^3 voxel sizes used in 1.5T and 3T fMRI.

This improvement in spatial resolution means that neural activity in fine-scale structures such as cortical layers and hippocampal subfields can be imaged. The cortical grey matter is approximately 3 mm thick and has 6 distinct cytoarchitectural layers (Brodmann, 1909; von Economo & Koskinas, 1925). Although we are not yet able to resolve activity separately in the 6 cytoarchitectural layers, it is possible to dissociate activity across 3 different cortical depths: superficial, middle and deep.

2.5.3. Geometrical distortions and signal drop out

As described previously, the stronger susceptibility field achieved by high field strength imaging can improve our ability to detect the BOLD signal. However, it also has a downside. The B0 field inhomogeneities caused by the susceptibility field result in dephasing of the hydrogen protons. This dephasing can interfere with the phase encoding of the MRI signal. This in turn results in a mis-mapping of signal in k-space and subsequently in image space, making the brain appear deformed (Andersson et al., 2003; Gallichan et al., 2010; Jezzard & Balaban, 1995; Jones & Cercignani, 2010). These deformations generally manifest as stretching (signal spreading into non-brain voxels) or compression (signal piling up in brain voxels) along the phase encoding axis.

Dephasing also results in faster decay of the transverse magnetisation. In brain areas with a very inhomogenous B0 field, the high level of dephasing can mean that the MRI signal is very low or has even reduced to zero by the time it is measured at TE. This results in areas of the image with very low intensity or no signal, known as signal drop out.

The brain regions most severely affected by geometrical distortions and signal drop out are the orbitofrontal cortex and inferior temporal cortex. This is because, these areas are close to the sinuses and ear canals, where biological tissues (bone and air-filled cavities) have very different magnetic susceptibilities.

2.5.4. Signal artefacts

Several types of artefact are more common at higher field strengths. For example, susceptibility-induced artefacts related to increases in B0 field inhomogeneity, and motion artefacts due to the increased sensitivity to motion.

Because B0 field inhomogeneities cause a mis-mapping of signal into k-space, Gibb's ringing or truncation artefacts can emerge (Czervionke et al., 1988; Wood & Henkelman, 1985). These are caused by abrupt changes in signal intensity between adjacent data points in k-space, which can occur when high intensity signal is mispositioned around the edges of k-space where the signal typically becomes zero.

It is important to minimise head motion during fMRI scanning so that each data point in k-space and each voxel in image space correspond to the same portion of brain tissue throughout the fMRI time course. Motion artefacts arise when the brain moves during acquisition of k-space, often manifesting as “ghosting” (repeating of the brain or parts of the brain) in the direction of phase encoding. Even sub-voxel motion can severely degrade the quality of the scan. At higher field strengths, voxel sizes are typically very small which reduces the margins for head motion.

Head motion also causes the susceptibility field to change and thereby alters the nature of distortions (Andersson et al., 2001). This means that the shape of the brain is inconsistent across different volumes of the fMRI time series, making it difficult to bring voxels back into alignment post-acquisition. At higher field strengths, the increased susceptibility fields increase the impact of susceptibility-motion interactions.

Another pervasive signal artefact that poses significant challenges for laminar fMRI studies is the so-called superficial layer bias effect or venous signal artefact (Kay et al., 2019; Norris & Polimeni, 2019; Ugurbil, 2016). The majority of laminar fMRI studies acquire data with GE-BOLD sequences as these are highly sensitive to changes in neural activity (Olman & Yacoub, 2011). However, in GE-BOLD fMRI, most of the measured BOLD signal arises from large veins (Boxerman et al., 1995). Because there are more large veins in the superficial layers, and blood drains from the deep to superficial layers (Figure 15; Duvernoy et al., 1981), the measured BOLD signal is highest in the superficial layers.

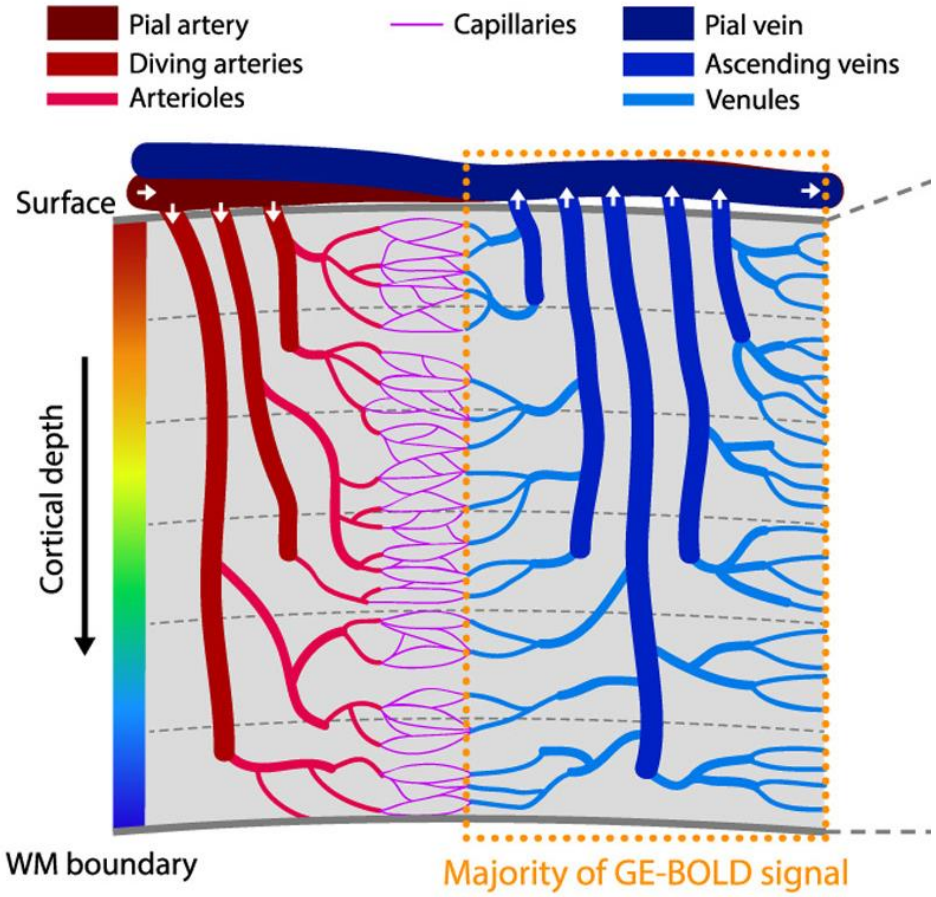

Figure 15. Vasculature of the cortical grey matter. Arteries (red) supply oxygenated blood to the cortical grey matter. These arteries branch into smaller capillaries (purple) that run parallel to the cortical layers. Deoxygenated blood is transported out of the grey matter by veins (blue). The measured BOLD signal in GE fMRI comes from deoxygenated blood in the veins. The concentration of large veins, running perpendicular to the cortical layers, increases from the deep to the superficial layers, and consequently so does the BOLD signal. Additionally, the movement (white arrows) of deoxygenated blood towards the pial surface causes leakage of BOLD signal from the deep to the superficial layers. Figure from Havlicek & Uludağ (2020).

Several methods that attempt to correct or mitigate the superficial layer signal bias have been developed and are explored in Experiment 1B (Chapter 3). These include: masking out voxels that are likely to measure signal from large veins (e.g., tSNR, t value, and T2 signal masks; Jia et al., 2021; Olman et al., 2007); using ratios rather than subtractions when contrasting task conditions (Huang et al., 2021; Kashyap et al., 2018; Liu et al., 2021); and regressing out signals that are correlated across layers in the GLM (Kok et al., 2016).

2.6. MRI scanner and sequences

2.6.1. MRI scanner

For my 7T experiments, MRI data were acquired on a Siemens Magnetom Terra 7T MRI system (Siemens Healthcare GmbH, Erlangen, Germany) with an 8-channel head coil for localised transmission, operating in a quadrature-like ("TrueForm") mode, with a 32-channel head coil insert for reception (Nova Medical, Wilmington, USA).

In Experiment 2 (Chapter 4), the 3T MRI data were open access (Clark & Maguire, 2023) and acquired using three Siemens Magnetom TIM Trio 3T MRI systems with 32 channel head coils. The three scanners used the same software and MRI sequences.

2.6.2. MRI sequences

2.6.2.1. Functional MRI

The 7T fMRI data were high resolution with a voxel size of 0.8 × 0.8 × 0.8 mm^3, which was sufficient spatial resolution to extract signal from 3 cortical depths - superficial, middle and deep - as well as the hippocampal subfields. To achieve this high resolution, and maintain a reasonable TR, only part of the brain was imaged. The partial volume was always positioned to include the hippocampus and mPFC, which was important to address the research questions of all 3 experiments. By virtue of the anatomy of the brain, it also always covered the retrosplenial, parahippocampal, occipital, lateral temporal and lateral frontal cortices. The use of isotropic voxel dimensions provided equal resolution in all three imaging planes, which was important for imaging convoluted structures like the cortical layers and hippocampal subfields.

The specific sequence used to acquire the fMRI data was a T2*-weighted 3D GE echo-planar imaging (EPI) sequence (Mansfield, 1977). In GE-EPI, the whole of k-space is acquired continuously during a GE train following a single RF pulse (Albert Einstein College of Medicine, 2014a). In a GE train, rather than administering an RF pulse for each acquisition, a gradient is used to rephase the precessing hydrogen protons. This produces an "echo" of the MR signal. A series of echoes can be used to iterate through the phase encoding acquisitions. GE-EPI is an extremely fast acquisition method in MRI, making it a popular choice for functional imaging. We achieved a volume acquisition time of 3872 ms, a TR of 44 ms and a TE of 18.7 ms. However, this fast imaging comes at a cost. Because the entire MR signal is acquired during a long train of gradient echoes, there is more time for dephasing to interfere with phase encoding. This means that the GE-EPI images are highly susceptible to geometric distortion.

We used parallel imaging acceleration, meaning that specific data points in k-space were not sampled and were instead inferred from the surrounding data points (Deshmane et al., 2012). This speeds up the time taken to acquire k-space, thereby reducing the time over which B0 inhomogeneities can cause geometric distortion and signal drop out. However, the process of reconstructing images from under sampled data can amplify noise. As a result, the SNR in the reconstructed images can be lower compared to fully sampled acquisitions. We applied parallel imaging in both the direction of the partition (i.e., the slice direction, factor = 2) and the phase encoding direction (factor = 4).

Because k-space is, in theory, symmetrical, a way to save time is by only sampling (a minimum of) half of k-space, which would correspond to half the number of phase encoding iterations. This is called partial Fourier imaging (McGibney et al., 1993). We used partial Fourier 6/8 in the phase-encoded direction, meaning that only 75% of k-space is acquired.

In order to include the mPFC and hippocampus in all participants (each with unique brain anatomy) we used a field of view of 192 × 192 × 70.4 mm^3 and a slab angle of 15°. Another MRI parameter was a flip angle of 14°.

We used binomial (1331) water-selective excitation (Hore, 1983). This technique uses a specific type of RF pulse sequence known as a binomial pulse. The numbers in the name (1331) refer to the amplitude and phase configuration of the sub-pulses within the binomial sequence. These pulses are designed to selectively excite water protons while suppressing the

signal from fat protons by taking advantage of their different lamor frequencies (when placed within the same external B0 magnetic field).

We used advanced B0 shimming. Shimming is an approach is to homogenise the B0 field before image acquisition (Juchem & de Graaf, 2017). It involves the application of small gradients using a separate set of coils. We used an advanced algorithm for estimating the necessary corrective gradients.

Anterior-to-posterior phase encoding was used, meaning that the phase-encoding gradient was applied along the scanner axis that corresponded to the anterior-posterior axis of the brain. Four images, with the same parameters but the opposite polarity (posterior-to-anterior) of phase encoding, were acquired at the start of the sequence. Different polarities of phase encoding lead to different manifestations of geometric distortion in the images. These differently distorted images can then be used to correct geometric distortion in a method known as "unwarping" (see Section 2.7.2).

2.5.2.2. Anatomical MRI

In my 7T fMRI experiments, a whole brain magnetisation transfer (MT)-weighted EPI image was also acquired (Ranjeva et al., 1997). This image was acquired on the 7T MRI scanner using a T2*-weighted 3D GE EPI sequence with many of the same parameters used in the fMRI sequence.

An MT-based contrast was used. This means that tissues with significant interactions between water and macromolecules (e.g., proteins) are suppressed. Because grey matter, white matter and cerebrospinal fluid (CSF) are all suppressed to different degrees, the MT-weighting amplifies the contrast between these tissue types.

The volume acquisition time of the MT-weighted EPI image was 3 m 45 s, the TE was 16.97 ms, the flip angle was 8°, the field of view was 192 × 192 × 128 mm^3, and parallel imaging (factor = 4) was used.

Because it had many of the same parameters (e.g., TE and acceleration factor) as the fMRI images, the MT-weighted EPI image was matched to the fMRI images in terms of geometrical distortion. This meant it could be aligned to the functional images and then used as a high

contrast functional image template for coregistering the anatomical images to the functional images (see Section 2.9.2.3 for explanation of coregistration).

In my 7T experiments, a whole brain, submillimetre resolution (voxel size = 0.65 × 0.65 × 0.65 mm^3) T1-weighted image was acquired using a Magnetization Prepared 2 Rapid Acquisition Gradient Echo (MP2RAGE) sequence (Marques et al., 2010). The volume acquisition time of 8 m 42 s, TR of 5000 ms, TE of 2.54 ms, inversion time of 900 ms and 2,750 ms, flip angles of 5° and 3°, field of view of 208 × 208 × 156 mm^3, and in-plane GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) acceleration (factor = 3). This whole brain anatomical image was used to define the two cortical surfaces, the boundary between grey matter and CSF, and the boundary between white matter and grey matter, which were later used to define the three cortical depths.

In my 7T experiments, two partial brain, submillimetre resolution (voxel size = 0.52 × 0.52 × 0.5 mm^3) T2-weighted images were acquired using the 7T MRI scanner for the purposes of hippocampal subfield segmentation. These images were acquired using a T2-weighted 3D spin echo sequence with variable flip angles (SPACE) sequence (Mugler et al., 2000) with a volume acquisition time of 10 m 25 s, TR of 3,500 ms, TE of 229 ms and flip angle of 8°. The field of view of 200 × 169 mm x 56 mm and slab angle of 15° captured the hippocampus in all participants. The sequence used partition oversampling of 14.3%, and in-plane GRAPPA acceleration (factor = 2).

2.7. Preprocessing of fMRI data

2.7.1. Motion correction

To correct for any head movements during scanning, the functional images must be aligned to one another such that the same voxel in each image contains the same anatomy. This process is known as realignment. It involves estimating six parameters, 3 translation parameters and 3 rotation parameters, that describe how an image must be moved so that it is in alignment with a reference image. The realignment algorithm estimates these parameters by minimising the differences in BOLD signal intensity between the images. First, all functional images are realigned to the first functional image in the time series. Then an average functional image is produced by calculating the mean BOLD signal intensity across the time

series for each voxel. A second pass of realignment is then performed with respect to the average fMRI image.

2.7.2. Distortion correction

At the same time as realignment, fMRI data are "unwarped". Unwarping corrects for geometric distortion in the fMRI data. There are several methods for unwarping, two of which involve estimating a B0 field map. In the most common method, the B0 field is directly measured by acquiring a field map image (Hutton et al., 2002; Jezzard & Balaban, 1995; Zeng & Constable, 2002). In another approach, images are acquired with the opposite phase encoding direction to the functional images (Andersson et al., 2003). All of our functional images were acquired with anterior-posterior phase encoding and so we acquired up to 4 additional images with posterior-anterior phase encoding. By combining the posterior-anterior phase encoded images with the same number of anterior-posterior phase encoded images, it is possible to estimate the field that has produced these data (Chang & Fitzpatrick, 1992). This can be done using the "topup" algorithm implemented in the FMRIB Software Library (FSL; Andersson et al., 2003). The field map produced by either of these approaches can then be used to produce a model of the movement-by-inhomogeneity interactions, that is, how the B0 field changes with participant movement, which can then be applied to the functional images to correct distortion (Andersson et al., 2001; Hutton et al., 2002). In Malekian et al. (2023), we pitted these two approaches against each other using a subset of the 7T fMRI data that I acquired in Experiment 4 (Chapter 6). We found superior distortion correction using the opposite phase encoding method compared to the B0 field map method, particularly in the brain areas with most severe susceptibility-induced distortions, such as the mPFC. Therefore, in all experiments reported in this thesis, I used the opposite phase encoding approach to distortion correction.

In my 7T experiments, realignment and unwarping steps were applied to the fMRI data in a single step. This is particularly important for cortical layer and hippocampal subfield analyses. This is because preprocessing steps that require the resampling of fMRI data onto a new voxel grid require the signal intensities corresponding to each voxel to be interpolated. Each time the data are resampled and interpolated, the effective resolution of the data decreases (Wang et al., 2022). This spatial smoothing can blur the signal between layers and prevent detection of layer-specific neural activity.

Note that geometric distortion is particularly problematic for cortical layer analyses, where it can blur the BOLD signal between layers. It is also a major concern in one of my main brain regions of interest, the mPFC (see Section 2.5.3 for explanation of this). Therefore, additional preprocessing steps are taken prior to cortical layer analyses (described later in Section 2.9.2.3) to correct for any remaining distortion and to prevent this remaining distortion from impacting analyses.

2.8. Whole-brain fMRI analyses

Although the main analyses within this thesis are at the level of cortical layers and hippocampal subfields (described in Section 2.9.1), I will first describe the standard whole-brain (or in my case, partial brain) mass-univariate approach to fMRI analysis, which can be used to identify the brain areas that are active during certain cognitive processes.

First, some further preprocessing steps are required to move the functional images of all participants into anatomical alignment with one another. This alignment is crucial to allow us to find specific brain anatomy that activates across all participants. To achieve this, first, for each participant separately, the BOLD signal associated with specific cognitive processes is estimated for each voxel, using a general linear model (GLM; Frackowiak et al., 2004). The results from this first-level analysis are then summarised into a single contrast image for each participant and taken to the second level. At the second level, classical statistical methods are applied to each voxel.

2.8.1. Additional preprocessing

The functional images of all participants are brought into alignment with one another in two steps. First, the images are "warped" into approximate anatomical alignment. As each participant has unique brain anatomy, this involves geometrically distorting the images. Next, because it is not possible to precisely align the voxels between participants with warping, the images are blurred, or "smoothed", so that the anatomy is coarsely aligned.

Note that although the nominal spatial resolution of the images has not changed, this smoothing significantly reduces the effective spatial resolution and results in a blurring of the BOLD signal between the cortical layers and between the hippocampal subfields. Furthermore, it may remove task-relevant information that is represented in the unsmoothed

multivoxel patterns of signal, which is the basis of MVPA analyses. For this reason, all laminar and subfield analyses and all MVPA analyses were performed within the native space of each participant's fMRI data, as described in Section 2.9.1.

2.8.1.1. Spatial normalisation (warping)

To approximately align their anatomy, the functional images of all participants are warped to a standard brain template in Montreal Neurological Institute (MNI) space.

There are several steps to achieve this. First, for each participant, the functional images are brought into alignment with the MP2RAGE anatomical MRI image, which has higher spatial resolution and much clearer detail of the brain's anatomy. This alignment is performed using a process called coregistration, specifically, normalised mutual information (NMI) coregistration (described in detail in Section 2.9.2.3). Then the different tissue types (e.g., grey matter, white matter and CSF) are delineated, or segmented, on the anatomical image using the SPM segment tool. These tissue segmentations are used to warp the anatomical images of different participants into alignment with one another using an SPM tool called DARTEL (Ashburner, 2007). The result of this can then be transformed into MNI space. Finally, the same warp and transformation is then applied to the functional images.

2.8.1.2. Spatial smoothing

Finally, the functional images are spatially smoothed. This involves applying a Gaussian smoothing kernel to each voxel. The full-width at half maximum (FWHM) of the kernel governs the degree of smoothing and typically ranges from 4 to 12 mm.

As described previously, spatially smoothing the images accounts for the fact that the spatial normalisation process is not able to precisely align each participant's anatomy at the level of individual voxels. Smoothing the images, blurs the anatomy, thereby bringing it into alignment across participants. Smoothing is also essential for the group-level statistical inference (Section 2.8.2.2). This is because the group-level statistical inference uses Gaussian random field theory, which assumes that the fMRI images are smooth. Therefore, smoothing is necessary to ensure the validity of the results.

2.8.2. Mass-univariate fMRI analysis

The mass univariate approach to fMRI analysis involves analysing the BOLD signal time course at each voxel independently using statistical models. This allows for identification of clusters of voxels (i.e., areas of the brain) where the neuronal activity is caused by the specific experimental conditions.

2.8.2.1. Single-subject analysis

Single-subject fMRI analysis examines the relationship between experimental conditions and the observed BOLD signal of individual participants. A general linear model (GLM) is used to model this relationship in each voxel independently. The basic form of the GLM can be written as:

$$Y = X\beta + \varepsilon$$

In this equation, Y is the observed BOLD signal in the form of a vector of BOLD signal values over the time course of the experiment. X is the design matrix, representing the list of variables that might explain the changes in the BOLD signal. β is a vector of values, which are estimated when the model is fitted to the data. These beta estimates represent the estimated contribution of the explanatory variables to the observed BOLD signal. Finally, ε is the error term, also known as the residuals, which is also estimated when the GLM is fitted to the data. These residuals represent any noise in the data that is unaccounted for by the explanatory variables.

The first step in the GLM analysis is construction of the design matrix. Each column of the design matrix corresponds to a different explanatory variable, referred to as a regressor. These could be task conditions (e.g., memory retrieval or counting), task events that are of no interest (e.g., presentation of instructions), so-called nuisance regressors that could have confounding effects (e.g., head motion parameters) or drift terms to account for low-frequency noise. The rows of the design matrix correspond to each time point at which a functional image was acquired. The data points in the design matrix represent the on and off state of the regressor. For task-related regressors, each onset of the regressor corresponds to a task event (e.g., stimulus presentations, tasks, or participant responses or presentation of instructions). The events can either be represented as a continuous event of a specified duration, using a square wave or “boxcar” function, or as a transient event, using a stick (delta)

function. These event regressors are then convolved with the canonical HRF to account for the delayed and dispersed nature of the BOLD signal relative to the neuronal activity that caused it. This step transforms the discrete event timings and durations into the predicted BOLD signal changes associated with each regressor.

Using the specified GLM, the beta values are estimated for each voxel within the functional image. This involves finding the beta values that minimise the difference between the predicted and observed time course of BOLD signal, using a maximum likelihood estimation. A beta value is estimated for each regressor. It represents the degree to which the regressor explains the BOLD signal changes of the voxel. It can be thought of as the magnitude of neuronal activation caused by the task condition.

Hypotheses about the conditions are tested using contrasts, which are linear combinations of the beta values. For example, to test the difference between two conditions, memory retrieval and counting, for each voxel, a contrast vector of [1 -1] is applied to the beta estimates of memory retrieval and counting, respectively. The result is a contrast value that represents the weighted combination of the memory retrieval and counting beta values. To assess the statistical significance of the contrast values, a t-test is commonly used. The t-test evaluates whether the contrast value for each voxel is significantly different from zero. For more complex hypotheses involving multiple contrasts, an F-test can be used. The F-test evaluates whether a set of contrasts is jointly significant. This is particularly useful for testing whether there are differences between any of the task conditions, rather than testing for pairwise differences between specific conditions. Once the statistics have been computed for each voxel, the result is a statistical map, where each voxel has a t-value (or F-value) indicating the strength of the effect.

To identify significant voxels, the statistical map is thresholded. This involves selecting a threshold value for the t-statistics (or F-statistics) that corresponds to a desired significance level (e.g., $p < 0.01$). Given the large number of voxels, it is important to correct for multiple comparisons to control the false positive rate. The most common correction method is the Family-Wise Error (FWE) correction.

2.8.2.2. Group-level analyses

Group-level fMRI analysis, also known as second-level analysis, aims to generalise the findings from individual participants to a broader population. This involves combining the first-level results from multiple participants to make inferences about the average effect and variability across the group.

A new GLM is constructed for the group-level analysis and is applied to each voxel separately. The primary input to the group-level GLM are contrast maps from the first-level analysis of each participant. These images represent the effect sizes for specific condition contrasts (e.g., memory retrieval versus counting). The design matrix at the second level represents the experimental conditions and covariates for the group analysis. The setup depends on the research question and the structure of the data. For all of the experiments reported in this thesis, I used the most simple and common design matrix, a single column of ones. For each voxel, the second level GLM is fit to the participant contrast values. This involves estimating the beta value for the single regressor of the design matrix. The beta value will correspond to the mean contrast value. A t-statistic is then computed to test the whether the beta value (i.e., the mean contrast value of the group) is significantly different to zero. Note, more complex GLMs with multiple regressors can be used to account for additional explanatory variables (e.g., age, gender, clinical scores). Such complexity was not warranted in my case, as all the participants were healthy young adults, with little variance overall in terms of age or gender.

The output of the group-level analysis is a group-level statistical map of t-values. As with the single-subject analysis, this map can be thresholded to identify significant voxels, and correct for the fact that multiple statistical tests have been performed (one for each voxel). As with single-subject analysis, the most common multiple comparisons correction method is FWE correction, which controls the false positive rate.

2.9. Region of interest fMRI analyses

By focusing analyses on predefined regions of interest (ROIs), it is possible to examine directly the involvement of these specific ROIs in particular cognitive processes. There are several advantages to an ROI-based approach compared to the whole-brain mass-univariate method described in Section 2.8. First, by averaging signals within ROIs, ROI-based analyses can increase the SNR and enhance sensitivity to subtle effects, particularly in regions with

relatively low signal intensity. Second, compared to whole-brain analyses, ROI-based analyses reduce the number of statistical comparisons, which can help mitigate the issue of multiple comparisons and increase statistical power. Finally, and perhaps most importantly for the experiments reported in this thesis, ROI-based analyses can be performed in the native space of the functional images of each participant. This yields a single result per ROI (e.g., per mPFC layer) for each participant, which can then be analysed at the group-level. This differs from the mass-univariate approach in which the individual voxels must be aligned across participants before performing statistical analyses for each specific voxel. This requires significant smoothing of the functional images, resulting in a substantial loss of spatial resolution which precludes the investigation of cognitive processes at the level of microstructures, that are of particular interest to me in this thesis.

The ROIs are typically defined based on (1) the aims and hypotheses of the experiment, (2) previous findings and/or (3) results from whole-brain analyses performed in the same experiment. Caution is required if using analyses from the same experiment. It is important that either a different set of data is used to select the ROI (e.g., from a functional localiser task) or, as I have done here, that structurally defined ROIs are used, to avoid “double dipping” (Kriegeskorte et al., 2009). Double dipping refers this improper use of the same dataset and condition contrast to both define the ROI and perform statistical inference on the ROI. This practice can lead to biased and inflated results because it violates the independence assumption of statistical tests. In all of my experiments, the ROIs were guided by all 3 of these selection criteria noted above. I set out to investigate the involvement of the mPFC layers and hippocampal subfields in event processing, and therefore included the 3 mPFC layers and 6 hippocampal subfields. Furthermore, in exploratory analyses, I included the 3 layers of other cortical areas within the core network, drawing upon the wealth of existing literature reporting their functional significance.

In this section, I first describe the steps taken to define the ROIs in the native functional space of each participant. Once the ROIs have been defined, they can be used in univariate and multivariate analyses; mine are described in detail in Sections 2.9.2 and 2.9.3. Finally, in Section 2.9.4, I describe a method for investigating functional connectivity between ROIs.

2.9.1. ROI definition

To delineate the ROIs, a segmentation process is performed on the anatomical images whereby anatomical brain structures are more easily distinguishable. The segmentations must then be moved into alignment with the functional images in a process called coregistration. Once in the same space as the functional images, the ROIs are refined and the cortical ROIs are divided into 3 cortical depths: deep, middle and superficial.

2.9.2.1. Segmentation of cortical areas and surfaces

The brain is first segmented into various cortical areas. Furthermore, in order to (later) divide the cortical grey matter into 3 cortical depths (see Section 2.9.2.5), I first defined the outer and inner boundaries of the grey matter, that is the pial surface, the boundary between grey matter and CSF, and the white matter surface, the boundary between white matter and grey matter. To achieve this, it is important that the anatomical image used to segment the cortical surfaces and areas has high contrast between grey and white matter and between grey matter and CSF. A "spliced" whole brain T2-weighted image was used, which combined several anatomical images. First the MP2RAGE T2-weighted image was moved into alignment with the two high resolution partial brain T2W images, using a NMI coregistration (see Section 2.9.2.3 for a full description of coregistration). These images were then averaged to produce a high resolution, high contrast T2-weighted image, with good contrast between grey matter and CSF. The MP2RAGE T1-weighted image was also used for segmentation, as it has good contrast between grey and white matter. To remove noise from the "spliced" whole brain T2-weighted image and the MP2RAGE T1-weighted image, the images were denoised and minimally smoothed (1 mm FWHM, see Section 2.8.1.2). Finally, the spliced whole brain T2-weighted image and the MP2RAGE T1-weighted image were passed through a segmentation algorithm called FreeSurfer recon-all, which, after running several preprocessing steps, delineates the two cortical surfaces and many cortical areas (Fischl et al., 2002; Fischl et al., 2004). These cortical surfaces are in the form of a mesh, a collection of vertices (points), edges (lines connecting vertices), and faces (triangles or polygons formed by edges) that create a 3D model of a surface. Two brain atlases were used for segmenting the cortical areas on the pial surface mesh: the Desikan-Killiany Atlas and the Destrieux Atlas (Desikan et al., 2006; Destrieux et al., 2010).

2.9.2.2. Segmentation of hippocampal subfields

In all experiments, I manually segmented the hippocampal subfields according to an established protocol that is accurate to our current understanding of hippocampal anatomy (Dalton et al., 2017).

First, the high resolution T2-weighted images were visually inspected and poor-quality images were discarded. Then, they were denoised and averaged to produce a single T2-weighted image with good SNR. This image was used to manually delineate the hippocampal subfields using ITK Snap software version 3.2.0. For each participant the whole hippocampus, bilaterally, was manually segmented into 6 subregions: DG/CA4, CA3/2, CA1, subiculum, pre/parasubiculum and uncus (Figure 16). On average, it took eight hours to segment the hippocampi of a participant.

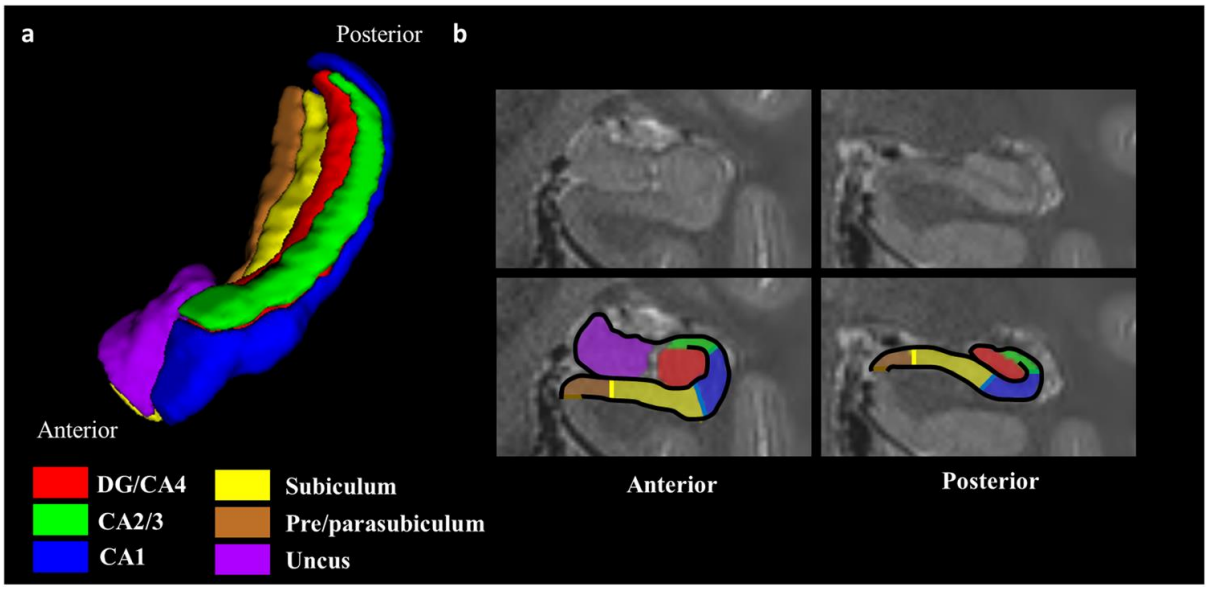

Figure 16. Segmentation of the hippocampal subfields. (a) A 3D representation of the segmented hippocampal subregions. (b) Two coronal slices of T2-weighted structural MRI scans, showing the anterior (left) and posterior (right) hippocampus. The raw image (top) and the same image with hippocampal subfield segmentations (lower) are shown. DG = dentate gyrus, CA = Cornu Ammonis. Figure from Clark & Maguire (2023).

In Experiment 2 (Chapter 4), I used these manually segmented anatomical scans to train an algorithm to automatically segment the hippocampal subfields in 7T MRI data. I also trained an algorithm to automatically segment the hippocampal subfields in 3T MRI data that had been previously acquired and manually segmented by another researcher. In so doing, I hope

that the hippocampal community will find this a useful way to save a substantial amount of time. The method for automatic segmentation is described in full in Chapter 4.

2.9.2.3. Coregistration

Having segmented the relevant brain areas on the anatomical scans, it is next necessary to align these anatomical images and their segmentations with the functional data in a process called coregistration. A choice must be made about whether to move the anatomical image into the space of the functional image or vice versa. In all experiments, I moved the anatomical images to the space of the functional images, to avoid unnecessary interpolation and resampling the functional data. In coregistration, the two images being aligned have different properties (e.g., an MP2RAGE image and a functional image). This differs from realignment, where all images are acquired with the same MRI sequence, giving them the same properties (e.g., functional images, see Section 2.7.1).

There are two main types of coregistration method, NMI and boundary-based registration (BBR), each with their strengths and limitations. The choice between them depends on the specific application and the characteristics of the images being coregistered. Both methods estimate the translations and rotations in each axis that are required to move the one image into alignment with the other. In the more commonly-used method, NMI, the transformation is estimated by maximising the statistical dependence, or mutual information, between the intensity values of the corresponding voxels in the two images. Whereas, in BBR, the focus is on matching the edges or boundaries of structures in different images. In order to do BBR, cortical surfaces must have been defined on one of the images being coregistered. BBR then maximises the signal intensity gradients across the cortical surfaces in the other image. As such, BBR is more effective at matching the edges or the images, in particular the cortical grey matter, while NMR is more accurate at matching the subcortical structures like the hippocampus.

Recall that the MT-weighted EPI image has the same geometric distortion and signal drop out as the functional data. However, the MT-weighting provides it with a considerably higher contrast-to-noise ratio. This makes it a good template, or reference, for moving the anatomical images and their segmentations into alignment with the functional data. However, first, the MT-weighted EPI image must be moved into alignment with the functional data. To do this, it

is first distortion corrected in the same way that the fMRI data was, that is, using the anterior-to-posterior and posterior-to-anterior images to estimate a field map, and then using the field map to create a single distortion corrected image (see Section 2.7.2). It is then coregistered, using NMI, to the average functional image that was produced during realignment of the functional data.

In all experiments, NMI was used to coregister the T2-weighted structural image with the MT-weighted image (and consequently the fMRI data). The same transformation was then applied to the hippocampal subfield segmentations such that they were also brought into alignment with the functional data.

The cortical surfaces and cortical area segmentations were coregistered to the functional data in 3 steps. First, using NMI, coregistration of the MP2RAGE T1 image to the MT-weighted image was estimated. Then, using BBR, the cortical surfaces were coregistered to the MT-weighted image, estimated using FreeSurfer BBRegister (Greve & Fischl, 2009). In a final step, BBR was applied recursively to improve coregistration of the cortical surfaces to the MT-weighted image. This works by dividing the cortical surface mesh into increasingly smaller subsections and performing BBR iteratively on each section of mesh (van Mourik, Koopmans, et al., 2019). Six iterations of BBR were performed. In each iteration, scaling was allowed along the phase-encoding direction, because in GE-BOLD EPI, geometrical distortion is most severe in the phase-encoded direction (Jezzard & Clare, 1999; Mansfield, 1977; Pykett & Rzedzian, 1987). Allowing for scaling in the phase encoding direction and coregistering the surfaces to subsections of cortex corrects for any distortion remaining in the functional data by essentially distorting the cortical surfaces to fit the functional data.

2.9.2.4. Definition of cortical area ROIs

Now that the cortical surface and area segmentations are in the space of the functional images, they can be processed into masks for each ROI. A mask is a voxel map of ones and zeros, ones for the voxels in the ROI and zeros to all other voxels. To produce these masks, first, the cortical area segmentations must be translated from the space of the cortical surfaces (i.e., in the form of meshes) into the space of the functional volumes (i.e., in the form of volumes comprised of voxels). This is done by projecting them across the surface normal, a 3D vector that runs perpendicular to the surface. As the cortical grey matter has undulations,

gaps will exist in the cortical area masks. These can be filled in using FSL Maths Edge, Dilate and Erode operations. Where appropriate, manual edits were made to these structural masks.

2.9.2.5. Definition of cortical layer ROIs

Next, the cortical area masks needed to be divided into 3 cortical depths for laminar analysis. However, before I did this, some edits were made to them to remove voxels that may have contaminated the laminar results. For example, voxels that were expected to contain large veins were removed from the masks, to reduce the impact of the superficial layer bias effect (see Section 2.5.4). Furthermore, voxels that were in areas with remaining distortion were removed from the masks, as they would introduce noise into the results.

Voxels with a low tSNR are expected to contain large veins (Jia et al., 2021). Therefore, voxels with a tSNR lower than the 65th percentile of the tSNR distribution were excluded from the masks.

Despite performing distortion correction, there were still some areas of the brain, namely the mPFC and MTL, that remained somewhat distorted. To prevent these voxels from contaminating the analyses, they should be removed from the masks. To do this, a Jacobian map was estimated from the fMRI data using the HySCO tool in the ACID MATLAB toolbox (http://www.diffusiontools.org/). The Jacobian is a measure of the displacement of the signal contained within a voxel and captures the degree of compression or stretching that has occurred (Jezzard & Balaban, 1995). It can therefore be used as a quantitative measure of distortion (Clark et al., 2021). Voxels with Jacobian values lower than the 5th percentile and higher than the 95th percentile of the Jacobian distribution were considered to be highly distorted and were therefore removed from the mask.

Next, the cortical grey matter was divided into 3 equivolume depth bins or layers (deep, middle and superficial) using the Open fMRI Analysis toolbox (https://github.com/TimVanMourik/OpenFmriAnalysis). Equivolume layers were used for anatomical validity - within small sections of cortical grey matter, the ratio of layer volumes remains constant whereas the ratio of layer thicknesses depends on the curvature of the cortex (Waehnert et al., 2014). To produce the layer boundaries, two additional surfaces were produced between the pial and white matter surfaces using the level set method (van Mourik, van der Eerden, et al., 2019; Waehnert et al., 2014). The distance of each ROI voxel to 5 cortical

layer depths (white matter, CSF, and the 3 grey matter depth bins) was calculated, once again using the level method. From this, the distribution of each voxel's volume over the 5 depths was calculated. For example, 80% of a voxel's volume might be in the superficial layer and 20% in the middle layer.

In multivariate analyses it is necessary to maintain the multivariate activity patterns, so this estimation of each voxel's volume across the 3 cortical layer depths was used to bin each voxel into the superficial, middle or deep layer masks. Whereas, in univariate analyses, the signal in each layer can be averaged. This was done using the spatial-GLM method, described next.

2.9.2. Univariate laminar and subfield analyses

Univariate analysis of cortical layer and hippocampal subfield ROIs is somewhat similar to mass-univariate analysis – it involves estimating and fitting a GLM to the data and then conducting statistical tests to compare condition effects. However, a major difference between the ROI analyses and whole-brain mass univariate analysis is that the BOLD signal is combined across voxels in the ROI mask prior to estimating the GLM. In Section 2.9.2.1. I explain two methods for combining BOLD signal across voxels, namely the spatial-GLM approach and simply averaging the signal, which were used for cortical layer and hippocampal subfield analyses, respectively. In Section 2.9.2.2, I explain the single-subject statistical analysis and, in Section 2.9.2.3, I explain the group-level analysis.

2.9.2.1. Preprocessing of input data

Spatial-GLM for obtaining layer time courses

As the cortical grey matter is on average 3 mm thick and our voxel size is 0.8 x 0.8 x 0.8 mm^3, some voxels will contain signal from more than one layer. This issue is known as the partial volume problem. Recall that to create the layer masks, voxels are binned into one of the layers if they at > 80% within the layer. This means that, within each layer, there may be up to 20% signal contamination from other layers. For this reason, the layer masks are not used to combine the signal from each layer.

Instead, a spatial-GLM method is used to extract signal from each cortical layer (van Mourik, van der Eerden, et al., 2019). This method is similar to the GLM framework (described in Section 2.8.2.1) for extracting activations relating to experimental conditions. For each cortical area ROI, the GLM is applied to the multivoxel pattern of BOLD signal intensities for each time

point (i.e., functional image) within the time course. It is used to estimate the BOLD signal that came from each cortical layer depth.

$Y = X\beta + \varepsilon$

In this equation, Y is the multivoxel pattern of BOLD signal intensities. X is the spatial design matrix, where each regressor column is a cortical layer depth (white matter, deep, middle, superficial and CSF) and each row corresponds to each voxel in the ROI. The data points in the design matrix represent the distribution of each voxel's volume over the cortical layer depths. β is a vector of BOLD signal estimations for each cortical layer depth. Finally, ε is the estimated residuals. Once the GLM has been applied to all time points, the result is a BOLD signal time course for each cortical layer.

Averaging BOLD signal across voxels for obtaining hippocampal subfield time courses

In the case of hippocampal subfields, voxels are manually classified as belonging to one of the 6 subfields (see Section 2.9.2.2). As with the cortical layers, there are likely to be some partial volume effects. However, the proportion of each voxels volume that belongs to each subfield is difficult to estimate. Therefore, for each subfield ROI, the signal is simply averaged across all voxels to obtain a BOLD signal time course for each subfield.

2.9.2.2. Single-participant analysis

A temporal GLM is then applied to the layer and subfield time courses using the same model described in Section 2.7. However, this time, instead of applying the GLM and statistical tests per voxel, they were applied per layer or per subfield.

2.9.2.3. Group-level analysis

At the group level, condition contrasts can be compared with zero using one-sample t tests, to indicate whether there was a difference between the conditions in the ROI. Furthermore, differences between condition contrasts can be assessed using two-sample t tests. This may involve making multiple comparisons between several different condition contrasts and so the resulting p-values must be corrected for multiple comparisons using Bonferroni correction.

2.9.3. Multivariate fMRI analyses

MVPA is a technique for fMRI data analysis that focuses on patterns of brain activity rather than individual voxel activations or individual ROI activations. Specifically, unlike the standard

mass-univariate analysis, which examines each voxel independently, MVPA considers the joint activity of multiple voxels to decode information about experimental conditions. This allows for detection of groups of voxels that represent very specific and subtle experimental conditions, such as specific memories (e.g., Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici et al., 2013; Bonnici & Maguire, 2018; Chadwick et al., 2010). There are several reasons for the high sensitivity afforded by MVPA. First, MVPA can combine weak but consistent signals from multiple voxels (Haynes & Rees, 2006; Kamitani & Tong, 2005). A distributed network of very weakly activating voxels that show no univariate effect may be detectable by MVPA. Second, it is sensitive to multivariate effects. For instance, the relationship between the activity of voxel A and voxel B might be critical for a particular cognitive process, even if changes in A or B alone are not significant. Finally, it can make use of advanced, flexible modelling methods. For example, support vector machine (SVM) algorithms, which are commonly used in MVPA, can capture non-linear relationships between voxels, providing a more nuanced understanding of brain activity compared to linear univariate methods.

2.9.3.1. Preprocessing of MVPA input data

In MVPA, information about what the participant was doing or experiencing on each trial is decoded from the multivoxel activation patterns. To make this possible, I first needed to extract the multivoxel activation pattern associated with each individual trial. A GLM approach was adopted, where a GLM was constructed for each trial. The first regressor contained the onset and duration of the trial and the second regressor contained the onsets and durations of all other trials. This allowed me to infer the activation that was unique to the trial of interest. Head motion parameters and drift terms were also included as additional regressors.

To limit the contribution of noisy voxels to the multivariate activation patterns, multivariate noise normalisation was used. In multivariate noise normalisation, the beta estimates for each voxel were divided by the square root of the GLM residuals, thereby down-weighting noisier voxels (Walther et al., 2016).

2.9.3.5. Decoding states, categories and stimuli

Different types of information can be decoded from the multivoxel activation patterns of the individual trials, depending on the research questions of the experiment (reviewed in

Chadwick, 2012). The most general type of information that MVPA can decode is cognitive states or processes, such as memory retrieval, attention, or emotion (e.g., Rissman et al., 2010). This type of information is also easily detectable using univariate approaches. MVPA is also effective at decoding categories of stimuli, such as distinguishing between visually presented scenes or faces (e.g., Bonnici, Kumaran, et al., 2012; Read et al., 2024). For many categorical kinds of information, univariate analyses are also effective. For example, univariate analyses have identified the parahippocampal place area as being involved in processing of scenes and the fusiform face area as being involved in processing of faces (Epstein & Kanwisher, 1998; Kanwisher et al., 1997). The most specific type of information that MVPA can decode are individual stimuli (e.g., a specific face) or specific internal representations (e.g., a specific memory; e.g., Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici et al., 2013; Bonnici & Maguire, 2018; Chadwick et al., 2010). This is where univariate analyses fall short, because they generally cannot distinguish between two individual memories based on overall activation differences within a brain area. MVPA, however, is sensitive to differences in the multivoxel patterns of activation associated with different internal representations.

Experiment 3 (Chapter 5) involved decoding of individual autobiographical memories. Whereas, Experiment 4 (Chapter 6) involved decoding of stimulus categories, in this case distinguishing between visually presented scenes or abstract patterns.

2.9.3.2. Representational similarity analysis

Representational Similarity Analysis (RSA) is a method for MVPA, first introduced by Kriegeskorte, Mur, & Bandettini (2008) and Kriegeskorte, Mur, Ruff, et al. (2008), that examines the similarity between patterns of voxel activation elicited by different trials. For each pair of trials, the multivariate distance (or similarity) between their activation patterns is measured. This can be done using simple measures like Pearson correlation, Euclidean distance, or other multivariate distance metrics. High similarity (high correlation or low distance) indicates that the multivoxel patterns for the two trials are similar and may therefore involve the same set of neurons. Whereas, low similarity indicates the multivoxel patterns for the two stimuli are different and may involve distinct sets of neurons. The similarity measures can be used to assess which types of information can be decoded from the multivoxel patterns. For example, if there is high similarity amongst trials where the same stimulus was presented and low

similarity amongst trials of different stimuli, then this indicates that the specific stimulus is decodable from the multivoxel patterns.

An advantage of RSA over other methods is its simplicity, making it easy to implement and interpret (Chadwick, 2012). It is also very adaptable to different experimental paradigms. Because it involves pairwise comparisons between stimuli, it does not require multiple trials for each stimulus being decoded (Chadwick, 2012).

2.9.3.3. Classification

Classification is another widely used approach for MVPA of fMRI data (e.g., Bonnici, Chadwick, Kumaran, et al., 2012; Bonnici et al., 2013; Bonnici & Maguire, 2018). This method involves training a machine learning algorithm to distinguish between different stimulus categories based on patterns of brain activity. If the classifier is able to do this, then we can infer that the categories are decodable from the multivoxel activation patterns.

The dataset is divided into two subsets: a training set and a test set. The training set is used to train the classifier, while the test set is used to evaluate its performance. During training, the classifier algorithm learns to optimise the separation between two categories (e.g., scenes versus abstract patterns) by determining an optimal "decision boundary" in high-dimensional voxel pattern space, referred to as feature space. One of the most commonly used classifier algorithms in fMRI analysis is the SVM. The SVM finds a hyperplane that best separates the data points (voxel patterns) of different conditions in the feature space. The goal is to maximise the margin, which is the distance between the hyperplane and the nearest data points from each task condition (support vectors). After training, the classifier is tested on the independent test dataset. The classifier predicts the condition of each test trial by determining which side of the decision boundary the trial falls on. The performance of the classifier is assessed by calculating the proportion of correctly classified trials in the test dataset. The accuracy is given as a percentage, indicating how well the classifier can generalise to new, unseen, data.

The train-test approach is crucial for avoiding overfitting, a common issue with complex multivariate algorithms (Duda et al., 2001). Overfitting occurs when a model learns the noise in the training data instead of the underlying pattern. This leads to inflated performance

estimates of the training set, and is the reason why a separate test set is used to estimate the model's performance.

2.9.3.4. Group level analyses

At the group level, the decoding metric (e.g., correlation or classification accuracy) across participants can be evaluated, using one-sample t tests, to see if they are different from the value expected by chance. Furthermore, condition differences can be tested using two-sample t tests, with Bonferroni correction.

2.9.4. Informational connectivity analysis

Informational connectivity analysis in fMRI is a method used to examine the relationships between different brain regions based on the informational content of their multivoxel activity patterns (Coutanche & Thompson-Schill, 2013). This approach is performed in conjunction with MVPA, and provides insights into how information is shared and integrated across different parts of the brain. Unlike traditional functional connectivity analyses, which rely on correlations of signal intensity over time, informational connectivity focuses on correlations in decoding power (e.g., RSA correlation or SVM classification accuracy) across time. If two brain areas have similar time courses of decoding for specific task-related information, they are considered to be informationally connected to one another. In other words, they may be communicating this task-related information between each other.

3. Experiment 1: Development of 7T MRI data acquisition sequences and analysis methods

3.1. Introduction

As discussed in Chapter 2, the high (7T) magnetic field strength used in laminar and hippocampus subfield fMRI studies presents numerous challenges. Geometric distortion and signal dropout scale with the size of the magnetic field, and signal artefacts such as the superficial layer bias effect are common. These issues can limit the reliability of laminar and hippocampus subfield fMRI analyses.

In this chapter, I present the first experiment, in two parts, conducted during my PhD. In the first part (Experiment 1A), I collaborated with the Centre's Physics Group to develop a 7T GE-BOLD fMRI sequence that would minimise geometric distortion, signal dropout, and signal artefacts. In the second part of the experiment (Experiment 1B), I tested data preprocessing and analysis methods in an attempt to address any remaining geometric distortion and correct for the superficial layer bias effect.

The preprocessing results showed that using a 2-fold segmented acquisition, fMRI data unwarping, and recursive BBR significantly improved the geometric distortion and the cortical surface alignment in the mPFC. However, some signal dropout and distortion persisted in areas such as the orbitofrontal cortex, with extreme levels of distortion and drop out, necessitating the exclusion of highly distorted and low signal areas using Jacobian maps.

In relation to data analysis, I found that MVPA was resistant to the superficial layer bias effect when decoding more fine-grained information such as stimulus categories and individual stimuli, but not for broader information like cognitive states or processes.

3.2. Background

Geometric distortion arises when B0 field inhomogeneity causes dephasing of hydrogen protons, which interferes with the spatial coding of the MRI signal. This distortion manifests as stretching or compression of the image along the phase encoding axis. In addition, dephasing accelerates the decay of the MRI signal (resulting in a shorter T2*), which can lead

to a phenomenon known as signal drop out, where only a small amount of signal remains during readout. The increase in geometric distortion and signal drop out at 7T is of particular concern for my experiments because the inferior part of the mPFC, the orbitofrontal cortex, experiences significant B0 field inhomogeneity. This is because it is in close proximity to the sinuses, where high B0 field inhomogeneity arises due to the large difference in the magnetic susceptibility of bone and air.

The best way to avoid geometric distortion is to use fMRI acquisition methods that prevent it from emerging in the first place. For example, one approach is to homogenise the B0 field before image acquisition using an acquisition technique called shimming (Juchem & de Graaf, 2017). Another method is to shorten the TE (the time between excitation and readout, in which dephasing occurs) by either accelerating or segmenting the acquisition of k-space (Deshmane et al., 2012; Stirnberg & Stocker, 2021). However, these sequence adjustments come with trade-offs, as improvements in one area often lead to compromises in another. For instance, excessive image acceleration can decrease SNR and functional sensitivity because fewer data points are sampled. Due to the unpredictable outcomes of different combinations of sequence parameters, it is essential to conduct pilot experiments to find the optimal balance necessary for any one fMRI study.

Another way to address the issue of high distortion is during the preprocessing of fMRI data. For example, B0 field maps, which estimate the B0 field, can be used to correct geometric distortion (Hutton et al., 2002; Jezzard & Balaban, 1995; Zeng & Constable, 2002), or recursive BBR can be used to warp anatomical segmentations, such as cortical surfaces, to match the warped functional images (van Mourik, Koopmans, et al., 2019). However, data preprocessing decisions also involve trade-offs. Excessive resampling and interpolation of fMRI data can lead to a loss of effective spatial resolution, which can prevent the dissociation of signals in microstructures like the cortical layers and hippocampal subfields.

Another challenge with 7T fMRI is susceptibility to artefacts. One of the most challenging signal artefacts that needs to be addressed in laminar GE-BOLD fMRI studies is the superficial layer bias effect, which arises from the anatomy of the cortical veins (Kay et al., 2019; Norris & Polimeni, 2019; Ugurbil, 2016). Large arteries carrying oxygenated blood descend from the pial surface into the cortical grey matter. This oxygenated blood is delivered to layer-specific neurons by small capillaries that run parallel to the pial surface. Deoxygenated blood is then

carried back to the pial surface by large ascending veins, which drain into larger veins on the pial surface (Figure 15). To isolate the layer-specific GE-BOLD response, I needed to remove the overwhelming macrovascular signal from these large veins, ensuring that only the microvascular signal from the layer-specific capillaries remained (Huber, 2020).

Understanding how the macrovasculature affects the measured BOLD signal is a crucial first step in correcting for its effect (Havlicek & Uludag, 2020; Huang et al., 2021; Kashyap et al., 2018; Markuerkiaga et al., 2016). First, the large veins on the pial surface are very close to the superficial layer, potentially contaminating the measured signal from this layer due to insufficient spatial resolution to distinguish between these veins and the superficial layer grey matter. Second, there is a higher proportion of large ascending veins in the superficial layers compared to the deeper layers, which increases the measured superficial layer BOLD signal. Third, deoxygenated blood flowing from the deep to superficial layers causes a leakage of BOLD signal from the deep to the superficial layers. Collectively, these factors reduce the spatial specificity of the measured BOLD signal, causing a bias towards the superficial layers (Huang et al., 2021).

I conducted a two-part experiment to aid the development 7T MRI acquisition and analysis methods that minimised geometric distortion, signal dropout, and signal artefacts. In the first part of the experiment (Experiment 1A), I worked with the Department's Physics Group to develop a 7T GE-BOLD fMRI sequence. This sequence aimed to minimise geometric distortion, signal dropout, and signal artefacts in the mPFC and hippocampus while maintaining good SNR and task sensitivity. In the second part of the experiment (Experiment 1B), my goal was to develop a data preprocessing and analysis pipeline that corrected for any remaining geometric distortion, and the superficial layer bias effect, whilst enabling me to achieve accurate cortical surface segmentation for laminar analyses.

3.3. Experiment 1A. 7T fMRI sequence development

3.3.1. Methods

3.3.1.1. Participants

Two pilot participants were recruited and scanned multiple times for development of the MRI sequences. They were 2 females aged 33 and 41, with no history of neurological or psychiatric conditions, and with normal or corrected-to-normal vision.

3.3.1.2. Experimental protocol

Participants attended for one 7T MRI scanning session that included fMRI data acquisition. A scanning session involved either resting in the scanner while fMRI data were acquired, or a generic autobiographical memory task was performed as followed. One week prior to scanning, the participant was provided with a list of generic memory cues such as “birthday party” or “graduation”. They were instructed to prepare some of their own personal memories that aligned with each cue. In each trial of the fMRI task, they were shown a memory cue for 3 seconds, were presented with a visual cue to close their eyes for 1 second and then recalled the memory for 10 seconds (Figure 17). If they were unable to recall a memory associated with the generic cue, they had the option to imagine an event instead, which has been shown to activate the same core network of brain areas as memory retrieval (Buckner & Carroll, 2007; Mullally & Maguire, 2014; Schacter et al., 2007). An auditory tone (1 second) then cued them to open their eyes, followed by visual presentation of a question ("How vivid was the memory?") for which they could answer on a scale from 1 to 3 (3 = "highly vivid"… 1 = “not at all vivid”). If they imagined an event then they did not answer the vividness question, but indicated, with a button press that they had “imagined instead”. Counting trials were used as a baseline condition in the fMRI analysis, whereby a 2-digit number was presented instead of a memory cue. A participant was tasked with counting up in 3's silently in their head from that number for 10 seconds, and then they rated their focus during the counting exercise.

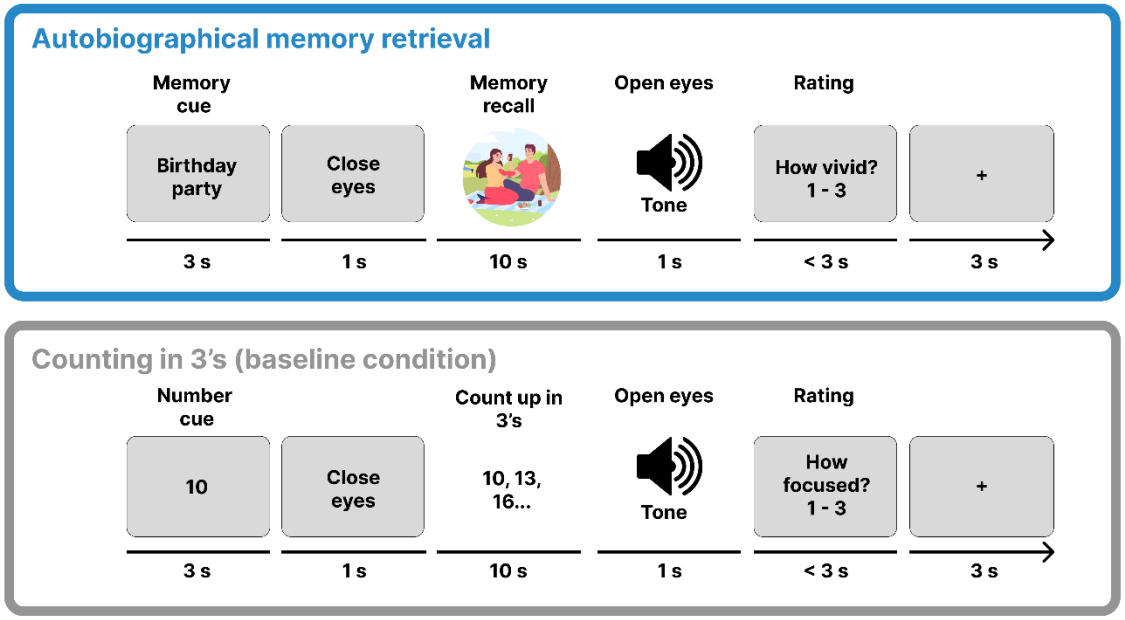

Figure 17. Generic autobiographical memory task performed during 7T fMRI. Autobiographical memory retrieval trials (blue box) consisted of a generic memory cue (3 s), then a cue to close eyes (1 s), followed by a period of vivid memory retrieval (10 s), after which there was an auditory tone to cue opening of the eyes (1 s), followed by a vividness rating (up to 3 s) and finally a fixation cross (3 s). Counting baseline trials (grey box) consisted of a number cue (3 s), followed by a cue to close eyes (1 s), then a period of counting (10 s), after which there was an auditory cue to open eyes (1 s), then a focus rating (up to 3 s) and finally a fixation cross.

3.3.1.3. MRI sequences

I acquired partial-brain, submillimeter resolution fMRI data using various different MRI sequences. All sequences that we tested used T2*-weighted 3D GE-BOLD EPI. The parameters that remained consistent across all sequences were the field of view of 192 x 192 x 70.4 mm^3 and the RF pulse sequence of binomial (1331) water-selective excitation.

Across a series of A-B tests, several sequence parameters were adjusted; these included: the direction of phase encoding, the angle of the slab, the use of a segmented acquisition, parallel imaging, partial Fourier imaging, and the shimming procedure. Many of these parameter decisions required other sequence parameters to be adjusted such that the total volume acquisition time was ~4 seconds, to achieve reasonable sensitivity in the autobiographical memory retrieval task, and the voxel size no more than 0.9 x 0.9 x 0.9 mm^3, to enable dissociation of signal in the 3 cortical layer depths.

3.3.1.4. fMRI data preprocessing and analysis

The functional images were visually inspected to understand the effect of the sequence parameter options on signal drop out and geometric distortion. For some of the parameter tests, we also quantified the effect on functional sensitivity by calculating tSNR and statistical maps (i.e., t-value maps) of task contrasts.

For each participant, the functional images were first realigned to correct for head motion during scanning. For each voxel, tSNR was calculated from the mean BOLD signal across the time course and dividing it by the standard deviation in order to compare the average tSNR value in the mPFC between participants. A coarse manual segmentation of the mPFC was performed on the mean functional image to facilitate this analysis.

To produce the t-value maps, a GLM was first fitted to the realigned fMRI data. It contained two regressors for task conditions: one for the 10 second memory retrieval/imagination period and another for the 10 second counting baseline period. Each of these regressors was convolved with the canonical HRF. The 6 motion parameters estimated during realignment were used as nuisance regressors in the model. A memory retrieval versus counting baseline contrast was applied to the beta estimates to yield a t-value map. This t-value map indicated the degree of activation during autobiographical memory retrieval of each voxel.

3.3.2. Sequence parameter tests and results

3.3.2.1. Signal drop out and geometric distortion

Phase encoding direction: In GE-EPI, phase encoding is achieved by applying a gradient pulse (sometimes referred to as a blip) along the phase encoding axis before acquisition of each line of k-space. The choice in polarity of this gradient can result in different manifestations of geometric distortion, due to differences in the magnetic field resulting from the interaction between the susceptibility field and gradient. By visually inspecting the functional images, we compared the nature of geometric distortion and signal drop out in functional images acquired with anterior-to-posterior phase encoding to images acquired with posterior-to-anterior phase encoding.

Angle of slab: The amplitudes of the x, y and z imaging gradients control the size, positioning and angle of the slab. The same susceptibility field will perturb the x, y, and z imaging gradients in different ways depending on their respective amplitudes and will therefore change the

nature of the distortion and signal drop out. The angle of the slab was altered within a range of 15° to 35°. The effect of the slab angle on geometric distortion was assessed by visually comparing the functional images.

Segmented acquisition: Typically, with 3D EPI sequences, the whole of k-space is acquired following a single excitation of the slab. With a segmented acquisition, the slab is excited n times and 1/n of k-space is acquired in each excitation, where n is the number of segments (Stirnberg & Stocker, 2021). By segmenting the acquisition, each of the n acquisitions is shorter, meaning that there is less time for the signal to decay and less time for dephasing due to B0 field inhomogeneity. Consequently, the distortion and drop out are reduced. However, using a segmented acquisition increases the total volume acquisition time, which could reduce the statistical power when detecting task-related activations as there will be fewer data points in the GLM. We compared geometric distortion and signal drop out in images acquired with (2 segments) and without segmented acquisition. We also compared memory retrieval activation maps (i.e., t-value maps of the memory retrieval versus counting contrast) of the two sequences.

Partial Fourier: Because, in theory, k-space is symmetrical, a way to reduce the time taken to acquire an image is by only sampling (a minimum of) half of k-space. This technique, known as partial Fourier imaging, decreases the number of phase encoding iterations and shortens the readout time (McGibney et al., 1993). As a result, there is less time for dephasing, leading to reduced signal dropout and distortion. However, in practice, k-space is not completely symmetrical due to the inherent phase resulting from B0 field inhomogeneity. This means that partially sampling k-space can reduce SNR of the image. We compared geometric distortion and signal drop out in an MRI sequence with 6/8 Partial Fourier, meaning 75% of k-space was sampled and the rest was filled with zeros, and a sequence without partial Fourier. To preserve the TR at close to 4 seconds, the sequence without partial Fourier also had a larger voxel size of 0.9 x 0.9 x 0.9 mm^3 as opposed to the 0.8 x 0.8 x 0.8 mm^3 voxel size used in all other sequence tests. We also estimated the impact of partial Fourier imaging on tSNR by comparing the mPFC tSNR values between sequences.

Advanced shimming: Some B0 field inhomogeneities can be corrected prior to data acquisition by applying small gradients using specially designed coils, a process known as active shimming (Juchem & de Graaf, 2017). Various algorithms can be used to estimate the

necessary corrective gradients. In our study, we compared geometric distortion and signal dropout in sequences acquired using the Siemens shimming algorithm and a bespoke advanced shimming algorithm.

The polarity of phase encoding blips noticeably affected distortion and dropout. Sequences with posterior-to-anterior phase encoding experienced similar levels of geometrical distortion in the mPFC as those with anterior-to-posterior encoding, but the nature of the distortion differed. Anterior-to-posterior encoding caused the front of the brain to appear compressed, whereas posterior-to-anterior encoding caused it to appear stretched (Figure 18a). Additionally, the posterior-to-anterior phase encoding direction resulted in increased signal dropout in the orbitofrontal cortex. Therefore, we decided to proceed with anterior-to-posterior phase encoding.

Among all the tested parameters, segmented acquisition yielded the most significant improvements in distortion and signal dropout. In the mPFC, segmentation visibly reduced the degree of deformation (Figure 18b). In addition, there was less empty space (i.e., signal dropout) in the images, specifically in the orbitofrontal cortex. Importantly, there was only a small effect on the memory retrieval activation maps, indicating that the reduced number of images produced by the segmented sequence did not limit our ability to detect task-related activity. Consequently, we decided to proceed with segmented acquisition.

The use of the advanced method for active shimming also visibly improved brain distortion (Figure 18c). Therefore, we decided to proceed with this method. The remaining sequence parameters (angle of the slab, partial Fourier imaging) had no noticeable effect on geometric distortion or signal dropout. Additionally, the impact of 6/8 partial Fourier imaging on memory retrieval activation maps and mPFC tSNR was marginal (Figure 18d; tSNR: 11.34 with partial Fourier and 11.92 without partial Fourier). Therefore, we selected these parameters based on other considerations. We chose a 15° angle to ensure mPFC and hippocampus coverage across a variety of brain shapes and sizes. We opted for partial Fourier imaging to decrease the TR.

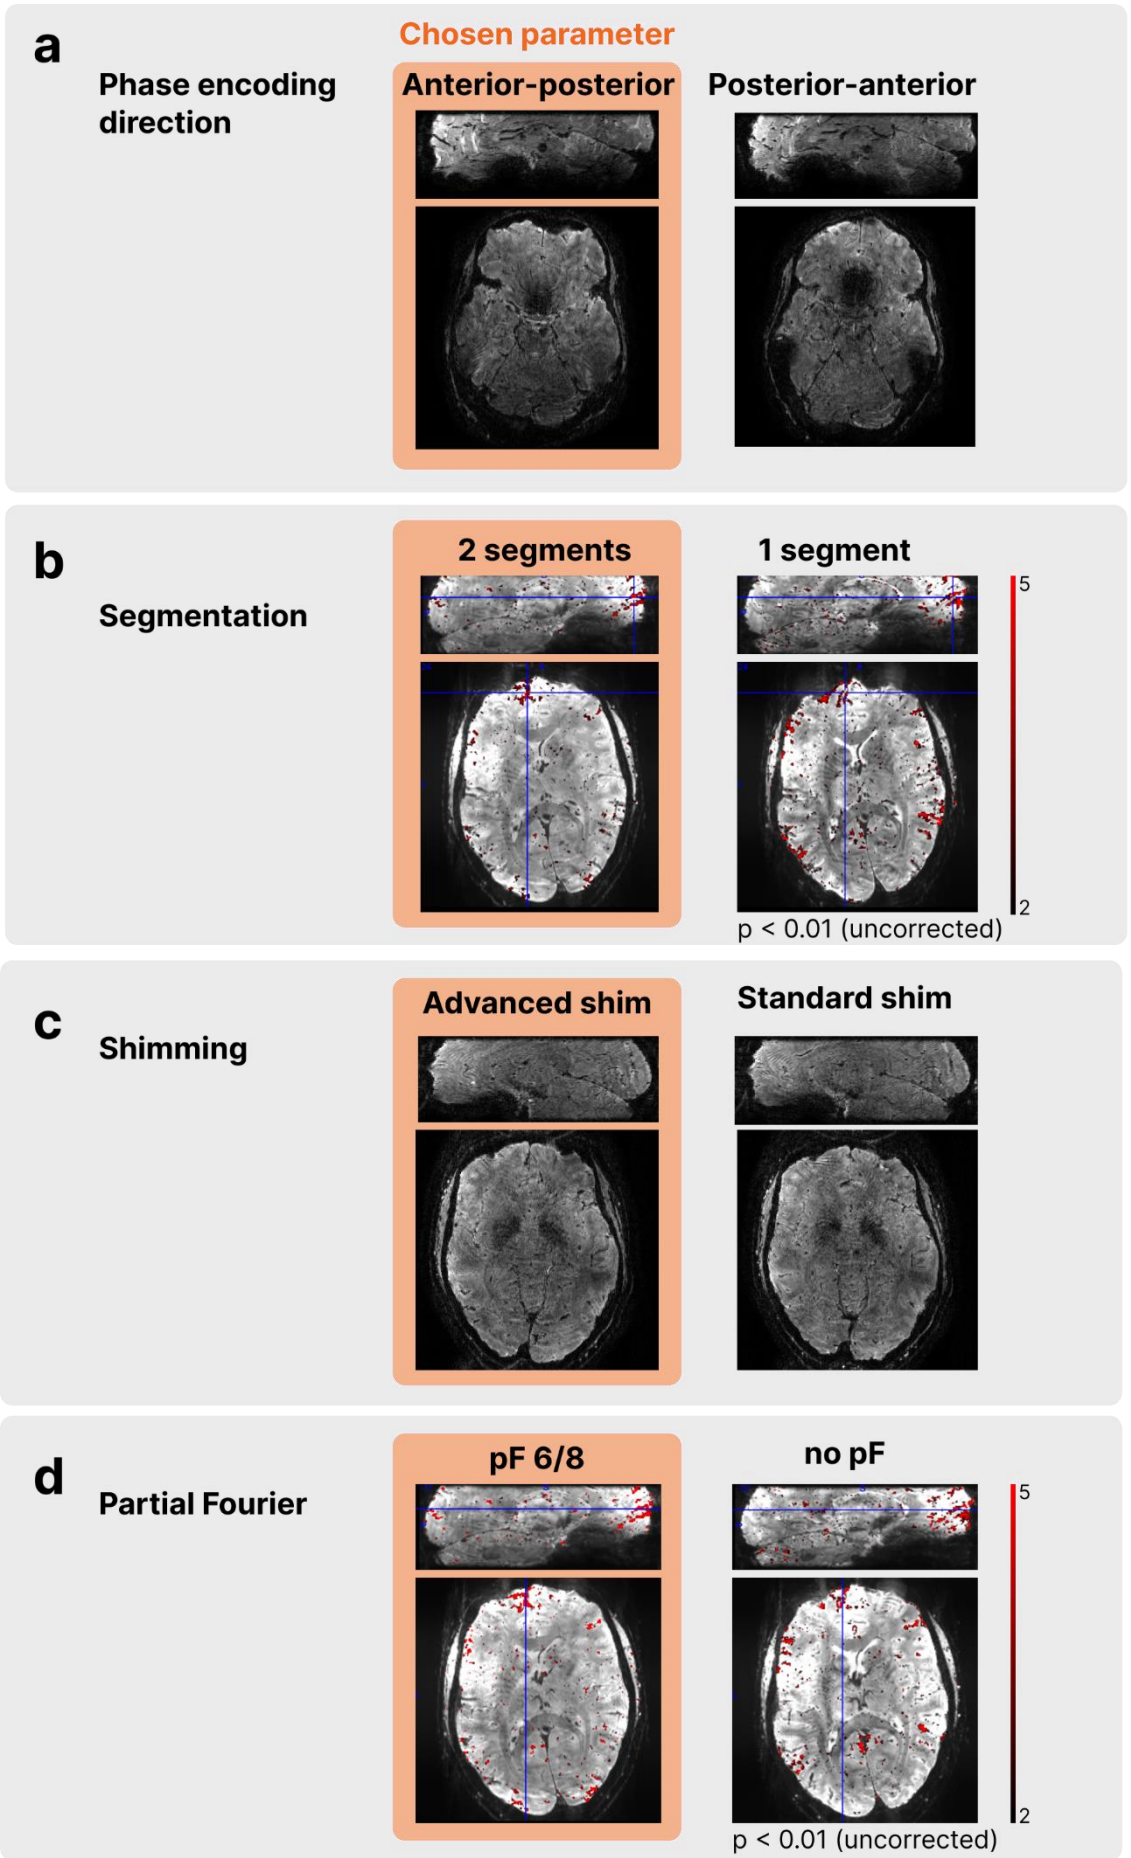
a
Phase encoding direction
Chosen parameter
Anterior-posterior
Posterior-anterior
b
Segmentation
2 segments
1 segment
5
2
p < 0.01 (uncorrected)
c
Shimming
Advanced shim
Standard shim
d
Partial Fourier
pF 6/8
no pF
5
2
p < 0.01 (uncorrected)

Figure 18. Optimising acquisition parameters. The effect of changing four acquisition parameters on signal distortion, dropout or functional sensitivity were assessed. (a) Setting the phase encoding direction to anterior-posterior resulted in less signal drop out in the mPFC than a posterior-anterior phase encoding direction. (b) A two-segment acquisition produced less distortion than a one-segment acquisition. (c) Advanced shimming marginally improved distortion and signal dropout. (d) Partial Fourier imaging did not markedly impact task-related activations.

3.3.2.2. Gibb's ringing artefact

During the above sequence tests, a "ringing" artefact was observed in the mPFC (Figure 19). We hypothesised that the artefact was caused by eye movement during acquisition of the reference data. Reference data are acquired at the start of the fMRI sequence before the task starts. It is the only full sample of k-space that is acquired during the fMRI sequence and it is used to fill in the gaps in the subsequent partial acquisitions of k-space (due to parallel imaging and partial Fourier imaging) that occur during the task. It is very important that the reference data are of good quality as this impacts all images in the time series. Subtle movements, including eye movements, can lead to motion artefacts that repeat throughout the time series. A series of tests was performed to test the hypothesis that eye movements during acquisition of the reference data were causing the ringing artefact, specifically:

Eyes fixed or moving: Two fMRI acquisitions were compared, one where the participant moved their eyes during acquisition of the reference data and a second where they fixated on a cross in the centre of the screen.

Field of View (FOV) including or excluding the eyes: Two FOVs were used, an inferior one that included the eyes and a superior one that did not.

Linear or interleaved acquisition: The reference data were previously acquired in an interleaved fashion, where k-space is acquired in two segments, the first covering even lines of k-space and the second covering odd lines (or vice versa). We suspected that interleaved acquisition may be more sensitive to "ringing" motion artefacts than a linear acquisition, where each line of k-space is acquired in order. Therefore, we tested two sequences, one with linear acquisition of reference data and the other with interleaved acquisition of the reference data.

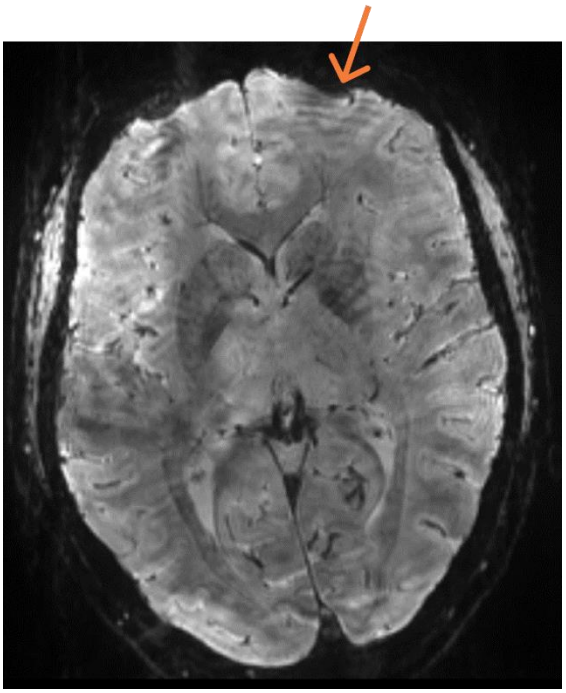

Figure 19. Gibb's ringing artefact. A Gibb's ringing artefact was observed in the mPFC during piloting. This was caused by an abrupt reduction in signal at the edge of k-space due to distortion of k-space.

None of the above tests impacted the artefact, leading us to reject our hypothesis that eye movements were the cause. Irrespective of this, we decided to use a fixation cross during acquisition of the reference data to avoid potential motion artefacts.

We pursued investigation of this artefact further. High intensity signal typically resides in the centre, as opposed to the edges, of k-space. On inspecting the raw k-space data, high intensity signal was observed around the edges of k-space, leading to an abrupt change in signal at the outermost perimeter of k-space where the signal becomes zero. We hypothesised that this abrupt change in signal was what was causing the artefact in the image, which is also known as Gibb's ringing or a truncation artefact (Gibbs, 1898). In a second pilot we altered some sequence parameters in an attempt to improve the quality of data in k-space.

Reduced acceleration: To reduce the volume TR, sampling of some of the data points in k-space can be skipped and later recovered by interpolating. This is referred to as acceleration. In previous pilots we used an acceleration in both the direction of the partition (factor = 2) and the direction of phase encoding (factor = 4). However, the mispositioning of signal in k-space may be compounded by this interpolation. Therefore, we removed the acceleration in

the direction of the partition. To maintain a TR of less than 4 seconds we also reduced the FOV substantially.

Altered the orientation of the slab: The susceptibility field impacts the imaging gradients differently depending on their relative magnitudes, meaning that different imaging gradient magnitudes can result in different mis-mappings of signal into k-space. As the imaging gradient magnitudes determine the angle of the slab, changing the angle could shift the high intensity signal around the edges of k-space further into the centre.

Apodisation during image reconstruction: Ideally, we aimed to avoid the displacement of high-intensity signal altogether, either by reducing acceleration or changing the orientation of the slab. However, as a backup plan, we also explored a reconstruction method to prevent the displaced signal from causing artefacts in the images. Apodisation, a technique for filtering the signal at the edges of k-space so that it gradually decreases to zero, can be applied after the acquisition of k-space during the image reconstruction process.

Neither reducing acceleration nor changing the angle of the slab impacted the artefact. Apodisation successfully removed the ringing artefact. This confirmed that the artefact was indeed a Gibb's ringing artefact caused by abrupt changes in signal intensity. A downside of apodisation is that it would remove some of the spatial information that is typically located in the edges of k-space. Furthermore, the high intensity signal that may reflect neural activations is lost, thereby reducing SNR. However, it was more important to us that we could remove the artefact as it would interfere with the coregistration of cortical surfaces. Therefore, we proceeded with an adjusted reconstruction of the image that included apodisation.

3.4. Experiment 1B. 7T MRI preprocessing and analysis pipeline development

3.4.1. Methods

3.4.1.1. Participants

MRI data from 5 participants that were scanned during Experiment 3 (Chapter 5) were used for developing the preprocessing and analysis pipeline. They were 3 females and 2 males aged between 18 and 35 years old (mean = 25.60, SD = 3.44), with no history of neurological or psychiatric conditions, and with normal or corrected-to-normal vision.

3.4.1.2. Experimental protocol

Participants attended the Department for two visits. During the first visit they underwent structural MRI scanning followed by an Autobiographical Interview (Levine et al., 2002), in which they described 12 autobiographical memories, 6 memories from the past month (recent memories) and 6 memories from 2-5 years ago (remote memories). They then decided on a short title for each memory.

They returned for a second visit 1-2 weeks later. The autobiographical memory retrieval task was very similar to the general autobiographical memory retrieval task described in Section 3.3.1.2, except instead of using generic memory cues, the participant-specific memory titles were used. The timings of the memory cue, close eyes cue, memory recall period, open eyes cue and vividness ratings was the same as for the general autobiographical memory retrieval task. This time, because the memory cues related to specific memories chosen by the participant, they did not have the option to imagine an event instead. Therefore, the vividness question was answered on a scale from 1 to 4 (4 = "highly vivid"… 1 = "not at all vivid") and there was no “imagined instead” option. The counting trials were also exactly the same as in the generic autobiographical memory task. There was a total of 6 runs of the task. Each memory cue and number cue was presented once per run in a random order (both within and across runs), yielding 24 trials per run. For more information on this task, see Chapter 5, Section 5.3.2.2.

3.4.1.3. MRI sequences

MRI data were acquired on a Siemens Magnetom Terra 7T MRI system (see Chapter 2, Section 2.6.1).

Functional MRI: Partial brain, submillimetre resolution (voxel size = 0.8 × 0.8 × 0.8 mm^3) fMRI images were acquired using a T2*-weighted 3D GE EPI sequence. The sequence parameters used were chosen based on the results of Experiment 1A. They were as follows: volume acquisition time = 3872 ms, TR = 44 ms, TE = 18.7 ms, flip angle = 14°, field of view = 192 × 192 × 70.4 mm^3, slab angle = 15°, direction of phase encoding = anterior-to-posterior, parallel imaging acceleration in both the direction of the partition (factor = 2) and the direction of phase encoding (factor = 4), partial Fourier 6/8 in the phase-encoded direction of the EPI readout, and a binomial (1331) water-selective excitation. Four images, with the same parameters but the opposite (posterior-to-anterior) phase encoding direction, were acquired at the start of the sequence. Advanced B0 shimming (WIP 1441) was used.

Anatomical MRI: Four anatomical images were acquired. One whole brain MT-weighted EPI image was acquired per subject using a T2*-weighted 3D GE EPI sequence with many of the same parameters used in the fMRI sequence. Some parameters were modified: MT-based contrast = on, volume acquisition time = 3 m 45 s, TE = 16.97 ms, flip angle = 8°, field of view = 192 × 192 × 128 mm^3, parallel imaging (factor = 4). MT-weighted EPI images were matched to the fMRI images in terms of geometrical distortion due to use of the same parameters (e.g., echo spacing and acceleration factor) but they had better contrast between grey and white matter due to the MT weighting. Therefore, they were useful for coregistering cortical surfaces to the fMRI data.

Submillimetre resolution (voxel size = 0.65 × 0.65 × 0.65 mm^3) T1-weighted images were acquired using a single acquisition of an MP2RAGE sequence (Marques et al., 2010) with a volume acquisition time = 8 m 42 s, TR = 5000 ms, TE = 2.54 ms, inversion time = 900 ms and 2,750 ms, 5° and 3° flip angles, field of view 208 × 208 × 156 mm^3, and in-plane GRAPPA acceleration (factor = 3). The MP2RAGE sequence produces 4 images: (1) the first read out with 900 ms inversion time (i.e., the time between the inversion pulse and the readout), (2) the second readout with 2,750 ms inversion time, (3) a T1-weighted image produced by

combining the two readouts, (4) an inverted T1-weighted image produced by combining the two readouts.

Two partial brain, submillimetre resolution (voxel size = 0.52 × 0.52 × 0.5 mm^3) T2-weighted images were acquired using a T2-weighted 3D SPACE sequence (Mugler et al., 2000) with a volume acquisition time = 10 m 25 s, TR = 3,500 ms, TE = 229 ms, flip angle = 8°, field of view = 200 × 169 mm x 56 mm, partition oversampling = 14.3%, and in-plane GRAPPA acceleration (factor = 2).

3.4.2. Preprocessing and analysis tests and results

Various preprocessing and analysis methods were tested with the ultimate goal of obtaining layer-specific univariate and multivariate results. The process involved several sub-goals. First, I optimised the definition of the cortical surfaces (the cortical grey matter boundaries) in the anatomical images, which is crucial for accurately defining the cortical layers. Second, I tested various methods for removing residual geometric distortion in the fMRI images. Third, I explored methods for defining the cortical layer ROIs to ensure a reasonably even distribution of voxels across the cortical layers. Finally, I tested methods for correcting the superficial layer bias effect, including further preprocessing of the ROIs and evaluating several options for calculating condition-specific responses. Success was measured differently for each sub-goal, as explained in the relevant sections.

3.4.2.1. Cortical surface segmentation

Accurate segmentation of cortical surfaces is crucial for achieving precise coregistration of the cortical surfaces to the functional images, which in turn ensures the accurate delineation of cortical layer depths. In laminar fMRI studies, FreeSurfer's recon-all pipeline is typically used for defining the pial and white matter cortical surfaces on the anatomical image(s) (e.g., Kok et al., 2016; Lawrence et al., 2019). This pipeline requires a whole brain T1-weighted image as an input. It performs several preprocessing steps on the image, including normalisation of signal intensities and removal of voxels corresponding to the skull. It then segments the brain into various tissue types (e.g., white matter, grey matter and CSF) and constructs two cortical surfaces (made up of vertices, edges, and faces) corresponding to the outer (pial surface) and inner (white matter surface) boundaries of the grey matter. Finally, it divides the cortex into different anatomical regions based on predefined atlases.

I tested 4 different combinations of input images to the FreeSurfer recon-all pipeline.

T1-weighted MP2RAGE image: FreeSurfer requires a T1-weighted image for extracting the cortical surfaces. Therefore, I used the raw T1-weighted MP2RAGE image as an input.

T1-weighted MP2RAGE and inverted T1-weighted MP2RAGE image: FreeSurfer also accepts an optional whole brain T2-weighted image as a second input image for improving definition of the pial surface. The MP2RAGE sequence produces an inverted T1-weighted image, which has inverted T1 contrast giving it the same appearance as a T2-weighted image. Although it is not technically a T2-weighted image, I tested whether it could be used in place of one in the FreeSurfer recon-all pipeline for improvement of the pial surface.

T1-weighted MP2RAGE and inverted T1-weighted MP2RAGE image (denoised and smoothed): The two images were minimally smoothed (FWHM = 1.5 mm) and denoised using a Structural Averaging Toolbox before being input into FreeSurfer recon-all.

T1-weighted MP2RAGE and a "spliced" whole brain T2-weighted image: To further improve definition of the pial surface, I wanted to take advantage of the high resolution (0.56 mm isotropic) T2-weighted partial brain images. It was not possible to input these images directly into FreeSurfer recon-all as the pipeline does not accept partial brain volumes as inputs. Therefore, I created a "spliced" whole brain T2 image, which was produced by combining the whole brain MP2RAGE inverted T1-weighted image with the partial brain T2-weighted images. To create this image, the whole brain inverted T1-weighted MP2RAGE image was coregistered (using NMI) with the partial brain high resolution T2-weighted images and the aligned images were then averaged to produce a "spliced" whole brain T2-weighted image. The resulting image had very high SNR within the central part of the volume that was an average of all 3 images. Whereas the most superior and inferior parts of the volume, which were comprised of only the whole-brain T1-inverted image, had lower SNR. Both the "spliced" T2-weighted image and the MP2RAGE T1-weighted image were smoothed and denoised (as described above) before being inputted to FreeSurfer recon-all.

The impact of each method on the quality of the cortical surface definition was assessed by visually inspecting the cortical surface segmentations. Using the raw T1-weighted MP2RAGE image resulted in areas of the pial and white matter surfaces overfitting the noise in the image, causing jagged edges in the surfaces (Figure 20a). Using the inverted T1-weighted MP2RAGE

image as an additional input improved definition of the surfaces, smoothing out some of the jagged edges (Figure 20b). Smoothing and denoising the input images further improved the smoothness of both of the surfaces (Figure 20c). Finally, use of the T2 "spliced" image further improved pial surface registration, whilst still meeting the FreeSurfer requirements of a whole brain scan (Figure 20d). The pial surfaces extracted on the "spliced" image were of good quality and so we proceeded with this method (orange box in Figure 20d).

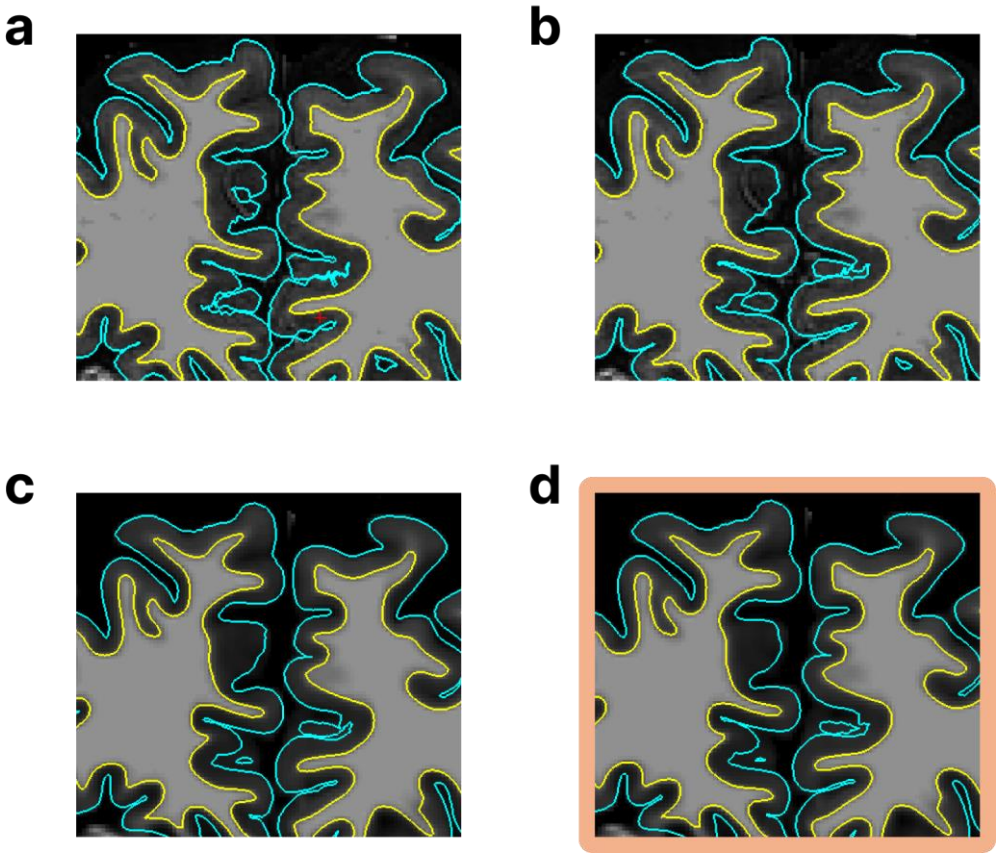

Figure 20. Optimising the definition of cortical surfaces. (a) Pial (turquoise) surfaces extracted on the T1-weighted MP2RAGE image were jagged and inaccurate. (b) Using the T2-weighted MP2RAGE image as an additional input improved definition of the pial surface. (c) Denoising and smoothing the T1- and T2-weighted images further smoothed and improved the pial surface. (d) A "spliced" whole brain T2 image, produced from the MP2RAGE T2-weighted whole brain image and the T2-weighted partial brain images, was used to take advantage of the high-resolution partial brain volumes.

3.4.2.2. Distortion correction

There are several methods for dealing with geometrically distorted functional images during data preprocessing. Perhaps the most commonly used is unwarping using a B0 field map. Here

a map of the B0 field can be used to estimate the displacement of voxels that would have been caused by field inhomogeneities. The B0 field map can either be directly acquired (Hutton et al., 2002; Jezzard & Balaban, 1995; Zeng & Constable, 2002) or estimated from images acquired with the reverse phase encoding (Andersson et al., 2003; Hedouin et al., 2017; Holland et al., 2010; Morgan et al., 2004). In Malekian at al. (2023), we found that the reverse phase encoding method produced superior distortion correction as it resulted in unwarped functional images that more closely resembled anatomical images.

Another, less common, approach to dealing with geometric distortion is using recursive BBR (van Mourik, Koopmans, et al., 2019). This method does not perform any unwarping of the functional images. Instead, the functional images remain distorted and the cortical surfaces are warped to fit them, so that they are distorted in the same way that the functional images are distorted. The logic behind this approach is that the defined cortical layers will also be distorted meaning that the true BOLD signal can be retrieved for each layer.

I tested two different preprocessing pipelines for distortion correction. In both approaches, the data was first realigned to correct for head movement during the course of fMRI scanning.

Recursive BBR only: In the first preprocessing pipeline, SPM Realign was used to realign the functional images and no further preprocessing of the functional data was performed. Instead, recursive BBR was relied upon to deal with any distortion in the functional images. The MT-weighted functional image was used for this coregistration instead of a functional image. This is because the MT-weighted image has superior contrast to the functional images, while still having equivalent distortion and signal drop. The MT-weighted image was coregistered to the mean functional image. Then, standard NMI coregistration and BBR were performed to bring the MP2RAGE T1-weighted image and the cortical surfaces into alignment with the MT-weighted image. Finally, recursive BBR was performed. In recursive BBR, BBR is applied 6 times to increasingly smaller partitions of the cortical surface mesh (van Mourik, Koopmans, et al., 2019). Importantly, in each iteration, scaling was allowed along the phase-encoding direction, because in GE-BOLD EPI, geometrical distortion is most severe in the phase-encoded direction (Jezzard & Clare, 1999; Mansfield, 1977). It is the combination of scaling in the phase encoding direction and coregistering the surfaces to subsections of brain that results in the non-linear transformation (i.e., warping) of the cortical surfaces into alignment with the functional data. If recursive BBR could successfully distort the cortical

surfaces to fit the functional images, then unwarping of functional images would be unnecessary. It was important to test whether this is possible because unnecessary preprocessing of the functional images should always be avoided. This is because, each time the voxel data are interpolated, there is some smoothing of the data which results in a reduction in the effective resolution and potentially blurring of the BOLD signal across the cortical layers and hippocampal subfields.

Unwarping and recursive BBR: In the second preprocessing pipeline, FSL topup was used to estimate a B0 field map from the 4 posterior-anterior phase encoded functional images and 4 anterior-posterior phase encoded functional images (Andersson et al., 2003; Smith et al., 2004). The SPM Field Map toolbox was then used to calculate a voxel displacement map from the estimated B0 field map, which was subsequently used to unwarp the functional image using SPM Realign and Unwarp. SPM Realign and Unwarp estimates both rigid body realignment and unwarping, while considering susceptibility-motion interactions by estimating changes in the B0 field with participant movement (Andersson et al., 2001). Importantly, the functional images were only interpolated and resampled once, in an attempt to limit smoothing of the data. Recursive BBR was then used to warp the cortical surfaces to fit the unwarped functional images in order to correct for any remaining distortion in the functional images. The MT-weighted image was first unwarped so that its distortion was still equivalent to the functional images. Then, the coregistration steps were performed as described above.

I assessed the degree of distortion correction achieved by each method, by viewing the coregistered cortical surfaces on the preprocessed MT-weighted images. I also calculated tSNR in the mPFC in each of the processing methods by dividing the mean BOLD signal by the standard deviation of the BOLD signal and averaging across voxels in the mPFC (using a manually defined mPFC mask).

The first method, recursive BBR only, was not successful in aligning the cortical surfaces to the functional images, particularly in parts of the brain that were highly distorted such as the mPFC (Figure 21, left column). Whereas, unwarping brought the functional images into closer alignment with the true anatomy, defined by the raw cortical surfaces (Figure 21, top right quadrant), which resulted in good alignment of the recursive BBR cortical surfaces (Figure 21, bottom right quadrant). Both methods have comparable tSNR (no unwarping: mean = 14.74,

SD = 2.58, unwarping: mean = 14.61, SD = 2.54). Therefore, I decided that unwarping was a necessary preprocessing step for laminar analyses in the mPFC.

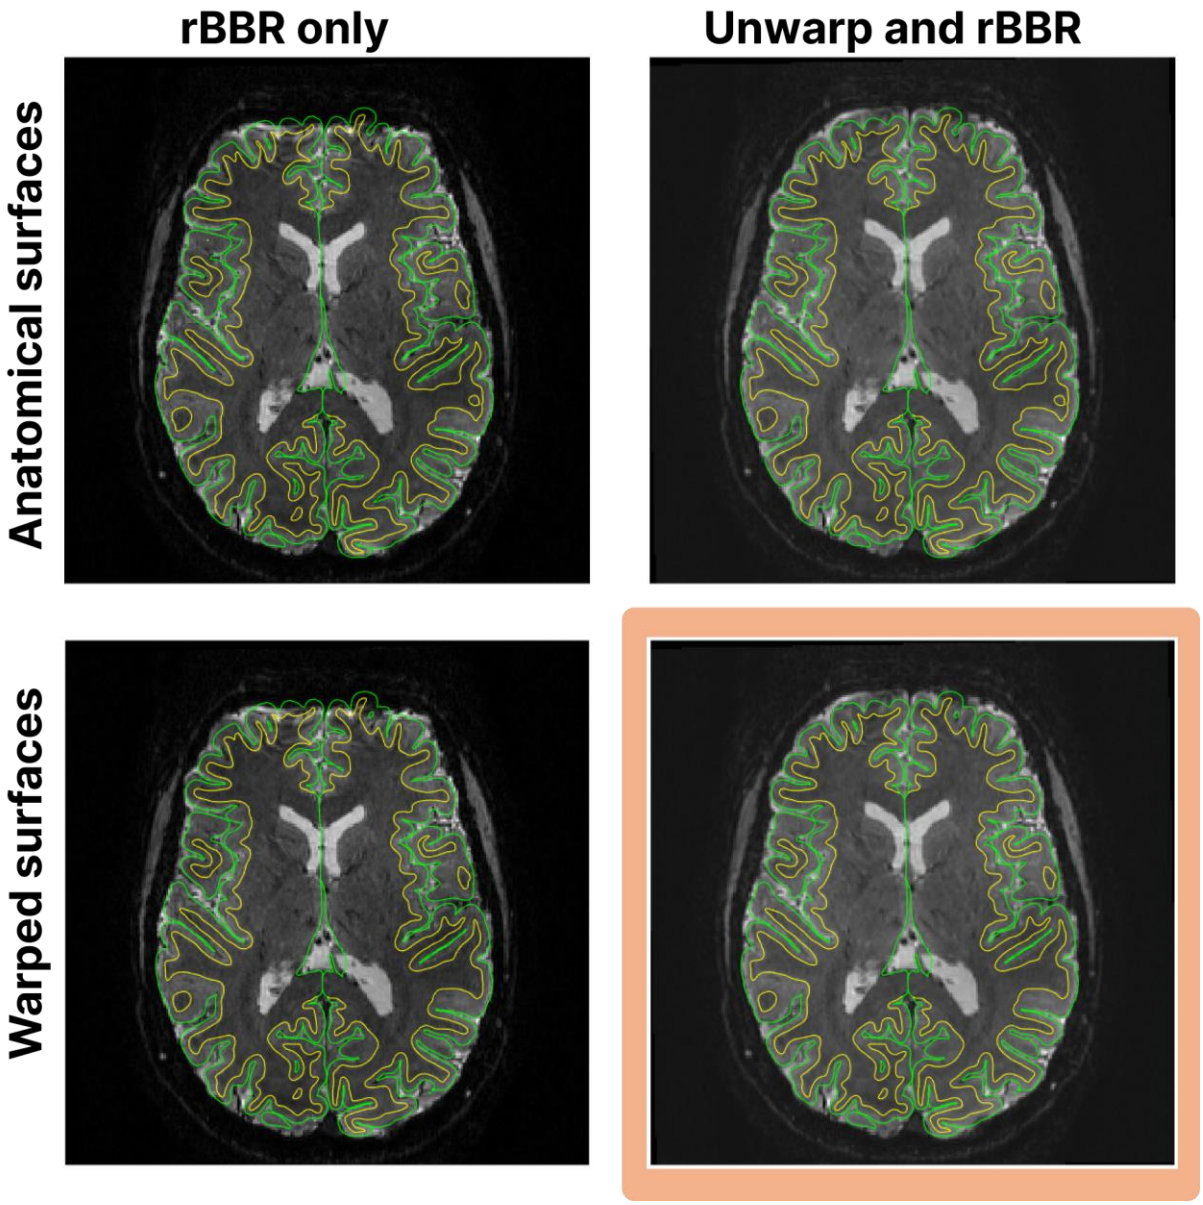

Figure 21. Optimising the distortion correction of functional images. The preprocessed MT-weighted images are overlaid onto the anatomically-faithful cortical surfaces (top row) and warped cortical surfaces (bottom row). Realigned and unwarped images (second column) are considerably closer to the true anatomy than images that were only realigned (first column), which in turn results in improved alignment of the warped cortical surfaces.

Jacobian mask: Upon visually inspecting the unwarped functional images, I noticed that there were areas of the brain, particularly around the orbitofrontal cortex, that appeared spatially

smoothed (i.e., blurred) or contained low or no signal due to signal drop out. As a final precaution against distortion and signal drop out, I decided to try a method for removing these brain areas from analysis. Using the SPM HySco toolbox, I produced a Jacobian map, a map of the derivative of the magnetic field that represents the changes in the B0 field across the image. By thresholding this map, areas with significant B0 field inhomogeneities (i.e., high levels of signal drop out and distortion) were masked out of the ROIs used for analyses. The masks were overlaid onto the functional images and visually inspected.

The Jacobian mask successfully excluded voxels that were in brain areas with significant spatial blurring and lack of signal. Therefore, I decided to proceed with this masking method.

3.4.2.3. Correction of superficial layer bias

I tested a total of 7 different methods for correcting the superficial layer bias effect. I focused on the mPFC, given it was my main cortical area of interest.

Before I could test these methods, I needed to produce mPFC ROIs. First, the FreeSurfer recon-all pipeline was used to segment the brain according to several brain atlases. These brain area masks were coregistered to the fMRI data by applying the same transformation estimated by BBR of the cortical surfaces to the MT-weighted image. The 4 anatomical segmentations in the region of the mPFC (the medial orbitofrontal, frontal pole, rostral anterior cingulate and superior frontal cortices) were then translated from surface space into volume space by projecting them across the surface normal, which is a 3D vector that runs perpendicular to the surface. Gaps in the ROIs were filled in using FSL Maths Edge, Dilate and Erode operations. The 4 ROIs were then combined to make an mPFC ROI.

Three methods that were used for correcting the superficial layer bias involved attempting to exclude voxels that contained large veins from the laminar analyses. They were as follows.

Removing voxels with low tSNR: Voxels in close proximity to large veins tend to have a low tSNR as they experience large variations in signal intensity due to large fluctuations in blood oxygen level (Jia et al., 2021). A tSNR map was computed from the fMRI data and voxels with the lowest 35% of tSNR values were excluded.

Removing voxels with low T2 signal intensity: Voxels near large veins have low T2 signal intensity due to the short T2 decay time of deoxygenated haemoglobin. The spliced high

resolution T2-weighted image was coregistered to the fMRI data and then thresholded. Voxels with the lowest 30% of T2 signal intensity were excluded.

Removing voxels with high t-values: Previous work has found higher t-values in the large veins compared to smaller capillaries (Kashyap et al., 2018; Polimeni et al., 2010). A GLM analysis was performed on unsmoothed fMRI data in SPM. The model included two regressors of interest which were the 10 s recall periods of memories and the 10 s periods of counting. The six motion parameters were also modelled. For each participant, a t-value map of memory recall was produced by contrasting memory beta estimates with counting beta estimates. I then excluded voxels with the top 1% of t-values for the memory versus counting contrast.

For the ROIs produce by each of these methods, the laminar response profiles were then computed. First, I divided the cortical grey matter into 3 equivolume depth bins or layers (deep, middle and superficial) using the Open fMRI Analysis toolbox (https://github.com/TimVanMourik/OpenFmriAnalysis). Two additional surfaces were produced between the pial and white matter surfaces using the level set method (van Mourik, van der Eerden, et al., 2019; Waehnert et al., 2014). The distance of each ROI voxel to 5 compartments (white matter, CSF, and the 3 grey matter depth bins) was calculated, once again using the level method. From this, the distribution of each voxel's volume over the 5 compartments was calculated. For example, 80% of a voxel's volume might be in the superficial layer and 20% in the middle layer.

Next, the signal time course for each layer was estimated by means of a spatial GLM (van Mourik, van der Eerden, et al., 2019):

$$Y=XB+\varepsilon$$

Where X is an [n x k] spatial design matrix (n = number of voxels in the ROI, k = 5 cortical depths) describing the distribution of each voxel's volume over the 5 depths (WM, CSF and the three cortical layer depths). Each row of the design matrix X corresponds to the proportion of the voxel's volume within each of the 5 depths. Y is an [n x 1] vector of voxel signal intensities at a given time point (e.g., within a functional image) that is regressed against the spatial design matrix to obtain B, a [k x 1] vector of signal estimations for each layer. For each ROI, the GLM was applied to each functional volume within the time series to yield a signal time course for each layer.

A temporal GLM was then applied to the layer time courses. The model included two regressors of interest which were the 10 s recall periods of memories and the 10 s periods of counting, along with the six motion parameters. For each layer, the counting baseline beta estimate was subtracted from the memory retrieval beta estimate to obtain layer-specific responses to memory retrieval.

To measure the success of each of the methods for removing large vein voxels, I compared the laminar activation profiles produced by each method. A steep positive slope of increasing activation from the white matter to the CSF indicated that there may be a superficial layer bias remaining. Whereas, if the superficial layer bias had been corrected, there may be a non-linear relationship between cortical depth and response.

With all masking methods, except the low T2 signal intensity mask, there was an increase in response between each cortical depth from the white matter to CSF, indicating that a superficial layer bias effect may still exist in the response data (Figure 22).

The low T2 signal intensity mask contained more inter-subject variability than the other masking methods (Figure 22). This may be because the high resolution T2 weighted images contained severe signal drop out in several brain areas, particularly the right lateral temporal lobe and right occipital lobe, so many non-vein voxels were also removed with the low T2 signal intensity mask. Furthermore, I was concerned that rigid-body coregistration would not align the T2-weighted structural image to the functional images with the extremely high spatial precision that would be required to align veins. This would also result in non-vein voxels being removed. Removing non-vein voxels instead of vein voxels would reduce statistical power and increase noise. Therefore, I decided not proceed with this method.

The removed voxels using the high t-values mask did not correspond with the low tSNR and low T2 signal intensity voxels, and many of the removed voxels appeared to be in the deep layers of the cortical grey matter. Therefore, I decided not to proceed with this mask.

Although the low tSNR mask did not have a substantial impact on the laminar activation profile, I decided to proceed with it as there was no downside. However, clearly further measures were required to address the issue of the superficial layer bias.

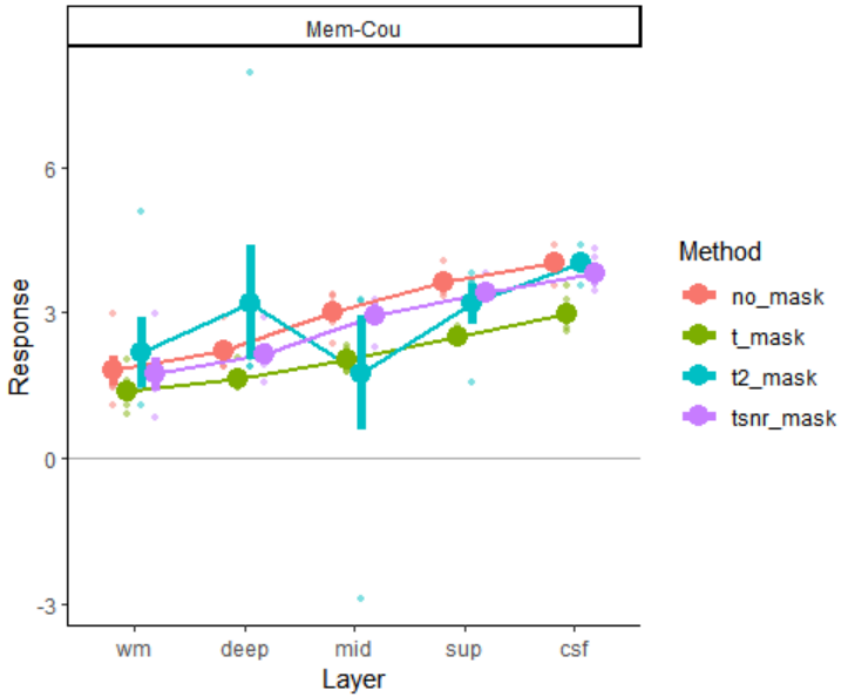

Figure 22. Exclusion of large vein voxels to address the superficial layer bias effect. Three masking methods were tested in the attempt to exclude large veins from the laminar analysis. Using no mask (red), the low tSNR mask (purple) or the high t-score mask (green) produced a continuous increase in response from the white matter to the CSF, indicating the existence of a superficial layer bias effect. The low T2 signal intensity mask (blue) produced a non-linear laminar response profile, suggesting that the superficial layer bias effect may be reduced, however, there was large variability (see error bars) between participants.

I tested 4 alternative ways of estimating laminar responses from the beta estimates, each with the aim of reducing the superficial layer bias effect.

Beta subtraction with tight contrasts: It is sometimes assumed that using a "tight" contrast subtraction, between two conditions that are both expected to activate the brain area, would remove the macrovascular component of the BOLD signal. This is based on the assumption that the signal from the large veins is of equal magnitude across all conditions that activate the brain area. To test this theory, a different temporal GLM was applied to the layer time courses, this time the recall periods of recent and remote memories were modelled separately, both of which are expected to activate the mPFC. The beta estimate for recent memory retrieval was subtracted from the beta estimate for remote retrieval, yielding the layer-specific responses to remote relative to recent retrieval.

Beta ratios: Another hypothesis is that the signal from the large veins scales with the degree of activation. To test this theory, instead of subtracting beta estimates for task conditions, the ratios of beta estimates were used to calculate the layer-specific responses (similar to Huang et al., 2021; Kashyap et al., 2018; Liu et al., 2021). For example, the memory retrieval beta estimate was divided by the counting beta estimate to yield the layer-specific response to memory retrieval. This method was tested for both the memory retrieval versus counting contrast and the remote versus recent contrast.

Normalising: Normalisation is another approach that is also based on the idea that the macrovascular signal scales with the degree of task-related activation. The estimated laminar responses (i.e., subtractions of condition beta estimates) were normalised by diving them by the mean BOLD signal of the layer (similar to Lawrence et al., 2019). By dividing by the mean BOLD signal of the layer, the scaling effect will be removed. This method was tested for both the memory retrieval versus counting contrast and the remote versus recent contrast.

Regression of other layer time courses: A regression method was used where, for each cortical depth, the time courses of the deeper cortical depths were used as nuisance regressors of in the temporal GLM (similar to Kok et al., 2016). Therefore, signal that correlated with the deeper layers was removed from the layer-specific beta estimate. This method is based on the assumption that signal leaks from the deep layers to the superficial layers due to the large draining veins. This method was tested for both the memory retrieval versus counting contrast and the remote versus recent contrast.

MVPA: An MVPA method was used where, for each voxel, a series of GLMs were fitted, one per trial, which modelled the trial compared to all other trials. The beta estimate for each trial was then noise-normalised by dividing by the square root of the GLM residuals, thereby down-weighting noisier voxels and yielding a response estimate per voxel (Walther et al., 2016). Then, I used RSA analysis of layer ROIs. Layer ROIs were produced by binning voxels with at least 80% of their volume in a given layer into that layer's ROI. For each layer ROI, the correlation between the multivoxel response patterns of trials within the same condition was calculated and separately, the correlation between multivoxel response patterns of trials of different conditions was calculated. The between-condition correlation was subtracted from the within-condition correlation to yield a value corresponding to the decodability of the condition. In MVPA, we are interested in the consistency of activation patterns across task

trials rather than the overall amplitude of activation. Therefore, if there was a large macrovascular signal in the superficial layers and weaker macrovascular signal in the deep layers, this difference in overall magnitude of signal would not impact the decoding results. Furthermore, MVPA affords the opportunity to use even tighter contrasts to investigate the decoding of individual stimuli. Therefore 3 contrasts were used for this method, memory versus counting, remote versus recent, and individual memory versus all other memories.

In order to compare the laminar response profiles between these different response metrics, the responses calculated by beta ratios, normalisation, and MVPA were normalised to the same scale as the responses calculated by beta subtraction (which was used in all other methods). To do this, I calculated the ratio of the subtraction and normalisation response scales and then multiplied the normalisation response values by this ratio. I then repeated this for the beta ratio responses.

Using beta ratios instead of subtractions resulted in extreme outliers. This may be because this method has high sensitivity to counting or recent beta values that are close to zero. Divisions by these close to zero values resulted in extreme outliers. For this reason I decided not to pursue this method and the results are not included in Figure 23.

For the memory versus counting contrast, the only method that substantially impacted the laminar response profile was the layer regression method, which resulted in a decrease (rather than increase) in response amplitude from the white matter to the CSF (Figure 23a). However, I was concerned that this method may remove signal that is correlated across layers due to the fact that each layer is functionally connected to the other layers. Regressing out this correlated signal in the superficial but not deep layers may result in a biasing of signal towards the deep layers. Therefore, I decided not to pursue this method. Normalising by the mean BOLD signal of the layer had little effect on the laminar response profile of memory retrieval versus counting baseline, nor did multivariate decoding of the memory retrieval cognitive state.

Using a tight condition contrasts (i.e., remote memory retrieval versus recent memory retrieval) produced laminar response profiles with shallower slopes than those produced by the memory retrieval versus counting contrast (Figure 23b). However, for all methods except the MVPA method, responses still clearly increased incrementally from white matter to CSF

(Figure 23b). When using MVPA to perform stimulus category decoding of recent versus remote memory retrieval, the responses decreased from the white matter to the deep layer and from the superficial layer to the CSF, indicating that a superficial layer bias may be somewhat corrected (Figure 23b).

Finally, when using MVPA to decode individual stimuli, there was a negative slope from the white matter to the CSF, indicating that this very tight contrast produced little to no superficial layer bias (Figure 23c).

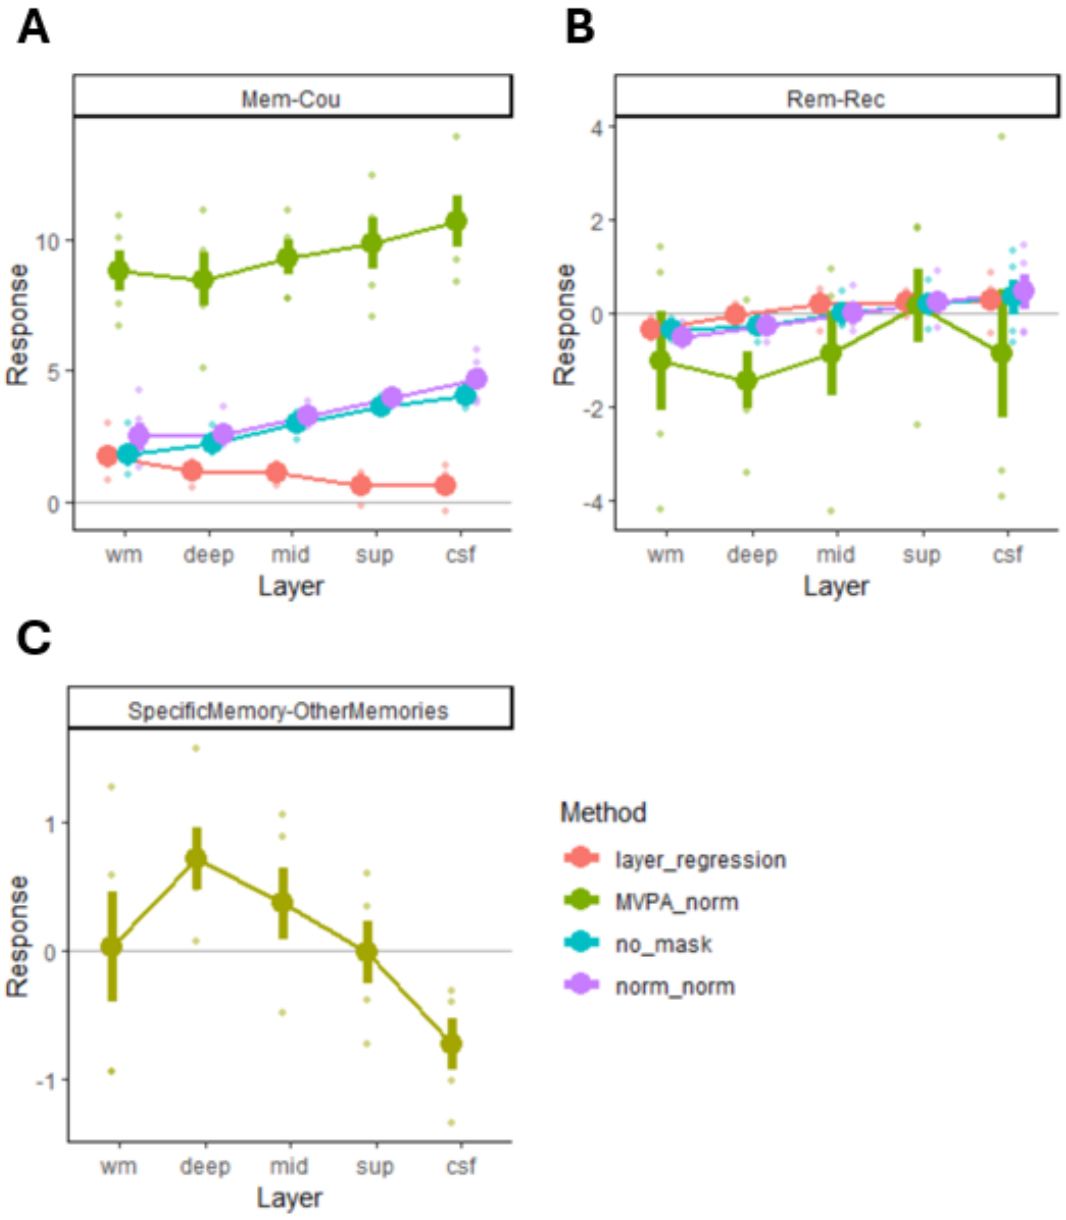

Figure 23. The effect of analysis method on superficial layer bias effect. The laminar response profiles produced by four different analysis methods were compared. (a) When using a broad condition contrast of memory retrieval versus counting, layer regression was the only analysis method that did

not produce an increase in response from the white matter to the CSF. When using a tighter condition contrast of (b) remote versus recent memory retrieval or (c) specific memory versus all other memories, MVPA was the only method that did not produce an increase in response from the white matter to the CSF.

3.5. Discussion

In this study, I first worked with the Physics Group to optimise the fMRI sequence with the aim of mitigating geometric distortion, signal drop out and signal artefacts in the functional images (Experiment 1a). The fMRI sequence parameter that made the most drastic improvements to distortion and signal drop out was use of a two-fold segmented acquisition, while apodisation of k-space during reconstruction effectively removed a Gibb's ringing artefact. However, it was not possible to entirely remove geometric distortion, signal drop out and signal artefacts during image acquisition. Therefore, I next tested methods for correcting the remaining distortion during image preprocessing (Experiment 1b). I found that unwarping of functional images with a B0 field map, estimated from reverse phase encoded images, combined with recursive BBR resulted in the most accurate coregistration of cortical surfaces to the functional images. I also tested methods for correcting a signal artefact coming from cortical veins, known as the superficial layer bias effect (Experiment 1b). I found that attempting to remove voxels containing large veins minimally reduced the bias of signal towards the superficial layers. Whereas, analysis decisions such as examining multivoxel response patterns using tight condition contrasts were successful in addressing the bias effect. Other optimisations were made to the preprocessing pipeline, for example, using a "spliced" T2-weighted image to accurately segment the cortical surfaces (Experiment 1b).

3.5.1. Signal drop out and geometric distortion

Segmentation is an acquisition method that was developed many years ago to address geometric distortion (McKinnon, 1993). Since then, it has been successfully used for 7T fMRI, specifically with 3D GE-BOLD EPI imaging (Zwanenburg et al., 2011). Our sequence parameter tests provide further support for 2-fold interleaved segmentation as an acquisition method for dealing with distortion and drop out.

In terms of preprocessing methods, I found that recursive BBR alone was insufficient to align the cortical surfaces with the functional images in the mPFC. Unwarping the functional images

prior to recursive BBR produced considerably better alignment. Unwarping is particularly necessary in the mPFC, which contains more distortion than other brain areas. To minimise interpolation and resampling of the data, I applied realignment and unwarping in one step using SPM Realign and Unwarp. However, even with this approach, the most distorted areas of the functional images, with severe pile-up of signal, appeared to be spatially blurred after unwarping. This is because, with significant signal pile-up, the unwarping algorithm becomes poorly conditioned and is unable to accurately divide and reposition the signal. These highly distorted areas were successfully excluded from laminar analyses using a Jacobian mask.

While geometric distortion can be corrected during image preprocessing, signal dropout cannot as the missing signal cannot be retrieved. While we do want the cortical surfaces to warp to distorted parts of the image, we do not want them to warp around areas with signal drop out. This is because it may lead to inaccurate allocation of signal to each of the cortical layer depths. For instance, signal dropout in the superficial layers could cause the pial surface to warp to the middle layer, leading to incorrect attribution of the middle layer's signal to the superficial layer. In attempt to only warp the surfaces to distorted areas of the image, during recursive BBR, we only allowed scaling of the cortical surfaces in the phase encoding direction, where distortion is expected to be most severe. Furthermore, the areas of the brain that contained the most significant signal drop out were excluded from the laminar analyses by using a thresholded Jacobian map.

3.5.2. Correcting for the superficial layer bias

In an attempt to remove voxels containing large veins, while retaining the voxels containing smaller layer-specific capillaries, I tried using several different masking methods: a tSNR mask, t-value mask, and T2 signal intensity mask. However, none of these methods had a substantial impact on the superficial layer bias effect. One explanation could be that the large vein voxels were not effectively removed due to inadequate thresholds. However, using more lenient thresholds caused large portions of the brain to be removed by the mask. This was particularly the case for the T2 signal intensity mask, which contained a considerable amount of signal drop out. Another explanation is that we do not have the required spatial resolution to delineate the large veins that penetrate the cortical grey matter from the smaller capillaries that are also within the cortical grey matter. Therefore, it is possible that we only removed the

very large veins on the pial surface, which would only accounts for one of the physiological causes of the superficial layer bias effect.

I tested out several different methods for computing laminar-specific responses: beta subtractions, beta ratios, normalisation, layer regression, and MVPA. For the very broad memory versus counting contrast, I found that layer regression was the only method that removed the superficial layer bias, reversing the laminar response profile such that it was skewed towards deeper layers. However, this may be due to the removal of functionally relevant signal that is correlated across layers. Removing signal only in superficial layers while retaining it in deeper layers introduces a new bias. A previous study using layer regression regressed out all other layers' time courses, avoiding this bias but still removing informative signal (Kok et al., 2016). I decided not to pursue this as an option as removal of any functionally relevant signal from the layers is not desirable.

When performing tighter condition contrasts, MVPA was the only method resistant to the superficial layer bias effect. However, MVPA was not resistant when decoding the broader memory retrieval state. This discrepancy can be understood by considering how different types of information are represented in the brain. Different cognitive states involve different brain areas, but the multivoxel activation patterns within these areas are probably less relevant. In contrast, processing of different individual stimuli likely involves the same brain areas but the specific ensembles of neurons within those areas may be specific to the stimulus. Therefore, for cognitive states, any correlated multivoxel response across recalls would be driven by the voxels containing less spatially specific, more diffuse signal, such as those containing macrovascular signal. Whereas, the voxels containing more spatially precise microvascular signal would be inconsistent across memory retrieval trials. Conversely, when decoding individual stimuli or categories, the spatially specific microvascular signal would be correlated across stimulus repeats, while the correlation driven by the less spatially specific macrovascular signal would be cancelled out by subtracting the correlation among trials of other stimuli. As such, for tight contrasts, such as decoding of specific stimuli and stimulus categories, MVPA would be more sensitive to the microvascular signal making it less susceptible to the superficial layer bias effect.

Previous studies have produced mixed results regarding MVPA's resistance to the superficial layer bias effect (Huang et al., 2021; Vizioli et al., 2020). Huang et al. (2021) found that MVPA

was not resistant, while Vizioli et al. (2020) found it was. This discrepancy may be due to the type of information being decoded. Huang et al. (2021) decoded attentional states, whereas Vizioli et al. (2021) decoded individual scene images. A question arises as to whether using MVPA to decode individual stimuli could lead to a bias towards the deep layers due to the more spatially precise microvascular signal in these layers. However, this appears not to be the case. Vizioli et al. (2020) artificially misaligned the multivariate response patterns by one voxel and found an equal decrease in decoding accuracy across the cortical layers, indicating comparable spatial precision of the multivoxel activation patterns in each of the cortical layers.

Clearly, it is very difficult to fully and accurately correct the superficial layer bias artefact, likely due to a combination of vascular causes, which may vary between brain areas depending on the vasculature. Some modelling methods have been developed that I did not test out as they are based on the anatomy of the visual cortex and so their generalisability to other cortical areas is unknown (Havlicek & Uludag, 2020; Heinzle et al., 2016; Markuerkiaga et al., 2016).

Furthermore, alternative MRI sequences have been developed that measure blood flow, rather than changes in blood oxygen, which do not have a superficial layer bias effect. There are two main types of blood flow-based sequence, vascular-space-occupancy (VASO) sequences that measure cerebral blood volume (CBV; Huber et al., 2017; Lu et al., 2013), and arterial spin labelling (ASL) sequences that measure cerebral blood flow (CBF; Huber et al., 2017; Petcharunpaisan et al., 2010). However, these acquisition methods have less functional sensitivity and long acquisition times, and/or small fields of view (Huber et al., 2017).

Given the inability to satisfactorily correct the superficial layer bias using univariate methods and the downsides of alternative acquisition methods, we decided against pursuing univariate laminar analyses and focused on multivariate analyses of individual stimuli or stimulus categories.

A limitation of using the autobiographical memory task to assess the superficial layer bias is the lack of clear expectations about the true laminar response profile. We assumed that a steeper positive slope indicated a stronger superficial layer bias effect, which means that a true superficially skewed laminar response could be misinterpreted as bias. Typically, methods for correcting superficial layer bias are validated using simulations or tasks with well-understood laminar activation profiles, such as feedforward visual input into the visual cortex

(e.g., Vizioli et al., 2020). However, as this is the first laminar fMRI study of the mPFC, there are no established tasks that can predict the true laminar response profile in this region.

In these experiments, we explored various MRI data acquisition, preprocessing and analysis methods to address the challenges of geometric distortion, signal dropout, and the superficial layer bias effect in laminar GE-BOLD fMRI. Using a 2-fold segmented acquisition, fMRI data unwarping, and recursive BBR significantly improved distortion and cortical surface alignment in the mPFC. However, signal dropout and extreme levels of distortion remained an issue, necessitating the exclusion of highly distorted and low signal areas using Jacobian maps. We tested multiple techniques to correct for the superficial layer bias artefact. Despite some success, none of the univariate analysis methods entirely eliminated the bias, highlighting the complexity of this artefact. Therefore, we concluded that univariate laminar analyses were not viable. MVPA was successful in addressing the bias when being used to decode stimulus categories and individual stimuli. Therefore, we decided to proceed with multivariate laminar analyses to understand event processing in the mPFC and wider network of brain areas. These experiments underscore the need for continued refinement of acquisition, preprocessing and analysis techniques to accurately interpret laminar fMRI data.

4. Experiment 2: Automated protocols for delineating human hippocampal subfields from 3T and 7T MRI data

4.1. Introduction

Researchers who study the human hippocampus are naturally interested in how its subfields function. However, many researchers are precluded from examining subfields because their manual delineation from MRI scans (still the gold standard approach) is time consuming and requires significant expertise. To help ameliorate this issue, I present here two protocols, one for 3T MRI and the other for 7T MRI, that permit automated hippocampus segmentation into six subregions, namely DG/CA4, CA2/3, CA1, subiculum, pre/parasubiculum and uncus along the entire length of the hippocampus. These protocols are particularly notable relative to existing resources in that they were trained and tested using large numbers of healthy young adults (N = 140 at 3T, N = 40 at 7T) whose hippocampi were manually segmented by experts from MRI scans. Using inter-rater reliability analyses, I showed that the quality of automated segmentations produced by these protocols was high and comparable to expert manual segmenters. I provide full open access to the automated protocols, and anticipate they will save hippocampus researchers a significant amount of time. They could also help to catalyse subfield research, which is essential for gaining a full understanding of how the hippocampus functions.

4.2. Background

The hippocampus is composed of distinct subfields that have different functions (Bonnici et al., 2013; Neunuebel & Knierim, 2014; Zeidman, Lutti, et al., 2015) and connectivity patterns (Dalton et al., 2022; Dalton et al., 2019). Studying human hippocampal subfields typically requires them to be delineated from structural MRI scans. Manual segmentation of subfields from such scans remains the gold standard approach (Dalton et al., 2017; Wisse et al., 2017). However, there are currently numerous protocols available that differ on a range of parameters. These include the location of borders between subfields, the granularity with which specific subfields can be individually delineated, and whether nor not the full length of the hippocampus is segmented (Yushkevich, Amaral, et al., 2015). Automated methods for subfield segmentation have also been developed, which essentially recapitulate these

protocol differences (e.g., Augustinack et al., 2013; DeKraker et al., 2022; Fischl et al., 2009; Hadar et al., 2018; Iglesias et al., 2015; Pipitone et al., 2014; Romero et al., 2017; Van Leemput et al., 2009; Wisse et al., 2016; Yushkevich, Pluta, et al., 2015; Yushkevich et al., 2010). This variety of approaches hinders the ability to integrate and interpret findings across studies.

To address this issue, the Hippocampal Subfield Group (HSG: http://www.hippocampalsubfields.com) was convened. It includes colleagues from around the world who are working towards the production of a reliable, validated, harmonised protocol for manual segmentation of hippocampal subfields and associated medial temporal lobe regions. This huge undertaking is progressing, and will eventually be an invaluable and unifying resource for the field of hippocampus neuroscience (Olsen et al., 2019; Wisse et al., 2017; Yushkevich, Amaral, et al., 2015).

In the meantime, the Maguire Lab previously devised a detailed manual subfield segmentation protocol that is faithful to our current understanding of hippocampal anatomy (Dalton et al., 2017). I favour this protocol because it allows the whole hippocampus to be segmented into 6 subregions: DG/CA4, CA2/3, CA1, subiculum, pre/parasubiculum, and the uncus. It takes approximately eight hours for an expert to segment a person's two hippocampi from T2-weighted isotropic voxel MRI scans using this protocol. But such expertise is not available in all research groups, and even when it is, the time sink involved often prohibits conducting experiments at scale, such as those examining individual differences.

Consequently, the goal of this experiment was to devise an automated subfield segmentation protocol based on the Dalton et al. (2017) approach. I could pursue this aim thanks to two open access resources. The first is the recently released open access dataset from Clark and Maguire (2023) that includes the largest set of manually segmented hippocampal subfields of healthy young adults that is currently available using the Dalton et al. (2017) approach. The second is the software package Automatic Segmentation of Hippocampal Subfields (ASHS) (Yushkevich, Pluta, et al., 2015) which is free, open source and editable. ASHS can be retrained using one's own MRI data and segmentation protocol to produce an ASHS "atlas package" that can then be used to automatically segment subfields from new MRI scans. We aimed to produce two ASHS atlases, one for 3T MRI data and another for 7T MRI data.

4.3. Materials and methods

4.3.1. Participants

For building and testing the 3T ASHS atlas, I availed of previously acquired and open access 3T MRI scans from 140 healthy adult participants (81 females) aged between 20 and 41 years old (mean age = 29.09, SD = 5.61; Clark & Maguire, 2023). For building and testing the 7T ASHS atlas, I recruited and scanned using 7T MRI 40 healthy adult participants (26 females) aged between 18 and 33 years old (mean age = 23.8, SD = 3.99).

4.3.2. Experimental protocol

All 140 3T MRI participants undertook the scanning in a dedicated structural MRI session. Of the 140 3T MRI participants, I randomly allocated 125 participants to a 3T training group, which was used to build the 3T ASHS atlas package. This provided an unprecedented opportunity for the ASHS algorithm to train on a large number of expertly segmented hippocampi. The remaining 15 participants were assigned to a 3T testing group that was used to assess the 3T ASHS atlas package.

Of the 40 7T MRI participants, 22 participants undertook the structural imaging session in a dedicated structural scanning session (as part of visit 1 of Experiment 3, Chapter 5), 12 participants undertook the structural imaging immediately following an fMRI session (as part of Experiment 4, Chapter 6), and 6 participants were scanned in a dedicated structural imaging session as part of a different 7T MRI study on which I am a collaborator, but that is not reported in this thesis.

Of the 40 7T MRI participants, I randomly allocated 30 of the participants to a 7T training group, which was used for training the 7T ASHS atlas package. The remaining 10 participants were assigned to a 7T testing group that was used to assess the quality of automatic segmentations produced by the 7T ASHS atlas package.

Participants were not required to do anything during structural scanning, other than rest and remain as still as possible.

4.3.3. MRI sequences

All 3T data were collected using Siemens Magnetom TIM Trio 3T MRI systems with a 32 channel head coil. For each 3T MRI participant, 3 separate T2-weighted partial volume images were acquired using a 3D turbo spin echo sequence (Mugler et al., 2000). During this sequence, a single axial slab was excited and the T2-weighted signal was read out. The sequence used the following parameters: voxel size = $0.52 \times 0.52 \times 0.5$ mm^3, matrix = 384 × 328, partitions = 104, partition thickness = 0.5 mm, partition oversampling = 15.4%, field of view = 200 × 171 mm, TE = 353 ms, TR = 3200 ms, GRAPPA × 2 in phase-encoding direction, bandwidth = 434 Hz/pixel, echo spacing = 4.98 ms, turbo factor in PE direction = 177, echo train duration = 881, averages = 1.9. To minimise signal bias (e.g., caused by spatial variation in coil sensitivity profiles), the images were normalised using a pre-scan and a low-intensity filter was applied. Each image took 13 minutes to acquire.

ASHS also requires a T1-weighted image for initial localisation of the whole hippocampus. For each 3T MRI participant, a whole brain T1-weighted GE image was acquired as part of a Multi-Parameter Mapping quantitative imaging protocol (Callaghan et al., 2015; Callaghan et al., 2019). The sequence had the following parameters: voxel size = $0.8 \times 0.8 \times 0.8$ mm^3, TR = 25 ms, flip angle = 21°, field of view = 256 mm × 224 mm × 179 mm, TE = 8 equidistant read outs ranging from 2.34 to 18.44 ms in steps of 2.30 ms, bandwidth = 488 Hz/pixel, GRAPPA x 2 in each phase encoding direction (anterior-to-posterior and right-to-left). The image took 25 minutes to acquire.

All 7T MRI data were acquired on a Siemens Magnetom Terra 7T MRI system (see Chapter 2, Section 2.6.1). For each 7T MRI participant, 2 partial volume T2-weighted images were acquired with the same sequence that was used at 3T, but adapted for 7T. The sequence had a voxel size of 0.52 x 0.52 x 0.50 mm^3, TR = 3,500 ms, TE = 229 ms, flip angle = 8°, field of view = 200 × 169 mm x 56 mm, matrix = 384 x 324 x 112, partitions = 112, partition thickness = 0.5 mm, partition oversampling = 14.3%, GRAPPA × 2 in phase-encoding direction, bandwidth = 868 Hz/Px, echo spacing = 3.83 ms, turbo factor = 176, echo train duration = 548 ms, averages = 1.4. Each image took 10 m and 25 s to acquire.

For each 7T MRI participant, a whole brain T1-weighted image was acquired using an MP2RAGE sequence (Marques et al., 2010), with the following parameters: voxel size = 0.65 x

0.65 x 0.65 mm^3, TR = 5000 ms, TE = 2.54 ms, TI = 900 ms and 2,750 ms, 5° and 3° flip angles, field of view = 208 × 208 × 156 mm^3, and GRAPPA x 3 in phase encoding direction. The image took 8 minutes 42 seconds to acquire.

Of note, all 3T and 7T images, such as those here, had isotropic voxels, meaning there was equivalent resolution in all orientations, which is particularly important for convoluted structures like the hippocampus, which can contain undulations in the subfields.

4.3.4. Data preprocessing

The T2-weighted images were visually inspected and those with poor image quality were discarded. For each participant, the remaining T2-weighted images were realigned, denoised and averaged. This averaging and denoising method improved the SNR of the T2-weighted image used for hippocampal subfield segmentation. Although, note that individual 7T images already had sufficient SNR for subfield segmentation, which could mean shorter structural imaging sessions for future studies. The T1-weighted image was then coregistered to the mean T2-weighted image using NMI.

4.3.5. Manual segmentation of hippocampal subfields

ITK Snap software version 3.2.0 was used to manually delineate the hippocampal subfields on the averaged and denoised T2-weighted images in line with the Dalton et al. (2017) protocol. For each participant the two hippocampi were manually segmented into 6 subregions. The DG/CA4 was segmented first, followed by CA2/3, CA1, subiculum, pre/parasubiculum and then the uncus. The subfields were segmented on coronal slices, starting at the most anterior portion of the hippocampus and moving in the posterior direction. Anatomical markers, described in Dalton et al. (2017), were used to define the boundaries between subfields. Examples of these included: (1) the vestigial hippocampal sulcus at the outermost boundary of the DG/CA4, which is visible as a “C” shape in the right hemisphere (or inverted “C” shape in the left hemisphere) of high intensity voxels, and (2) the increase in white matter projections from the perforant pathway at the boundary between the subiculum and pre/parasubiculum, which is visible as a slight increase in T2 signal intensity. I performed the manual segmentations on all 7T MR images and the 3T MR images had been previously manually segmented by another researcher.

Inter-rater reliability analysis was used to assess the reliability of the manual hippocampal segmentations. To perform this analysis, a subset of images were segmented for a second time by a different researcher. The subset included 20 out of the 125 participants in the 3T training group, and 9 out of the 30 participants in the 7T training group. The Dice similarity coefficient (Dice, 1945), which ranges from 0 (no overlap) to 1 (complete overlap), was used to calculate the degree of overlap between the subfield masks produced by the first and second researchers. The results are shown on Table 2, and for both 3T and 7T MRI they were comparable to previous studies using this protocol (Barry et al., 2021; Clark et al., 2023; Dalton et al., 2019) and other approaches (Berron et al., 2017; Bonnici et al., 2013; Chadwick et al., 2014; Lee et al., 2014; Palombo et al., 2013; Yeung et al., 2019).

Table 2. Inter-rater reliability results for the two manual segmenters.

Subfield	Dice similarity coefficient (mean ± SD)			
	3T MRI		7T MRI	
	Left	Right	Left	Right
DG	0.84 ± 0.02	0.85 ± 0.03	0.87 ± 0.02	0.86 ± 0.03
CA2/3	0.67 ± 0.04	0.68 ± 0.04	0.69 ± 0.02	0.68 ± 0.04
CA1	0.78 ± 0.03	0.79 ± 0.02	0.78 ± 0.03	0.80 ± 0.02
Subiculum	0.81 ± 0.02	0.79 ± 0.02	0.81 ± 0.03	0.79 ± 0.03
Pre/parasubiculum	0.71 ± 0.03	0.69 ± 0.03	0.70 ± 0.04	0.72 ± 0.02
Uncus	0.82 ± 0.03	0.84 ± 0.02	0.88 ± 0.03	0.86 ± 0.04

Note: Two researchers independently segmented the hippocampi of 20 out of the 125 participants in the 3T training group, and 9 out of the 30 participants in the 7T training group.

4.3.6. Automatic segmentation of hippocampal subfields

Having established that the manual segmentations were reliable, I next used them to produce the ASHS atlases. I trained ASHS on the manual hippocampal subfield segmentations of the 3T training group sample (125 participants) and separately on the 7T training group sample (30 participants) to produce a 3T and 7T ASHS atlas, respectively.

The ASHS training pipeline was as follows (Yushkevich, Pluta, et al., 2015):

1. **Preprocessing**: For each participant, the T2-weighted image was coregistered to the T1-weighted image using NMI. Next, the T1-weighted images from all participants were normalised to a template using the Advanced Normalization Tools (ANTs) deformable coregistration tool. This template is provided in the atlas package. The template image was then cropped around the left and right hippocampi (separately). To achieve this, for each hemisphere, subfield masks were combined to produce a whole hippocampus mask. This mask was then transformed into the template space and used to crop the template image ensuring that the whole hippocampus was included. Finally, the same transformation and cropping was applied to all the T1-weighted images, T2-weighted images and segmentation images so that all images were in the template space.
2. **Joint label fusion (JLF)**: For each participant, an initial automatic segmentation was produced. To achieve this, the T2-weighted images (and corresponding segmentations) of all other participants were warped to the T2-weighted image of the participant using ANTs deformable coregistration. Then a JLF algorithm was applied to all of the warped segmentations (the following explanation is based on Wang et al., 2013; Wang & Yushkevich, 2013; and Yushkevich, Pluta, et al., 2015). JLF combines segmentations from multiple participants into a single segmentation. For each input segmentation, a measure of similarity between it and the other segmentations was computed, based on the relative T2 signal intensities of a patch centred on a voxel. The similarity measure helps to determine how much weighting (i.e., trust) to place in the segmentation at each voxel. The final subfield label for each voxel was determined by a weighted voting scheme. The consensus label for each voxel was the one that receives the highest total weighted vote. JLF calculates the correlation between pairs of segmentations, and reduces the weights of correlated images. This reduces a bias of the segmentation towards more typical anatomy, allowing the final segmentation to also represent the more atypical cases.
3. **Corrective learning classification**: Corrective learning classifiers were then used to correct any mislabelling of subfields by the JFL method (the following explanation is based on Wang & Yushkevich, 2013; and Yushkevich, Pluta, et al., 2015). For each subfield (in each separate hemisphere), a classifier was trained to predict whether a given voxel belonged to the subfield based on various features of the voxel. Each voxel

that was labelled as belonging to the subfield by the JLF method in each training image was used as a training instance, as were voxels surrounding the JFL segmented subfield. Rather than including all other voxels in the hippocampus, only the voxels surrounding the subfield were included as these are the most likely voxels to have been mislabelled by JLF. They were selected by slightly dilating (expanding) the JFL segmented subfield for each participant image. Several voxel features were used for classification. These included the T2 signal intensity values in a patch centred on each voxel, the posterior probability maps produced by the JFL algorithm for the patch, and the location of the voxel in relation to the centre of the subfield (as defined in the initial JFL segmentation). Using spatial and contextual information from the JFL segmentation as features gives these classifiers more power than most machine learning methods used for medical image analysis, which typically only use features extracted from the images (Morra et al., 2010; Tu & Bai, 2010; Tu et al., 2007; reviewed in Wang & Yushkevich, 2013). Each subfield classifier algorithm was iteratively trained using a type of corrective learning called AdaBoost ("adaptive" boosting). On each iteration a "weak" classifier is trained and then tested using a leave-one-out approach (in each iteration a single participant is excluded from training and is used for testing instead). Then, the weights of the training instances (i.e., the voxels) were updated to give more emphasis to the voxels that were misclassified, thereby guiding the next weak classifier to focus more on these difficult cases. This process is repeated for a specified number of iterations. A classifier weight is also assigned to each weak classifier based on its accuracy. These weights are used to construct a final "strong" classifier, which is the weighted sum of all the weak classifiers from each iteration.

4.3.7. Evaluating performance

Two approaches were used to evaluate performance of the 3T and 7T ASHS atlases. The first approach used cross-validation tests. Here, ASHS atlases were iteratively retrained using a leave-K-out approach, in which K participants were excluded from the atlas build and were used to test the atlas. We used leave-one-out cross-validation to test performance of the 7T ASHS atlas and, due to the larger 3T training group sample and for computational efficiency, we used leave-five-out cross-validation to test performance of the 3T ASHS atlas.

In the second, and more reliable, approach I used the final 3T and 7T ASHS atlas packages to segment the hippocampi of participants in the new 3T testing group (15 participants) and new 7T testing group (10 participants) respectively. This allowed me to test the complete and final 3T and 7T ASHS atlases that will be used by future scientists wishing to segment subfields.

The ASHS segmentation pipeline for a given testing group participant is as follows (Yushkevich, Pluta, et al., 2015):

1. **Preprocessing**: First, the participant's T2-weighted image is coregistered to their T1-weighted image using NMI. Then, the T1-weighted image is warped (using ANTS deformable registration) to the group template (provided in the atlas package) and the same warp is applied to the T2-weighted image. All subsequent steps are performed in template space.
2. **JLF**: Each of the training group T2-weighted images and corresponding manual segmentations (provided in the atlas package) are warped (using ANTS deformable registration) to the testing participant's T2-weighted image. JFL is used to produce an initial segmentation of the new participant's T2-weighted image.
3. **Corrective learning classification**: The corrective learning classifiers (provided in the atlas package) are applied to the initial segmentation (the following explanation is based on Wang & Yushkevich, 2013). For each subfield, voxels are selected by dilating the JFL segmentation. Each voxel is then classified by the corrective learning classifier as belonging to or not belonging to the subfield. If multiple corrective learning classifiers determine that a voxel belongs to their subfield, the subfield whose classifier has the highest probability is assigned to the voxel. This results in a refined segmentation where mislabeled voxels from the JLF method are corrected based on the more accurate classifications provided by the corrective learning classifiers.
4. **Bootstrapping:** It is important that the warping between the participant's T2-weighted image and all of the training participant's T2-weighted images is accurate in order to achieve high quality segmentations. Therefore, Steps 2 and 3 are repeated, but this time the segmentation is used to improve the deformable coregistration.

The automatic segmentations were then compared to their corresponding manual segmentation using the Dice similarity coefficient (Dice, 1945) between the manual and automatic segmentations was computed. When Dice similarity coefficients are calculated

between manual segmentations performed by independent expert segmenters (i.e., in inter-rater reliability analyses), > 0.83 would be considered good for the DG/CA4, > 0.67 for CA2/3, > 0.77 for CA1, > 0.77 for the subiculum, > 0.68 for the pre/parasubiculum, and > 0.81 for the uncus (Barry et al., 2021; Clark et al., 2023; Dalton et al., 2019). Note that Dice similarity coefficients are typically higher for the larger subfields (e.g., DG/CA4; the uncus) compared to the smaller subfields (e.g., CA2/3; pre/parasubiculum). Finally, I also visually inspected the automatic segmentations for systematic errors in the segmentation of each subfield, of which there were none.

4.4. Results

The results of the cross-validation tests are shown in Table 3. In all subfields, the ASHS atlases resulted in Dice similarity coefficients that were, on average, higher than those typically achieved by manual segmenters using the same segmentation protocol.

Table 3. Reliability results for the ASHS atlases based on cross-validation tests.

Subfield	Dice similarity coefficient (mean ± SD)			
	3T MRI ASHS atlas		7T MRI ASHS atlas	
	Left	Right	Left	Right
DG	0.87 ± 0.02	0.87 ± 0.02	0.89 ± 0.06	0.90 ± 0.06
CA2/3	0.72 ± 0.07	0.71 ± 0.06	0.77 ± 0.12	0.79 ± 0.06
CA1	0.81 ± 0.02	0.81 ± 0.02	0.83 ± 0.07	0.84 ± 0.06
Subiculum	0.82 ± 0.03	0.81 ± 0.03	0.85 ± 0.05	0.86 ± 0.04
Pre/parasubiculum	0.70 ± 0.05	0.68 ± 0.06	0.75 ± 0.07	0.77 ± 0.06
Uncus	0.84 ± 0.04	0.84 ± 0.04	0.88 ± 0.05	0.88 ± 0.05

Note: Segmentations performed in the cross-validation tests were compared with manual segmentations. For the 3T MRI atlas, 24 cross-validation tests were performed, each on 5 participants. For the 7T MRI atlas, 30 cross-validation tests were performed, each on 1 participant.

The results of the tests performed on the final and complete ASHS atlases are shown in Table 4 and are plotted in Figure 24. Again, both atlases achieved Dice similarity coefficients that were, on average, higher than is typically achieved with manual segmentations using the same segmentation protocol.

Table 4. Reliability results for the ASHS atlases when compared to an independent set of manually segmented images.

Subfield	Dice similarity coefficient (mean ± SD)			
	3T MRI ASHS atlas		7T MRI ASHS atlas	
	Left	Right	Left	Right
DG	0.88 ± 0.03	0.89 ± 0.01	0.87 ± 0.03	0.89 ± 0.01
CA2/3	0.74 ± 0.04	0.75 ± 0.05	0.75 ± 0.05	0.78 ± 0.03
CA1	0.81 ± 0.03	0.83 ± 0.02	0.81 ± 0.03	0.82 ± 0.02
Subiculum	0.83 ± 0.02	0.82 ± 0.02	0.83 ± 0.02	0.84 ± 0.02
Pre/parasubiculum	0.72 ± 0.05	0.72 ± 0.05	0.71 ± 0.05	0.74 ± 0.02
Uncus	0.85 ± 0.02	0.85 ± 0.02	0.84 ± 0.03	0.86 ± 0.03

Note: Our 3T and 7T ASHS atlas packages were used to segment the hippocampi of participants in a new 3T test group (n=15) and a new 7T test group (n=10) respectively.

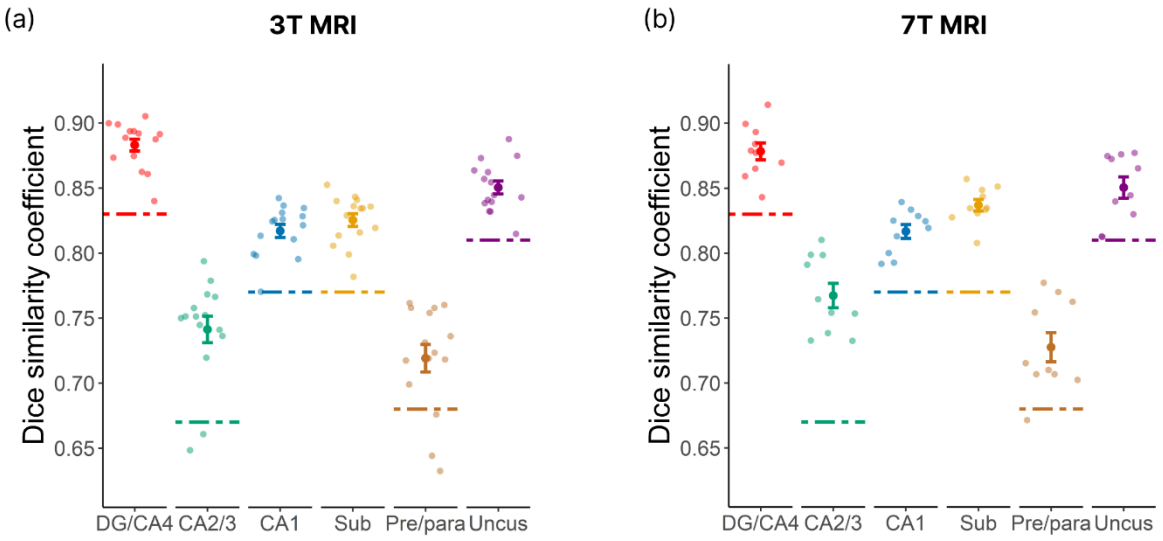

Figure 24. Similarity between automatic and manual segmentations. Dice similarity coefficients for the comparison between ASHS automatic segmentations and the corresponding manual segmentations, where (a) relates to our complete 3T ASHS atlas package used to segment hippocampi from new 3T MRI data (n = 15), and (b) relates to our complete 7T ASHS atlas package used to segment hippocampi from new 7T MRI data (n = 10). The overall mean Dice similarity coefficients are plotted with opaque points and the mean Dice similarity coefficients for each test participant's data are plotted with transparent points. Error bars indicate the standard error of the mean. To aid interpretation, the dashed lines provide an approximate indication of the typical inter-rater reliability thresholds between

human manual segmenters for the protocol used here. DG = DG; Sub = subiculum; Pre/para = pre/parasubiculum.

When visually inspecting automatic segmentations performed using the final atlases and their corresponding manual segmentations, the subregion masks appeared very similar (Figure 25). There were some minor local differences, which are also common in manual segmentations performed by different human segmenters.

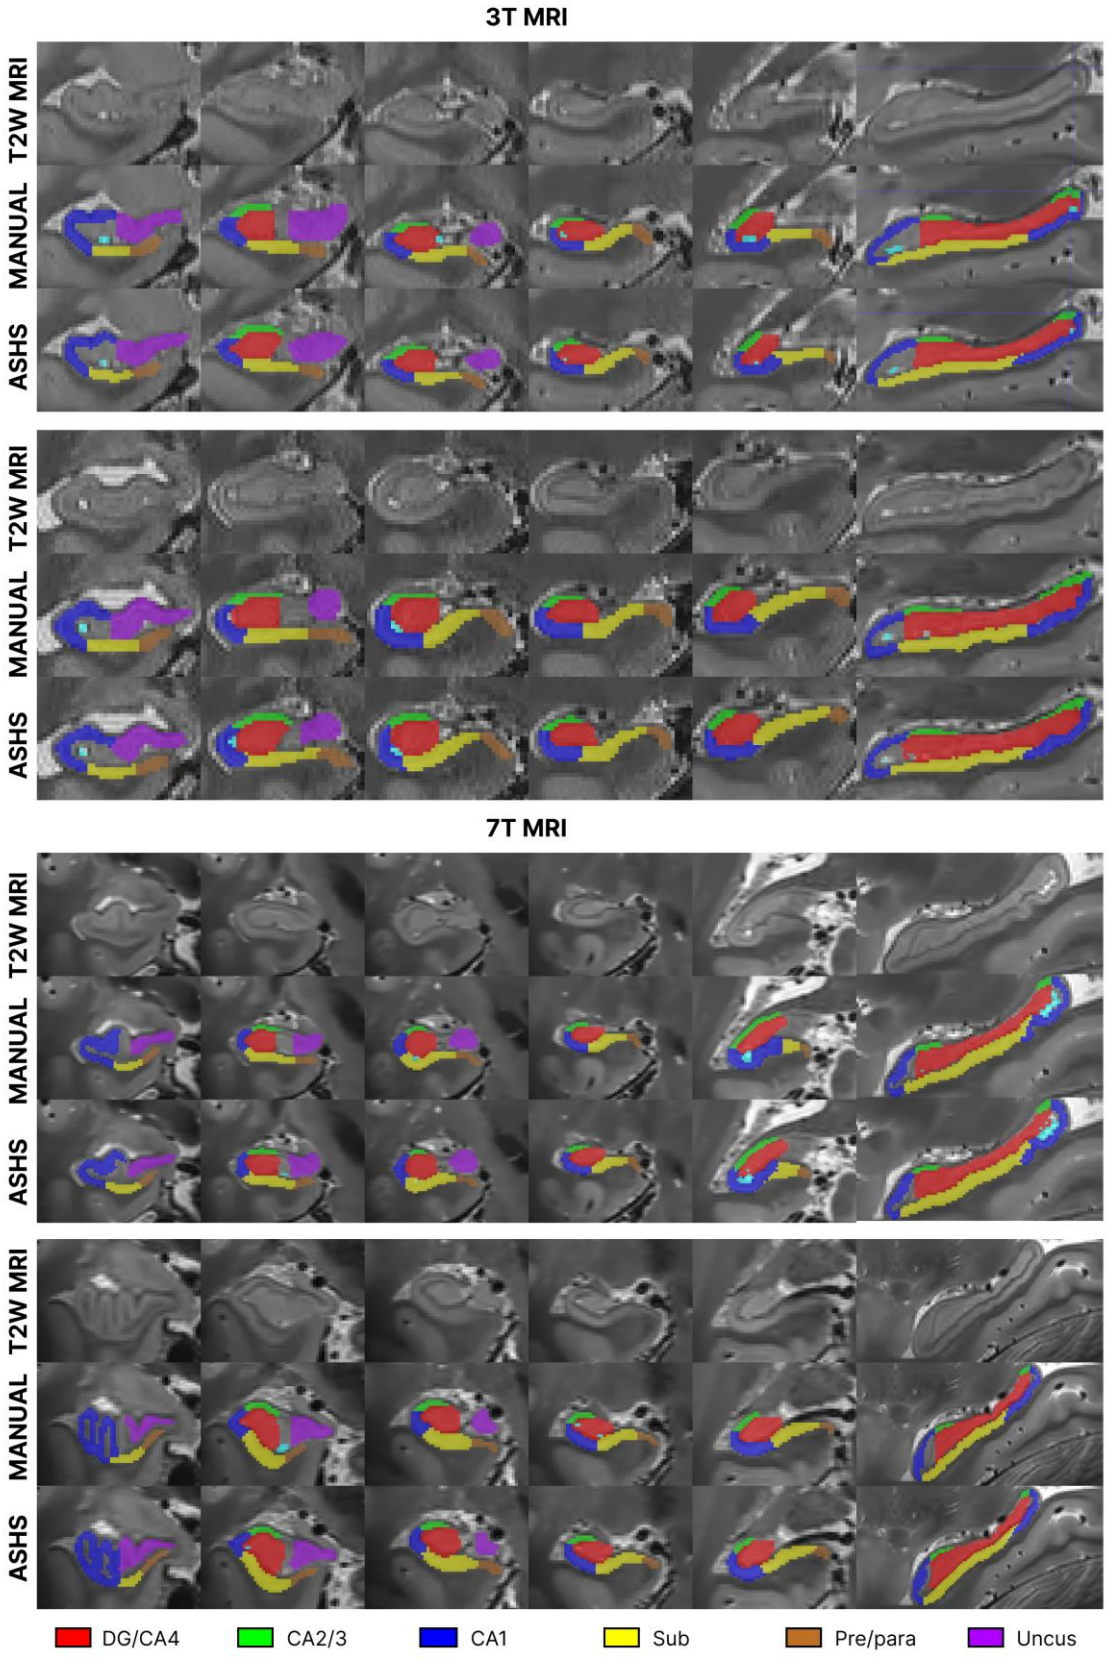

Figure 25. Example automatic and manual segmentations. Example segmentations, one from 3T MRI participants (upper panels) and 7T MRI participants (lower panels), performed using (a) the 3T ASHS atlas package to segment hippocampi from new 3T MRI data, and (b) the 7T ASHS atlas package to

segment hippocampi from new 7T MRI data. In both (a) and (b), for each participant, the top row displays the raw T2-weighted image, the second row shows a manual segmentation of the hippocampal subfields and the bottom row displays the ASHS automatic segmentation. Note that light blue areas are cysts. T2W = T2-weighted, DG = dentate gyrus, Sub = subiculum, Pre/para = pre/parasubiculum. The images shown for 3T and 7T are from approximately similar locations along the length of the hippocampus.

4.5. Discussion

I developed two protocols, one for 3T MRI and the other for 7T MRI, that permit automated hippocampus segmentation into six subregions, namely DG/CA4, CA2/3, CA1, subiculum, pre/parasubiculum, and uncus along the entire length of the hippocampus. This involved making use of pre-existing manual segmentations on 3T MRI images of 140 participants, and manually segmenting hippocampi on 7T MRI images of 40 participants. I produced 3T MRI and 7T MRI ASHS atlases using a subset of the data and tested the performance of these atlases on the remaining data. Using inter-rater reliability analyses, I showed that the quality of automated segmentations produced by these protocols was high and, on average, outperformed expert manual segmenters using the same protocol (Barry et al., 2021; Clark et al., 2023; Dalton et al., 2019). These two ASHS atlas packages can, therefore, be used to automatically segment subfields from new 3T and 7T T2-weighted structural MRI data with isotropic voxels. The atlases, along with a README file containing information about the IT requirements and a step-by-step guide for using the atlases, are open access and available on Figshare (https://doi.org/10.6084/m9.figshare.24298891).

In addition to providing superior segmentation, these automated protocols significantly reduce the time and expertise required to analyse high resolution MRI data. High resolution isotropic MRI data take a particularly long time to manually segment due to the large number of coronal slices, which is considerably more than in anisotropic images used for subfield segmentation. Training experimenters in detailed subfield segmentation protocols, like the Dalton et al. (2017) protocol, is particularly time-consuming, typically spanning several months. These ASHS atlases make automatic subfield segmentation feasible for everyone, including small, inexperienced research teams, or for the analysis of large datasets.

This is the first automated segmentation process using the Dalton et al. (2017) segmentation protocol. The segmentation protocol that we used is true to our current understanding of

hippocampal anatomy. As other automated methods used different segmentation protocols, subfield Dice scores cannot be directly compared. Most other protocols combine the subiculum and the pre/parasubiculum into one larger subregion (Hadar et al., 2018; Wisse et al., 2016; Yushkevich, Pluta, et al., 2015), despite these areas having distinct functional roles (Maass et al., 2014; Zeidman, Lutti, et al., 2015) and patterns of functional connectivity (Dalton et al., 2019). Many protocols delineate the uncus into subfields (Iglesias et al., 2015; Kulaga-Yoskovitz et al., 2015). However, given that the uncal subfields have unique cellular properties and patterns of connectivity, I believe that it is appropriate to segment the uncus as a separate subregion (Dalton et al., 2017). Finally, some protocols do not segment the tail of the hippocampus at all (Iglesias et al., 2015; Wisse et al., 2016), despite it having many important functions (Lee et al., 2020).

A second issue with the extant automated processes is the non-isotropic MRI data that they have been trained on (e.g., DeKraker et al., 2022; Yushkevich, Pluta, et al., 2015). To our knowledge only one protocol exists that is optimised for 7T isotropic MRI data (Wisse et al., 2016). It is important for an automated algorithm to be trained with (and indeed be used to segment) data with isotropic voxels. Such data have equivalent resolution in all orientations, which is particularly important for convoluted structures like the hippocampus, which can contain undulations in the subfields. Time of acquisition is often cited as a reason not to acquire isotropic data. However, T2-weighted isotropic data can be acquired with ever-increasing speed as demonstrated here (~10 minutes for 7T and ~13 minutes for 3T).

There are several caveats that should be borne in mind in relation to my ASHS atlases. They were built using the MRI scans of healthy young adults, and so would be best used to segment subfields in scans from similar people. As many participants in cognitive neuroscience experiments are from the age groups scanned here, I hope the atlases will be useful. ASHS also requires T2-weighted scans, and indeed such scans are, in our experience, best suited to delineating the hippocampal subfields. As already noted, I also strongly advocate having scans with isotropic voxels particularly for an undulating structure like the hippocampus. Atlas users should therefore acquire data with isotropic voxels in order to achieve maximum benefit. A final point to note is that my atlases focus solely on the hippocampus. For researchers interested in adjacent cortical areas such as the entorhinal or perirhinal cortices, the large

open access Clark and Maguire (2023) dataset and the editable ASHS software provide the opportunity to segment and build new atlases that include these regions in the future.

In summary, here I demonstrated that the whole hippocampus can be automatically segmented from 3T and 7T MRI scans with isotropic voxels into 6 subregions according to a detailed subfield segmentation protocol. I provide free, open access to the ASHS atlas packages I developed that makes this possible. I showed that the quality of automated segmentations produced by these ASHS atlas packages is high when assessed by inter-rater reliability analyses, which is the method typically used for checking the quality of manual segmentations. I anticipate that these atlases will save researchers a significant amount of time, especially when conducting subfield experiments at scale. Moreover, I hope that the ability to automatically segment subfields will allow more research groups to conduct subfield experiments, which will be essential for gaining a full understanding of how the hippocampus functions.

5. Experiment 3: The neural microcircuitry underpinning retrieval of remote and recent autobiographical memories

5.1. Introduction

Autobiographical memories are the ghosts of our past. Through them we visit places long departed, see faces once familiar and hear voices now silent. These, often decades-old, personal experiences locate us on life's timeline and underwrite our capacity for autonomy.

Previous neuroimaging studies have revealed the brain regions supporting this ability, which include the hippocampus and the medial temporal, lateral temporal, medial parietal, medial prefrontal, lateral prefrontal and occipital cortices. However, it is unclear how retrieval of remote and recent autobiographical memories is supported by cortical and subcortical microcircuits.

Having developed the 7T MRI sequences and processing pipeline that I described in the previous chapters, I used these to investigate, for the first time in humans, representations of individual remote and recent autobiographical memories in the cortical layers and hippocampal subfields. Participants vividly recalled a set of remote (2-5 year old) and recent (1-6 week old) autobiographical memories 6 times each. Using RSA, I detected representations of the individual remote memories in the deep cortical layers of two brain areas, the mPFC and the middle gyrus of the lateral temporal cortex (mLTC). Whereas individual recent memory representations were detected in the middle layers of the mLTC. I then sought to understand whether or not there was a relationship between these memory representations and representations in any other areas in the autobiographical memory network. To do this, I correlated the trial-by-trial representation strengths using a method known as informational connectivity. For remote autobiographical memories, I found that memory-specific information may be shared by brain areas using feedback signalling pathways, which are known to travel in the deep cortical layers. By contrast, recent memory representations seemed to involve feedforward signalling pathways. Overall, these findings provide a window

into the microcircuitry supporting the memory system as it works to remember experiences from the distant and recent past.

5.2. Background

The initial encoding of autobiographical events into memory is thought to involve the hippocampus (Dickerson & Eichenbaum, 2010; Spiers et al., 2001; Vargha-Khadem et al., 1997). Over time, autobiographical memories are then consolidated into neocortical areas, by means of interactions between the hippocampus and neocortex. This process is known as systems-level consolidation (Barry & Maguire, 2019; Frey & Morris, 1998; Marr, 1971; Redondo & Morris, 2011). The mPFC is one likely site for the consolidation of remote autobiographical memories, with fMRI studies showing that remote autobiographical memories are more strongly represented in the mPFC than recent memories (Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici & Maguire, 2018). Interestingly, despite remote and recent memory traces being stored in different brain areas, when they are retrieved from memory, the same core network of brain areas is active (Buckner & Carroll, 2007; Hassabis & Maguire, 2007; Maguire, 2001; Schacter et al., 2007; Svoboda et al., 2006). This includes the hippocampus, the medial prefrontal, lateral prefrontal, retrosplenial, lateral temporal, parahippocampal and occipital cortices. Although, there is some debate over the degree of involvement of these areas, particularly the hippocampus, after memories have been consolidated in the neocortex (Barry & Maguire, 2019; Moscovitch et al., 2016; Nadel et al., 2007; Squire et al., 2015).

The mPFC plays a critical role in autobiographical memory retrieval. Damage to this area results in severe impairments in memory recall ability. For example, patients with mPFC damage are unable to spontaneously recall autobiographical memories without heavy cueing (Kurczek et al., 2015). This led to the proposition that the mPFC initiates and coordinates the retrieval of autobiographical memories; a proposition that is supported by neuroimaging research. For example, MEG work found that the mPFC responds earlier than the hippocampus during retrieval of remote but not recent autobiographical memories (McCormick et al., 2020). Furthermore, DCM analysis of these data showed that the mPFC drove activity in the hippocampus, rather than vice versa. Remote memories may have a more

requirement for this top-down control, perhaps because the memory traces are less available in the hippocampus (Barry & Maguire, 2019; McCormick, Ciaramelli, et al., 2018).

There are no direct anatomical pathways from the mPFC to the hippocampus, but there are several indirect pathways via other cortical areas in the core network. Long range feedback connections from higher order cortical areas, such as the mPFC, to lower order cortical areas, such as the retrosplenial cortex and parahippocampal cortex, travel in the deep cortical layers (Felleman & Van Essen, 1991). This differs from feedforward pathways, which travel in the middle and superficial layers. The deep layer feedback pathways seem like a plausible pathway for top-down control by the mPFC in support of autobiographical memory retrieval.

The principal aim of this experiment was to leverage the high spatial resolution of 7T fMRI to investigate remote and recent autobiographical memory retrieval in the mPFC layers, the hippocampal subfields, and the layers of other cortical areas within the core network. A second goal was to understand the functional connectivity of the cortical layers and hippocampal subfields during remote and recent autobiographical memory retrieval. As far as I am aware, a laminar investigation has never been employed before when studying autobiographical memory recall. The two other 7T fMRI studies that I am aware of in the literature that have examined autobiographical memory recall did not interrogate cortical laminar responses, and so could not speak to my research questions of interest (Leelaarporn et al., 2024; Pfaffenrot et al., 2024).

Specifically, I asked:

1. Which brain areas were active during autobiographical memory retrieval?
2. Were there any differences in activation of the hippocampal subfields between remote and recent autobiographical memory retrieval? Note, it was not possible to investigate univariate activation differences in the cortical layers due to the existence of a superficial layer bias effect (see Chapter 3).
3. Were specific remote and/or recent memories represented in any of the mPFC layers or hippocampal subfields, or the layers of other cortical areas in the core network?
4. What was the nature of the connectivity between the mPFC layers, hippocampal subfields and other areas in the core network during autobiographical memory retrieval?

To address these research questions, participants underwent 7T fMRI scanning while performing an autobiographical memory retrieval task. During the task, participants were presented with a memory verbal cue followed by a retrieval phase (10 s) during which they vividly re-experienced the memory in their mind's eye. Two types of memory were recalled: remote memories (2 to 5 years old) and recent memories (less than one month old). To address question 1 above, I used mass whole brain (or partial volume in this case) univariate analysis to identify the brain areas that were active during memory retrieval. To address question 2, I used univariate analyses to compare the activity during remote and recent memory retrieval in each of the hippocampal subfields. To address question 3, I used RSA to test whether individual memories were represented stably in any of the mPFC layers, hippocampal subfields or the layers of other cortical areas. Finally, to address question 4, I used informational connectivity analysis to understand whether any cortical layers or hippocampal subfields with stable memory representations had correlated representation time courses with other areas, indicating that memory-related information might be shared between these brain areas.

I hypothesised that:

1. The core network of brain areas would be activated during autobiographical memory retrieval, including the mPFC, lateral prefrontal, lateral temporal, parahippocampal, retrosplenial, occipital cortices, and the hippocampus, which would corroborate countless previous research (Buckner & Carroll, 2007; Hassabis & Maguire, 2007; Maguire, 2001; Schacter et al., 2007; Svoboda et al., 2006).
2. There would be no difference in activation of the hippocampal subfields during remote and recent autobiographical memory retrieval. This hypothesis was based on a few previous fMRI studies that have found no effect of memory age on univariate activation of the whole hippocampus (Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici & Maguire, 2018). However, note that this question has not previously been examined in the separate subfields of the hippocampus.
3. Individual memories would be represented in the mPFC deep layers (the layers in which feedback pathways originate), and more strongly when the memories were more remote. Remote memory representations may also be detected in the deep layers of lower order cortical areas in the core network, such as the lateral temporal,

parahippocampal, retrosplenial and occipital cortices. This hypothesis is based on the proposition that the mPFC initiates and coordinates memory retrieval, via feedback signalling, and that remote memories may have a greater reliance on this process.

4. Memory-specific information would be shared between the mPFC deep layers and the deep layers of lower order cortical areas (e.g., the lateral temporal cortex, parahippocampal cortex, retrosplenial cortex and/or occipital cortex) and/or the hippocampus, again based on the idea that the mPFC controls the retrieval process.

This experiment could provide, for the first time in humans, an understanding of autobiographical memory retrieval at the level of cortical layers, allowing autobiographical memory retrieval to be characterised in terms of feedforward or feedback processing. These data also speak to systems level consolidation theories, offering a more mechanistic understanding, as the layer specificity provides some indication of the microcircuitry involved.

5.3. Methods

5.3.1 Participants

Twenty-four healthy participants were recruited and scanned. Two participants were excluded from the analysis due to considerable head motion (> 1.6 mm movement) during fMRI data acquisition. Another participant was excluded due to an artefact in the fMRI data, and a final participant for their inability to vividly recall memories when performing the fMRI task. The remaining 20 participants were 14 females and 6 males aged between 18 and 35 years old (mean = 22.75, SD = 3.16) with no history of neurological or psychiatric conditions.

5.3.2. Experimental procedure

Each participant paid two visits to the Department. During the first visit they underwent structural MRI scanning, followed by an autobiographical interview (Levine et al., 2002), an established procedure for sampling and selecting autobiographical memories that are vivid, easy to remember and distinct from one another. Seven to 19 days later (mean = 9.65, SD = 3.82), they returned for a second visit where they recalled the autobiographical memories during 7T fMRI scanning.

5.3.2.1. Autobiographical interview

During the first visit, a participant was asked to describe out loud, in as much detail as they could remember, autobiographical experiences from two time periods in their life: 2-5 years ago (remote memories) and less than one month ago (recent memories). An event needed to be retrievable as a story, with a beginning, middle and end. If a participant provided insufficient detail, the interviewer would give general probes (e.g., "Can you remember anything else from the event?"). They were asked to choose a section of the event that they could later recall within 10 seconds during the fMRI task. A one- or two-word phrase was agreed with the interviewer that would remind the participant of the specific section of the memory to focus on during the subsequent fMRI scan. During this session, autobiographical memories were scored approximately for Internal (episodic) and External (non-episodic) details according to an adapted version of the widely-used Levine Autobiographical Interview scoring method (Levine et al., 2002). The interview was terminated once a participant had provided 8 remote and 8 recent memories that met the criteria for detail (number of “internal” episodic details > 5; see Behavioural analysis section below for explanation of internal episodic details). To maximise the differentiability of the memories, I ensured that, for each participant, the memories were distinct from one another in terms of time, place, people and activities. Furthermore, I ensured that all memories were temporally specific and unique. For each event, participants were asked if they had experienced anything like the event before or since. Events that were similar in content to other past events were excluded. The autobiographical interview was audio recorded and transcribed to facilitate further objective scoring of the memory descriptions.

5.3.2.2. Autobiographical memory task during fMRI

Following a training session where they practiced recalling 4 memories that were not used in the fMRI experiment (2 remote and 2 recent), a participant underwent fMRI scanning while recalling the remaining 12 memories (6 remote and 6 recent). During each trial (Figure 26), a participant was shown a memory cue for 3 seconds, was then presented with a visual cue indicating they should close their eyes for 1 second, and then recalled the memory for 10 seconds. An auditory tone (1 second) then cued them to open their eyes, followed by visual presentation of a question ("How vivid was the memory?") for which they could answer on a scale from 1 to 4 (1 = "not at all vivid", 2 = “not very vivid”, 3 = “very vivid”, 4 = "extremely

vivid"). On average, both remote (mean = 3.16, SD = 0.301) and recent memories (mean = 3.32, SD = 0.364) were rated as very vivid to extremely vivid with ratings higher than 3 out of 4 in both cases.

Counting trials were used as a baseline condition in the mass univariate fMRI analysis and a control condition in the RSA analysis. Counting trials were the same as memory trials except that a 2-digit number was presented instead of a memory cue. A participant was tasked with counting up in 3's silently in their head from that number for 10 seconds, and then they rated their focus during the rating section of the trial. Counting trial timings were matched to the memory trials. Participants on average rated their focus as 3.22 (SD = 0.424) out of 4 on counting trials. Each memory or number cue was presented once per run in a random order (both within and across runs), yielding 24 trials per run. There were 6 runs meaning that each memory and number was presented a total of 6 times, yielding a total of 144 trials. The fMRI task was produced using the MATLAB 2021b toolbox, Psychtoolbox-3 or Psychophysics Toolbox Version 3 (PTB-3).

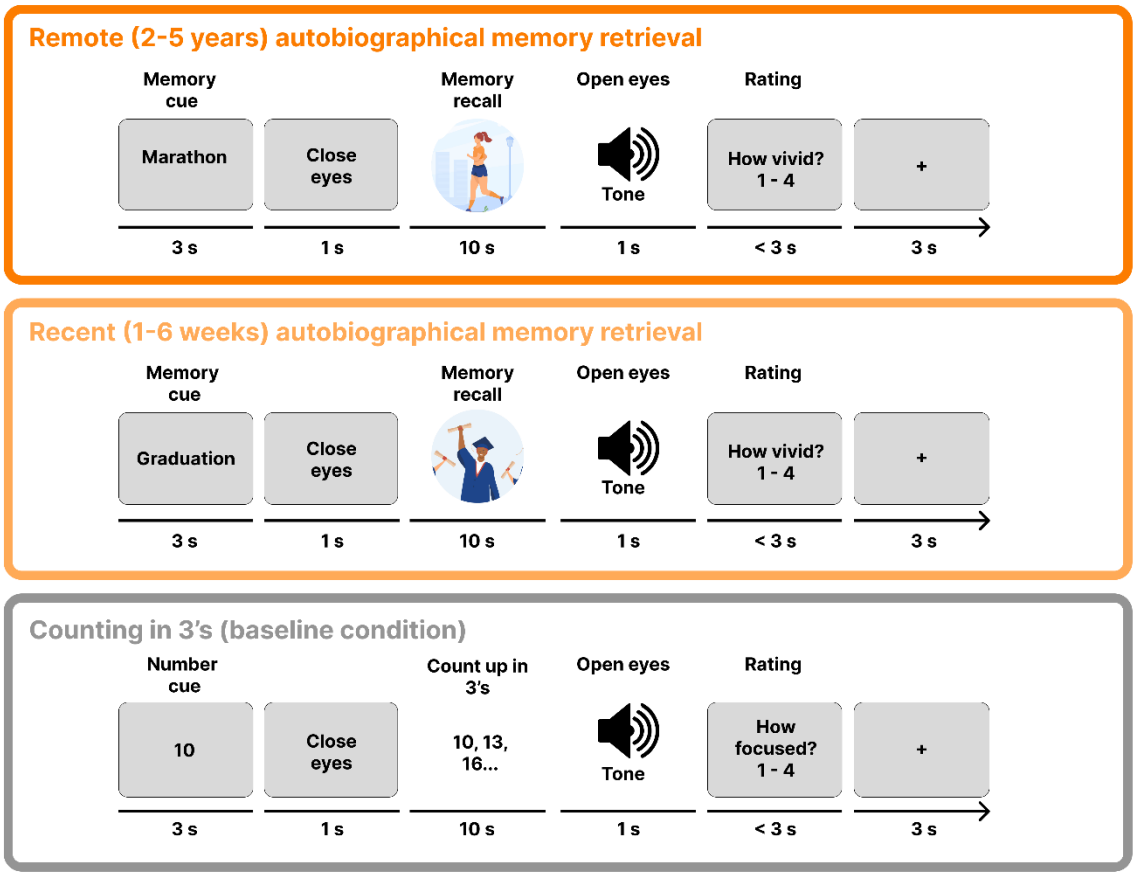

Figure 26. Autobiographical memory task performed during 7T fMRI. Remote (dark orange box) and

recent (light orange box) autobiographical memory retrieval trials consisted of a memory cue (3 s), then a cue to close eyes (1 s), followed by a period of vivid memory retrieval (10 s), after which there was an auditory tone to cue opening of the eyes (1 s), followed by a vividness rating (up to 3 s) and finally a fixation cross (3 s). Counting baseline trials (grey box) consisted of a number cue (3 s), followed by a cue to close eyes (1 s), then a period of counting (10 s), after which there was an auditory cue to open eyes (1 s), then a focus rating (up to 3 s) and finally a fixation cross.

5.3.3. Behavioural analyses

The memory descriptions provided in the autobiographical interview were objectively analysed using the autobiographical interview method (Levine et al., 2002). Each memory detail was categorised as belonging to one of the "internal" or "external" detail categories. Internal detail categories included specific events, places, perceptual observations and thoughts or emotions. External detail categories included unrelated events, semantic knowledge, repetition of details, or other more general statements. Because participants were using their calendars to remind them of events, I did not include temporal details as an internal details category. For each memory, the number of memory details in each category were summed to give a score per category. To check the reliability of memory scoring, I performed an interrater reliability analysis. A subset of 48 randomly-selected memories were scored by another experimenter. Intraclass coefficient estimates were calculated using SPSS statistical package version 22 (SPSS, Chicago, IL) based on a single-measure, absolute-agreement, two-way random-effects model (see Table 6 for interrater reliability of scores). For each category, the scores were averaged across memories of the same time period, remote and recent. Paired Wilcoxon signed rank tests were used to determine if there were differences in the number of details remembered for remote compared to recent memories.

5.3.4. MRI sequences

MRI data were acquired on a Siemens Magnetom Terra 7T MRI system (see Chapter 2, Section 2.6.1).

5.3.4.1. Functional MRI

Partial brain, submillimetre resolution (voxel size = 0.8 × 0.8 × 0.8 mm^3) fMRI images were acquired using a T2*-weighted 3D GE-EPI sequence, with a volume acquisition time of 3872 ms, TR = 44 ms, TE = 18.7 ms, flip angle = 14°, field of view = 192 × 192 × 70.4 mm^3, slab angle = 15°, direction of phase encoding = anterior-to-posterior, parallel imaging acceleration in

both the direction of the partition (factor = 2) and the direction of phase encoding (factor = 4), partial Fourier 6/8 in the phase-encoded direction of the EPI readout, and a binomial (1331) water-selective excitation. Four images, with the same parameters but the opposite (posterior-to-anterior) phase encoding direction, were acquired at the start of the sequence. Advanced B0 shimming (WIP 1441) was used.

5.3.4.2. Anatomical MRI

Four anatomical images were acquired. One whole brain MT-weighted EPI image was acquired per participant using a T2*-weighted 3D GE-EPI sequence with many of the same parameters used in the fMRI sequence. Some parameters were modified: MT- based contrast = on, volume acquisition time = 3 m 45 s, TE = 16.97 ms, flip angle = 8°, field of view = 192 × 192 × 128 mm^3, parallel imaging (factor = 4). MT-weighted EPI images were matched to the fMRI images in terms of geometrical distortion (due to use of the same parameters e.g., echo spacing and acceleration factor) but they had better contrast between grey and white matter (due to the MT weighting). Therefore, they were useful for coregistering cortical surfaces to the fMRI data.

One whole brain, submillimetre resolution (voxel size = 0.65 × 0.65 × 0.65 mm^3) T1 weighted image was acquired using an MP2RAGE sequence (Marques et al., 2010) with a volume acquisition time = 8 m 42 s, TR = 5000 ms, TE = 2.54 ms, TI = 900 ms and 2750 ms, 5° and 3° flip angles, field of view 208 × 208 × 156 mm^3, and in-plane GRAPPA acceleration (factor = 3). This image was used to define the boundary between grey matter and CSF, and the boundary between white matter and grey matter.

Two partial volume, submillimetre resolution (voxel size = 0.52 × 0.52 × 0.5 mm^3) T2 weighted images were acquired using a T2-weighted 3D SPACE sequence (Mugler et al., 2000) with a volume acquisition time = 10 m 25 s, TR = 3500 ms, TE = 229 ms, flip angle = 8°, field of view = 200 × 169 mm x 56 mm, partition oversampling = 14.3%, and in-plane GRAPPA acceleration (factor = 2). The two images were first coregistered using SPM Realign. Then they were denoised and averaged, using the Structural Averaging Toolbox, prior to being used for manual segmentation of the hippocampal subfields.

5.3.5. Preprocessing

5.3.5.1. Defining cortical surfaces

For each participant, the MP2RAGE image was denoised, using the Structural Averaging Toolbox, before using FreeSurfer version 7.3.2 to define the pial surface, the boundary between grey matter and CSF, and the white matter surface, the boundary between white matter and grey matter. To improve the definition of the pial surface, FreeSurfer has the option to additionally use a whole brain T2-weighted image. As my T2-weighted images were partial brain, I produced and additionally used a "splice" whole brain T2-weighted image by coregistering, denoising and averaging the MP2RAGE T2-weighted image with the two high resolution partial brain T2-weighted images. The surfaces were visually inspected and, where appropriate, manual edits were made.

5.3.5.2. Motion and distortion correcting fMRI data

Using FSL version 6.0, the susceptibility-induced off-resonance field was estimated from the four opposite phase encoded (i.e., posterior-to-anterior) images and the first four (i.e., anterior-to-posterior) images from the fMRI time series, which contained distortions going in opposite directions (Andersson et al., 2003; Smith et al., 2004). I then used SPM12 software, implemented in MATLAB, to calculate a voxel displacement map from the estimated B0 field map. The voxel displacement map could then be used to undistort the fMRI data. The SPM Realign and Unwarp function was used to estimate rigid body realignment and unwarping (Andersson et al., 2001). SPM Realign and Unwarp also corrects for susceptibility-motion interactions by estimating how the B0 field changes with participant movement. The fMRI data were resampled, correcting for motion and distortion in one step and a mean fMRI image was calculated. No further resampling of the fMRI data was performed to minimise smoothing of the image due to interpolation.

As noted previously, the data from two participants were excluded from further analyses as their head moved more than 2 voxels (1.6 mm) over the course of the scanning session, one participant moved 4.39 mm and the other moved 5.78 mm. The remaining 20 participants did not move their head more than 1.35 mm (mean = 0.65 mm, SD = 0.30).

The whole brain MT-weighted EPI image was also unwarped using FSL, using one image that had opposite phase encoding (i.e., posterior-to-anterior) to the fMRI time series and the other

that had the same phase encoding (i.e., anterior-to-posterior; Andersson et al., 2003; Smith et al., 2004). It was then coregistered to the mean fMRI image using SPM. Due to its high grey matter to white matter contrast, this image was used for coregistration of the cortical surfaces to the fMRI data.

5.3.5.3. Aligning cortical surfaces to fMRI data

The pial and white matter surfaces were coregistered to the fMRI data. The coregistration was estimated in three steps. First, a volume-based registration of the MP2RAGE T1 image to the MT-weighted EPI image was estimated by FSL Flirt using NMI. Then, BBR of the cortical surfaces to the high contrast fMRI image was estimated using FreeSurfer BBRegister, which works by maximising the signal intensity gradients across the pial and white matter surfaces (Greve & Fischl, 2009). Finally, recursive BBR was used to coregister the cortical surfaces to the high contrast fMRI image using the OpenFmriAnalysis toolbox. This works by dividing the cortical surface mesh into increasingly smaller subsections and performing BBR iteratively on each section of mesh (van Mourik, Koopmans, et al., 2019). Six iterations of BBR were performed. To correct for any distortion remaining in the fMRI data, in each iteration, scaling was allowed along the phase-encoding direction, because in GE-BOLD EPI, geometrical distortion is most severe in the phase-encoded direction (Jezzard & Clare, 1999).

5.3.6. Whole brain mass univariate analysis

The fMRI data were normalised to MNI space and then smoothed (FWHM = 6 mm) using SPM12. The GLM analysis was performed in SPM12. The model included three regressors of interest which were the 10 s recall periods of remote memories, the 10 s recall periods of recent memories, and the 10 s periods of counting. Despite memory recall and counting trails being rated, on average, as very vivid and very focused, respectively, I decided to use vividness and focus as additional regressors of no interest in the GLM to account for any potential effects. The six motion parameters were also modelled. For each participant, a t-statistic map of memory recall was produced by contrasting remote and recent beta estimates with counting beta estimates. Group level analysis was then performed in SPM12 to obtain a t-statistic map for each contrast.

5.3.7. Delineating regions of interest

5.3.7.1. Cortical layer ROIs

As well as defining the cortical surfaces, FreeSurfer segments the brain into cortical and subcortical areas. The brain areas that corresponded with the areas activated by memory recall, as determined by the brain area level analysis, were selected. For some brain areas, the FreeSurfer masks were combined to create a mask that encompassed the full area of activation. These brain area masks were coregistered to the fMRI data using BBR and then translated from surface space into volume space by projecting them across the surface normal, which is a 3D vector that runs perpendicular to the surface. Gaps in the ROIs were filled in using FSL Maths Edge, Dilate and Erode operations. Small masks were combined to give 7 cortical brain areas. Where appropriate, manual edits were made to the cortical area ROIs.

Within each cortical area ROI, I divided the cortical grey matter into 3 equivolume layers (deep, middle and superficial) using the OpenFmriAnalysis toolbox. Two additional surfaces were produced between the pial and white matter surfaces using the level set method (Waehnert et al., 2014). The distance of each ROI voxel to 5 compartments (white matter, CSF, and the superior, middle and deep cortical layers) was calculated, once again using the level method. From this, the distribution of each voxel's volume over the 5 compartments was calculated. Voxels that were 80% or more within a given layer were included in the layer ROI.

GE-BOLD sequences are susceptible to an effect known as the superficial layer bias effect or venous effect (Kay et al., 2019). Because there are more large veins in the superficial layers, and blood drains from the deep to superficial layers (Duvernoy et al., 1981), the measured BOLD signal is highest in the superficial layers. To address this, voxels with a tSNR lower than the 65th percentile of the tSNR distribution were removed from the cortical layer ROIs as they were expected to contain large veins (Jia et al., 2021). To remove voxels in areas with high distortion from the analysis, a Jacobian map was estimated from the fMRI data using the HySCO tool in the ACID MATLAB toolbox (Ruthotto et al., 2012). The Jacobian is a measure of the displacement of the signal contained within a voxel and captures the degree of compression or stretching that has occurred (Jezzard & Balaban, 1995). It can therefore be used as a quantitative measure of distortion (Clark et al., 2021). Voxels with Jacobian values

lower than the 5th percentile and higher than the 95th percentile of the Jacobian distribution were considered to be highly distorted and were therefore removed from the cortical layer ROIs.

5.3.7.2. Hippocampal subfield ROIs

The two partial brain T2-weighted images were denoised and averaged. ITK-SNAP version 3.2.0 was used to manually delineate the hippocampal subfields, according to an established segmentation protocol, on the resultant T2-weighted image (Dalton et al., 2017). The whole hippocampus, bilaterally, was manually delineated into 6 subregions: DG/CA4, CA3/2, CA1, subiculum, a region which combined the presubiculum and the parasubiculum known hereafter as the pre/parasubiculum, and finally the uncus. To assess the reliability of the manual segmentations, I used inter-rater reliability analysis. Twenty five percent of hippocampi (5 of 20 participants) were additionally segmented by a separate experimenter and the similarity between experimenter segmentations was assessed using the Dice similarity coefficient, which ranges from 0 (no overlap) to 1 (complete overlap; see Table 9 for the interrater reliability results).

5.3.8. Univariate hippocampal subfield analyses

A single BOLD signal time course was obtained for each subfield by averaging the BOLD signal across all voxels in the subfield. A temporal GLM was then applied to the subfield time courses using the same model described in Section 5.3.6. However, this time, instead of applying the GLM and statistical tests per voxel, they were applied per subfield. In each subfield, two condition contrasts were calculated using a condition subtraction: remote memory retrieval minus the counting baseline, and recent memory retrieval minus the counting baseline. Differences between these two contrasts were then tested using Bonferroni corrected paired Wilcoxon signed rank tests, thereby testing for a difference in activation during remote and recent memory retrieval.

5.3.9. Representational similarity analysis

Each trial was modelled with a separate GLM with one regressor of interest: the 10 second retrieval or counting period of that trial. For memory trials, the 10 second retrieval periods of all other memory trials were included as regressors of no interest; while for counting trials, the 10 second retrieval periods of all other counting trials were included as regressors of no

interest. This protocol attempted to isolate the activity relating to each specific memory recall or number trial. Vividness ratings, focus ratings and motion parameters were included as regressors of no interest. To limit the contribution of noisy voxels to the multivariate activation patterns, I used multivariate noise normalisation (Walther et al., 2016). In multivariate noise normalisation, the beta estimates for each voxel are divided by the square root of the GLM residuals, thereby down-weighting noisier voxels.

I performed RSA using the RSA toolbox, following a procedure that was used in a previous study on individual memory representations (Barry et al., 2018). For each memory trial, the similarity of activity patterns across trials of the same memory was calculated by averaging the Fisher Z-transformed Pearson product-moment correlation coefficient to give a "within-memory" similarity measure. Next, the average correlation of the trial with trials of all other memories was calculated to give a "between-memory" similarity measure. Both within- and between- memory similarity measures were calculated by comparing the trial to trials in different runs. To estimate the memory-specific representation strength on the recall trial, the between-memory similarity measure was subtracted from the within-memory similarity measure. These calculations were repeated for counting trials, first comparing the trial to other trials of the same number and then to trials of other numbers. The trial-by-trial memory- and number- specific representation strengths were then averaged across trials of the same condition. At the group level, a mean representation strength that was significantly higher than 0, evaluated using a one-sample Wilcoxon signed rank test, would suggest that on average across participants and recall trials, the activity pattern for the memory or number was distinct from activity patterns of all other memories or numbers. Bonferroni-corrected Wilcoxon signed rank tests were used to test for stronger representations of remote memories compared recent memories and numbers.

5.3.10. Informational connectivity analysis

For the ROIs that contained a difference in representation strength between conditions, I next sought to understand if there was a relationship between these representations and representations in other ROIs. To examine this, I used a method known as informational connectivity (Coutanche & Thompson-Schill, 2013). I estimated the covariation of trial-by-trial representation strengths between the ROI and each of the other ROIs by calculating the Fisher Z-transformed Pearson product-moment correlation. A positive correlation indicated that

whenever the ROI with detectable representations represented the memory or number, the other region was likely to as well. For each pair of brain regions and each experimental condition, I tested for a reliable positive correlation across participants using one-sample Wilcoxon signed rank tests.

5.4. Results

5.4.1. Behavioural results - no differences between remote and recent memories in terms of the number of overall internal and external details

There were no differences between remote and recent memories in terms of the number of overall "internal" or "external", which was assessed by analysing descriptions of the memories provided by the participants prior to the scan details (Table 5 for detail scores by subcategory; Table 6 for interrater reliability of scores). The most pertinent details here were the internal details as they reflected the episodic nature of the recalled experiences, which was my main interest.

In terms of the detail subcategories, for the external details (which were less relevant for my analysis), the number of general statements that were external to the event differed between remote and recent memories ($p = 0.0326$; Figure 27). However, this difference was very marginal – on average 0.47 details (approximately half a detail) difference between memories. This could reflect a slight increase in irrelevant statements as a participant filled the silence while they attempted to remember remote details that might have been less accessible. There were no differences between any other detail subcategories, particularly for the internal details subcategories (Figure 27; Table 5).

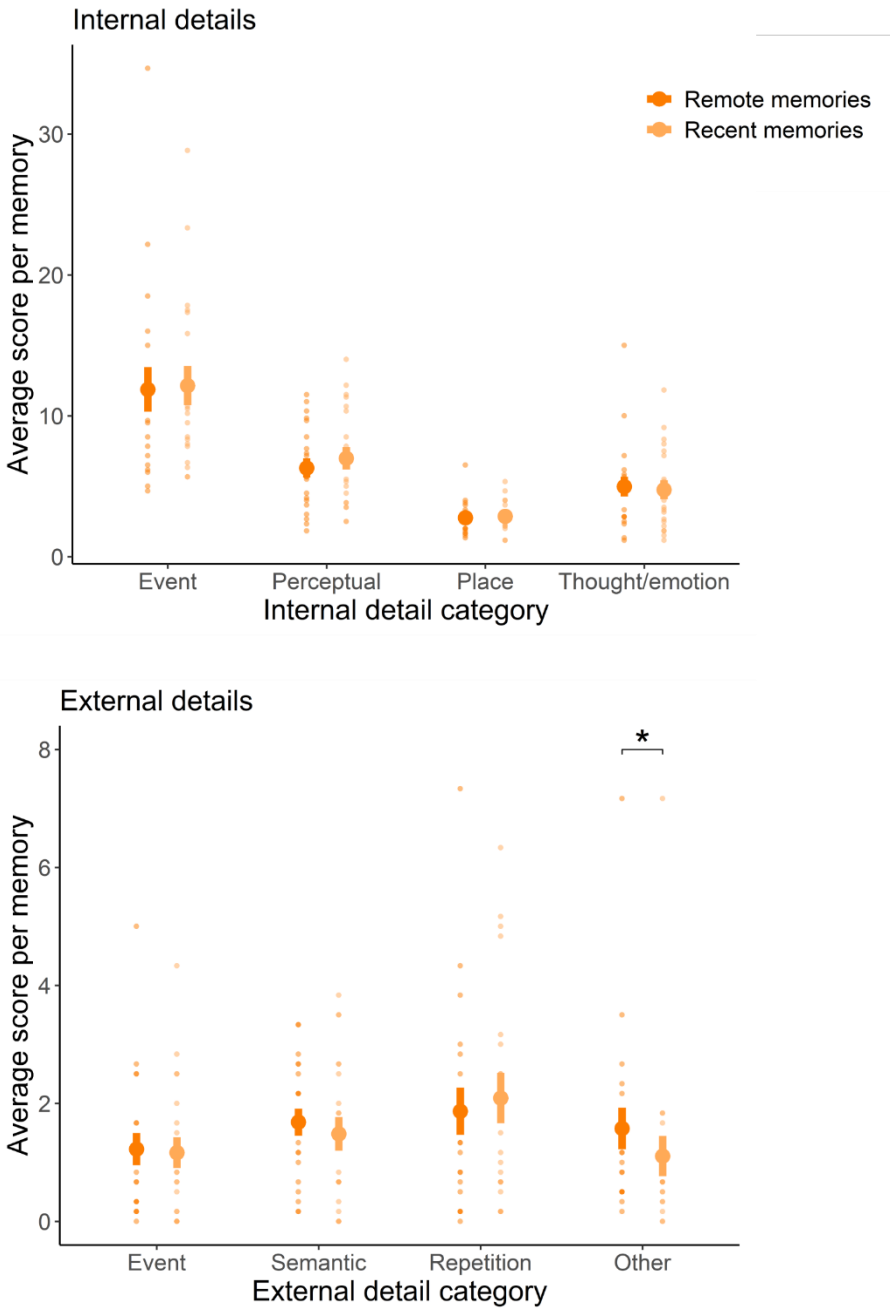

Figure 27. Comparison of remote and recent autobiographical memory detail scores. The average number of Internal (upper plot) and External (lower plot) details per memory were compared for remote and recent memories. The number of "External Other" details (i.e., general statements that were unrelated to the memory being recalled) were marginally higher for remote compared to recent memories. The heavy orange dots are the means, and the lighter dots are the individual participants' data.

Table 5. The data used to generate the graphs in Figure 27 - comparison of remote and recent autobiographical memory detail scores.

Detail category	Mean score ± SD		P value
	Remote memories	Recent memories	
Internal	25.92 ± 12.63	26.75 ± 12.56	0.467
Event	11.88 ± 7.09	12.14 ± 6.21	0.55
Perceptual	6.30 ± 3.07	6.98 ± 3.57	0.224
Place	2.77 ± 1.19	2.87 ± 1.04	0.468
Thoughts/emotions	4.98 ± 3.17	4.76 ± 3.00	0.478
External	6.35 ± 4.54	5.85 ± 4.97	0.104
External events	1.23 ± 1.22	1.17 ± 1.17	0.619
Semantic	1.68 ± 1.03	1.48 ± 1.27	0.138
Repetition	1.87 ± 1.80	2.09 ± 1.91	0.532
Other	1.58 ± 1.58	1.11 ± 1.53	**0.0326 ***

Table 6. Interrater reliability of memory detail scores. Intraclass correlation coefficients between the two researchers who independently scored a subset of 20% of the memories. Scores above 0.8 constitute very high agreement.

Detail category	Intraclass correlation
Internal	0.975
Event	0.940
Perceptual	0.866
Place	0.871
Thoughts/emotions	0.844
External	0.832
External events	0.683
Semantic	0.781
Repetition	0.767
Other	0.798

5.4.2. A core network of brain areas were activated by autobiographical memory retrieval

Next, I identified the brain areas were engaged during autobiographical memory retrieval by comparing all of the memory trials with all of the counting trials. I found significant ($p > 0.01$, uncorrected) activity in several brain areas (Figure 28a, Table 7), including the mPFC (peak

voxel at left: x = -11, y = 43, z = -7; right: x = 7, y = 52, z = -9), retrosplenial cortex (left: x = -7, y = -47, z = 6), lateral prefrontal cortex (left: x = -29, y = 30, z = -17; right: x = 35, y = 34, z = -17), mLTC (left: x = -45, y = -56, z = 17; right: x = 56, y = -4, z = -17), superior gyrus of the lateral temporal cortex (sLTC; left: x = -39, y = -6, z = -14; right = x = 22, y = 58, z = 5), parahippocampal cortex (right: x = 9, y = -41, z = 3), occipital cortex (left: x = 7, y = 52, z = -9; right: x = 15, y = -92, z = -1), and hippocampus (left: -19, -33, 10). This same set of brain areas were active when remote and recent recall trials were separately compared with counting trials (Figure 28b and 28c). These results show that ultra-high 7T resolution can reveal the autobiographical memory retrieval core network of brain areas, irrespective of memory remoteness, supporting previous 3T fMRI studies (Buckner & Carroll, 2007; Hassabis & Maguire, 2007; Maguire, 2001; Schacter et al., 2007; Svoboda et al., 2006).

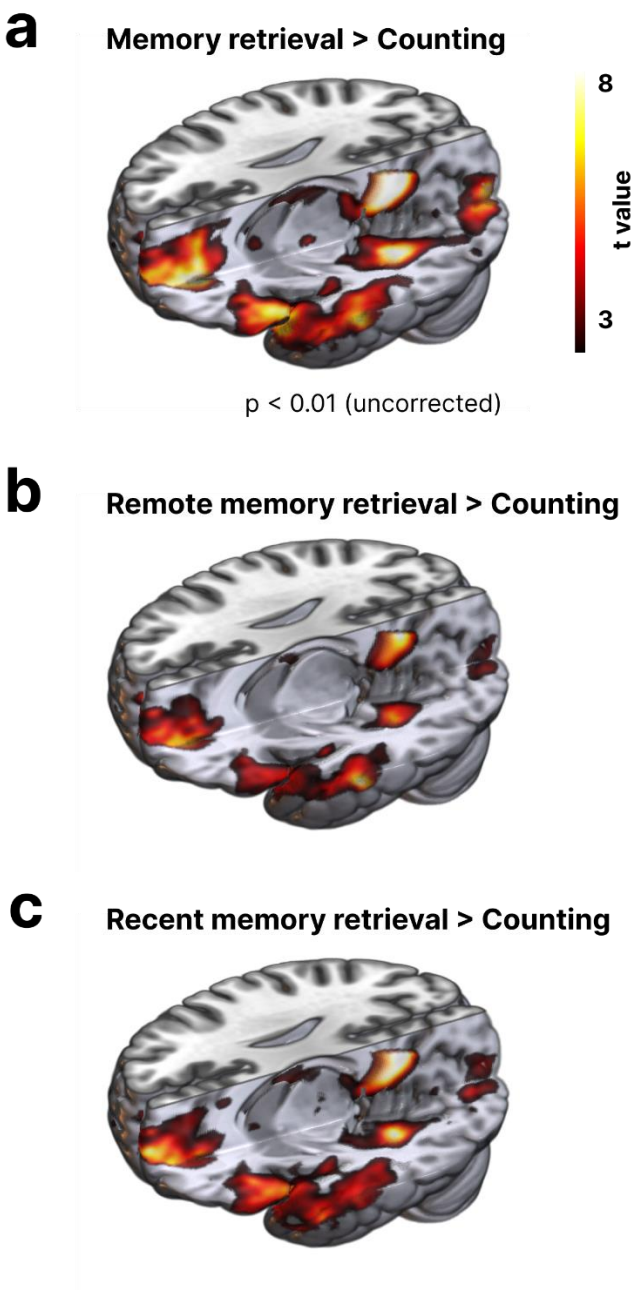

Figure 28. Core network of brain areas active during autobiographical memory retrieval. A core network of brain areas was activated during (a) autobiographical memory retrieval, (b) remote autobiographical memory retrieval and (c) recent autobiographical memory retrieval. T-scores are thresholded at $p < 0.001$ (uncorrected).

Table 7. The data used to generate the graphs in Figure 28 - activation peaks for the autobiographical memory retrieval > counting baseline contrast.

Brain region	Cluster size (voxels)	Z score	Coordinates		
			x	y	z
Retrosplenial cortex (L)	10963	6.72	-7	-47	6
Precuneus (L)		6.57	-5	-51	13
Parahippocampal cortex (R)		6.24	9	-41	3
Fusiform gyrus (L)	11390	6.57	-25	-40	-13
Pars orbitalis, lateral prefrontal cortex (L)	25036	6.00	-29	30	-17
Pars triangularis, lateral prefrontal cortex (L)		5.95	-52	18	3
Ventromedial prefrontal cortex, medial prefrontal cortex (L)	15022	5.97	-11	43	-7
Ventromedial prefrontal cortex, medial prefrontal cortex (R)		5.70	7	52	-9
Fusiform gyrus (R)	9520	5.81	29	-35	-17
Pars orbitalis, lateral prefrontal cortex (R)	4525	5.68	35	34	-17
Occipital cortex (L)	3243	5.48	-11	-98	-4
Cerebellum (L)		4.57	-18	-88	-18
Middle temporal gyrus, lateral temporal cortex (R)	3766	5.09	56	-4	-17
Occipital cortex (R)	1711	4.64	15	-92	-1
Olfactory cortex (R)	234	4.62	0	12	-2
Temporal pole, lateral temporal cortex (R)	208	4.28	52	12	-11
Insula (R)	1142	3.51	34	10	-14
Middle temporal gyrus, lateral temporal cortex (L)	701	4.12	-45	-56	17
Hippocampus (L)	336	4.11	-19	-33	10
Thalamus (L)		3.18	-10	-30	9
Superior temporal gyrus, lateral temporal cortex (L)	230	3.61	-39	-6	-14
Temporal pole, lateral temporal cortex (L)		3.53	-43	0	-18
Superior frontal gyrus, lateral prefrontal cortex (R)	91	3.46	22	58	5
Pars triangularis, lateral prefrontal cortex (R)	23	3.28	51	25	17
Superior temporal gyrus, lateral temporal cortex (R)	3	3.19	47	-19	4
Thalamus (R)	10	3.18	4	-5	10
Anterior cingulate & paracingulate gyri, medial prefrontal cortex (R)	2	3.17	7	36	7

Note: (R) refers to right hemisphere and (L) refers to left hemisphere.

5.4.3. DG and CA2/3 were more active during retrieval of recent compared to remote autobiographical memories

Next, I examined activation during remote and recent autobiographical memory retrieval in the hippocampal subfields. I divided the hippocampus into six subfields - the DG, CA2/3, CA1, subiculum, pre/parasubiculum and uncus. In each subfield, I calculated the percentage change in signal during remote and recent autobiographical memory retrieval. A signal change that

was higher than zero, assessed using a Wilcoxon signed rank test, indicated that remote or recent memory retrieval activated the hippocampal subfield. Differences in the degree of signal change between remote and recent memory retrieval were also compared using paired Wilcoxon signed rank tests.

Recent autobiographical memory retrieval activated all hippocampal subfields (Figure 29, Table 8; Table 9 for interrater reliability of scores), while remote autobiographical memory retrieval activated all subfields except hippocampal CA2/3 (Figure 29, Table 8). Furthermore, there was higher activation during recent compared to remote memory retrieval in two subfields: the DG ($p = 0.000734$, Figure 29) and CA2/3 ($p = 0.00639$, Figure 29). There were no differences in any of the other subfields (Table 8).

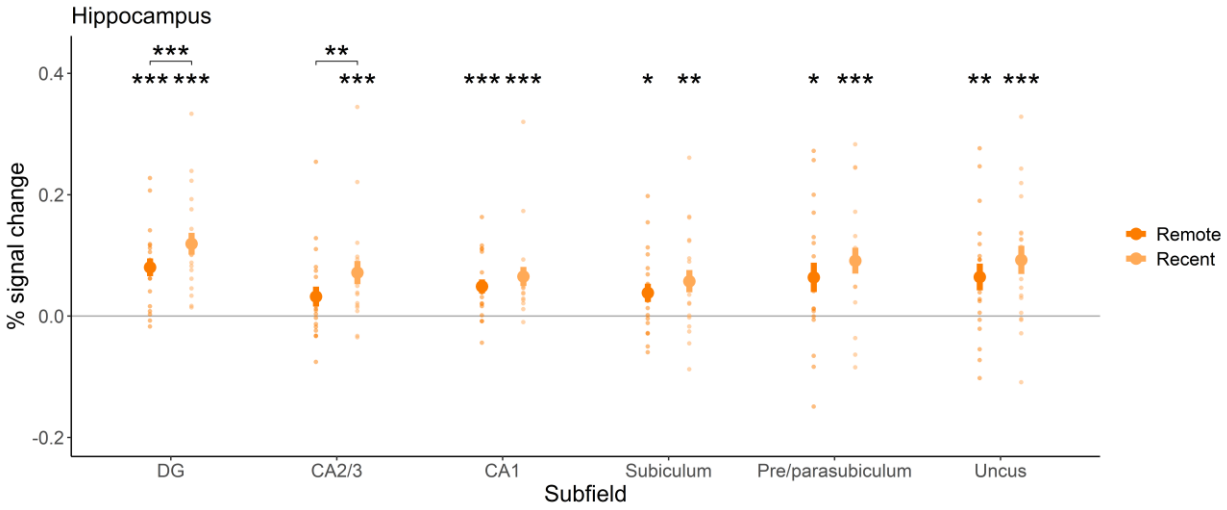

Figure 29. Hippocampal subfield activation during remote and recent autobiographical memory retrieval. The percentage signal change associated with remote and recent autobiographical memory retrieval was compared in each hippocampal subfield. The DG and CA2/3 had a higher response to recent compared to remote autobiographical memory retrieval. The heavy orange dots are the means, and the lighter dots are the individual participants' data.

Table 8. The data used to generate the graphs in Figure 29 - hippocampal subfield activation during remote and recent autobiographical memory retrieval.

ROI	Condition	% Signal change		p-value	p-value
		mean	SEM	(vs. zero)	(remote vs. recent)
DG	Remote	0.0803	0.0146	**0.0000263** ***	**0.000734** ***
	Recent	0.1191	0.0180	**0.00000249** ***	
CA2/3	Remote	0.0322	0.0163	0.114	**0.00639** **
	Recent	0.0715	0.0191	**0.000105** ***	
CA1	Remote	0.0490	0.0115	**0.000409** ***	0.245
	Recent	0.0653	0.0157	**0.00000381** ***	
Subiculum	Remote	0.0382	0.0149	**0.0191** *	0.396
	Recent	0.0575	0.0184	**0.0057** **	
Pre/parasubiculum	Remote	0.0637	0.0242	**0.0163** *	0.324
	Recent	0.0911	0.0213	**0.00041** ***	
Uncus	Remote	0.0643	0.0221	**0.00895** **	0.179
	Recent	0.0924	0.0232	**0.0008** ***	

Table 9. Interrater reliability results for hippocampal subfield segmentation. Dice similarity coefficients for the Two researchers independently segmented 5 of 20 participants. All subfields achieved Dice coefficients that were higher than the pre-established thresholds from previous work (Barry et al., 2021; Berron et al., 2017; Bonnici et al., 2013; Chadwick et al., 2014; Clark et al., 2023; Dalton et al., 2019; Lee et al., 2014; Palombo et al., 2013; Yeung et al., 2019).

Subfield	Dice similarity coefficient (mean ± SD)	
	Left	Right
DG	0.86 ± 0.03	0.86 ± 0.03
CA2/3	0.69 ± 0.01	0.70 ± 0.04
CA1	0.77 ± 0.02	0.79 ± 0.01
Subiculum	0.80 ± 0.03	0.78 ± 0.02
Pre/parasubiculum	0.70 ± 0.03	0.71 ± 0.03
Uncus	0.87 ± 0.03	0.85 ± 0.04

5.4.4. Individual remote memories were represented in the mPFC deep layers and lateral temporal cortex deep and middle layers

Having conducted all the univariate analyses that I could notwithstanding the superior layer bias, I next sought to understand whether individual remote memories could be decoded from the multivoxel activity patterns in any of the hippocampal subfields or layers of the seven cortical areas that were engaged during memory retrieval. I divided each of the seven cortical

areas (the bilateral mPFC, lateral prefrontal cortex, mLTC, sLTC, retrosplenial cortex, parahippocampal cortex and occipital cortex) into 3 cortical layers: superficial, middle and deep. This yielded a total of 27 bilateral cortical layers and hippocampal subfield ROIs.

For each ROI, I used RSA (Kriegeskorte, Mur, & Bandettini, 2008) to quantify the strength of the representations on each recall trial, following a similar procedure as a previous study that decoded individual autobiographical memories (Barry et al., 2018). For each trial, the correlation of voxel activation patterns within recall trials of the same memory were subtracted from the correlation of patterns between recall trials of all other memories to yield the representation strength of the trial (Figure 30). If a recall trial activated a similar voxel response pattern as trials in which the same memory was recalled, and activated a distinct pattern to trials in which different memories were recalled, then it would have a high representation strength. This same procedure was repeated for the counting. Then the trial-by-trial representation strengths were averaged within each condition: remote memories, recent memories, and counting. A positive representation strength across all 20 participants indicated that the memory or counting was decodable above chance. In other words, the brain area represented the memory or counting.

I detected representations of remote autobiographical memories in two of the ROIs, the mPFC deep layer ($p = 0.0328$) and the mLTC deep layer ($p = 0.0117$). Furthermore, in both the mPFC deep layer and the mLTC deep layer, the remote memory representations were significantly stronger than the counting representations (mPFC: $p = 0.046$; mLTC: $p = 0.025$), although there was no detectable difference in representation strength between remote and recent memories (mPFC: $p = 0.33$; mLTC: $p = 0.786$). There were also no detectable representations (Table 10) or condition differences (Table 11) in representation strengths in any of the remaining 25 ROIs.

I detected representations of recent autobiographical memories in one of the 27 ROIs, the mLTC middle layer. These recent memory representations were significantly stronger than representations of counting ($p = 0.000201$) and the remote memories ($p = 0.007$).

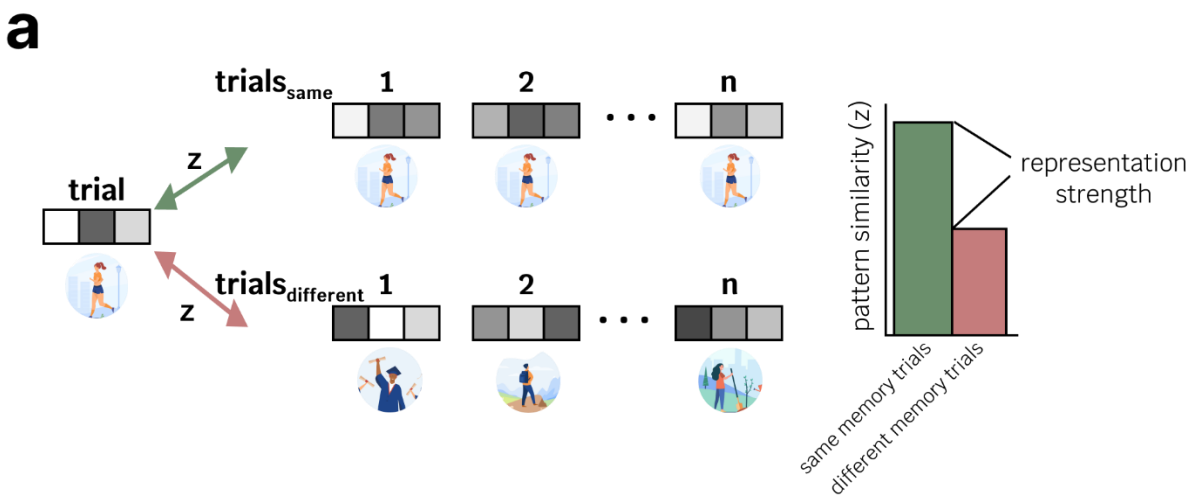
a
trials same
1
2
n
trial
z
z
trials different
1
2
n
pattern similarity (z)
representation strength
same memory trials
different memory trials

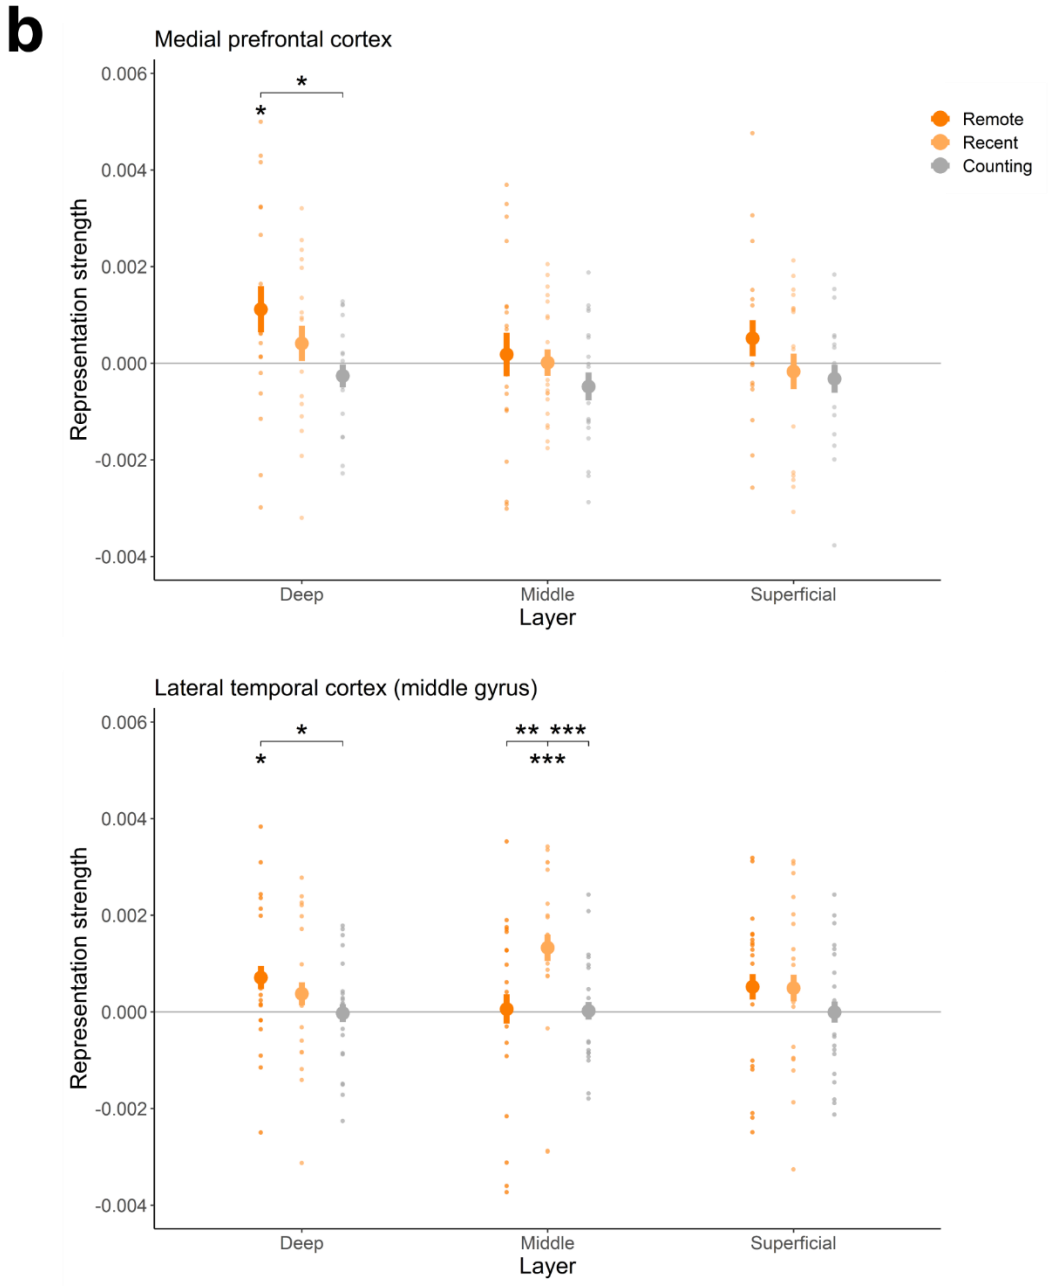
b
Medial prefrontal cortex
Remote
Recent
Counting
Representation strength
0.006
0.004
0.002
0.000
-0.002
-0.004
Deep
Middle
Superficial
Layer
Lateral temporal cortex (middle gyrus)
Representation strength
Deep
Middle
Superficial
Layer

Figure 30. Representations of individual remote memories, recent memories and counting. (a) For each memory, the average Fisher Z transformed correlation in multivoxel pattern of activation associated with recall trials of the same memory was calculated giving a "same memory similarity score". This was subtracted from its "different memory similarity score", the average correlation with multivoxel patterns of recall trials of different memories, to yield the representation strength of the memory. (b) Two cortical areas contained layers with representation strengths that were significantly greater than zero and different between remote memories, recent memories and/or counting. These were the mPFC deep layer, mLTC deep layer and mLTC middle layer.

Table 10. Data used to generate Figure 30 - representations of individual remote autobiographical memories, recent autobiographical memories, and counting.

Brain area	Cortical layer	Condition	Representation strength		W-statistic	p-value
			mean	SD		
Medial prefrontal cortex	Deep	Remote	0.001116	0.002147	162	**0.0328 ***
		Recent	0.000410	0.001634	138	0.2310
		Counting	-0.000260	0.001053	78	0.3300
	Middle	Remote	0.000182	0.002014	120	0.5960
		Recent	0.000015	0.001212	105	1.0000
		Counting	-0.000482	0.001302	61	0.1050
	Superficial	Remote	0.000517	0.001669	140	0.2020
		Recent	-0.000170	0.001642	94	0.7010
		Counting	-0.000322	0.001307	80	0.3680
Lateral prefrontal cortex	Deep	Remote	-0.000211	0.001175	78	0.3300
		Recent	0.000029	0.001112	109	0.8980
		Counting	-0.000169	0.000498	68	0.1770
	Middle	Remote	0.000079	0.001143	117	0.6740
		Recent	-0.000106	0.000693	99	0.8410
		Counting	-0.000063	0.000577	98	0.8120
	Superficial	Remote	0.000291	0.001255	126	0.4520
		Recent	0.000051	0.000816	109	0.8980
		Counting	0.000100	0.000536	129	0.3880
Lateral temporal cortex (superior gyrus)	Deep	Remote	0.000225	0.001733	465	0.4640
		Recent	0.000023	0.001503	426	0.8350
		Counting	0.000069	0.000503	439	0.7020
	Middle	Remote	-0.000034	0.001993	364	0.5410
		Recent	0.000193	0.001282	516	0.1560

		Counting	0.000327	0.001021	510	0.1810
	Superficial	Remote	-0.000196	0.001632	374	0.6330
		Recent	-0.000017	0.001316	398	0.8770
		Counting	-0.000025	0.000777	394	0.8350
Lateral temporal cortex (middle gyrus)	Deep	Remote	0.000710	0.001518	598	**0.0117 ***
		Recent	0.000375	0.001501	519	0.1450
		Counting	-0.000026	0.001165	413	0.9730
	Middle	Remote	0.000059	0.001926	464	0.4720
		Recent	0.001326	0.001731	703	**0.0000841 *****
		Counting	0.000020	0.001150	422	0.8770
	Superficial	Remote	0.000519	0.001654	546	0.0685
		Recent	0.000492	0.001736	535	0.0942
		Counting	-0.000005	0.001373	410	1.0000
Retrosplenial cortex	Deep	Remote	-0.000030	0.005202	105	1.0000
		Recent	-0.000122	0.004051	103	0.9560
		Counting	-0.001282	0.002817	60	0.0973
	Middle	Remote	0.000552	0.003727	129	0.3880
		Recent	0.000468	0.003919	112	0.8120
		Counting	-0.000612	0.003665	95	0.7290
	Superficial	Remote	0.000303	0.004832	105	1.0000
		Recent	-0.000139	0.004691	107	0.9560
		Counting	0.000449	0.003104	120	0.5960
Parahippocampal cortex	Deep	Remote	0.000018	0.004905	98	0.8120
		Recent	-0.000650	0.003757	66	0.1540
		Counting	-0.000471	0.002175	86	0.4980
	Middle	Remote	0.000997	0.003278	133	0.3120
		Recent	-0.000454	0.003330	95	0.7290
		Counting	0.000217	0.003367	118	0.6480
	Superficial	Remote	0.000848	0.003569	138	0.2310
		Recent	-0.000224	0.004472	90	0.5960
		Counting	-0.000302	0.002300	87	0.5220

Occipital cortex	Deep	Remote	-0.000264	0.001111	75	0.2770
		Recent	0.000058	0.000970	118	0.6480
		Counting	0.000042	0.000779	109	0.8980
	Middle	Remote	0.000008	0.001319	105	1.0000
		Recent	-0.000017	0.001046	108	0.9270
		Counting	0.000027	0.000901	113	0.7840
	Superficial	Remote	-0.000140	0.001362	90	0.5960
		Recent	-0.000534	0.001473	59	0.0897
		Counting	0.000097	0.001052	112	0.8120
DG	N/A	Remote	-0.000353	0.003229	93	0.6740
		Recent	-0.000929	0.002005	59	0.0897
		Counting	0.000022	0.001964	106	0.9850
CA2/3	N/A	Remote	-0.001795	0.005319	70	0.2020
		Recent	0.000067	0.006467	107	0.9560
		Counting	0.000574	0.003606	120	0.5960
CA1	N/A	Remote	-0.000638	0.002059	79	0.3490
		Recent	-0.000243	0.001910	91	0.6220
		Counting	0.000310	0.001964	116	0.7010
Subiculum	N/A	Remote	0.000067	0.002610	93	0.6740
		Recent	0.000266	0.003003	116	0.7010
		Counting	0.000603	0.002058	146	0.1330
Pre/parasubiculum	N/A	Remote	-0.001941	0.006136	61	0.1050
		Recent	-0.001687	0.005438	79	0.3490
		Counting	0.000048	0.002354	107	0.9560
Uncus	N/A	Remote	0.000875	0.003661	127	0.4300
		Recent	0.000496	0.003024	126	0.4520
		Counting	0.000317	0.002344	105	1.0000

Note: The hippocampal subfields were not divided into layers and so the Cortical Layers column is filled with N/A.

Table 11. Pairwise comparisons between representation strengths of remote autobiographical memories, recent autobiographical memories and, counting.

Brain area	Cortical layer	Remote vs. Recent		Remote vs. Counting		Recent vs. Counting	
		W-statistic	p-value	W-statistic	p-value	W-statistic	p-value
Medial prefrontal cortex	Deep	135	0.33	169	**0.046 ***	143	0.33
	Middle	121	0.588	151	0.269	134	0.588
	Superficial	129	0.776	152	0.248	116	0.776
Lateral prefrontal cortex	Deep	80	0.831	92	0.831	135	0.831
	Middle	121	1	117	1	100	1
	Superficial	127	1	129	1	96	1
Lateral temporal cortex (superior gyrus)	Deep	410	1	455	1	415	1
	Middle	294	0.245	280	0.245	381	0.702
	Superficial	382	1	355	1	396	1
Lateral temporal cortex (middle gyrus)	Deep	453	0.786	607	**0.025 ***	474	0.786
	Middle	191	**0.007 ****	428	0.814	707	**0.000201 *****
	Superficial	403	0.93	516	0.39	523	0.39
Retrosplenial cortex	Deep	119	1	121	1	133	0.936
	Middle	102	1	117	1	117	1
	Superficial	114	1	102	1	95	1
Parahippocampal cortex	Deep	116	1	106	1	93	1
	Middle	140	0.606	118	0.996	86	0.996
	Superficial	119	1	143	0.495	106	1
Occipital cortex	Deep	82	0.818	68	0.531	114	0.818
	Middle	95	1	106	1	111	1
	Superficial	126	0.904	95	0.904	61	0.315
DG	N/A	124	0.996	94	0.996	75	0.831
CA23	N/A	72	0.693	77	0.693	97	0.784
CA1	N/A	75	0.831	80	0.831	88	0.831
Subiculum	N/A	102	1	88	1	87	1
Pre/parasubiculum	N/A	80	0.736	59	0.269	86	0.736
Uncus	N/A	112	1	109	1	114	1

Note: All p-values are Bonferroni corrected. The hippocampal subfields were not divided into layers and so the Cortical Layers column is filled with N/A.

5.4.5. Remote memory representations are correlated with cortical deep layers and hippocampal CA3

I next explored the idea that the remote memory representations in the mPFC deep layers and mLTC deep layers are communicated with other brain areas in the core network. While I did not detect representations in any of the other brain areas when the representation

strengths were averaged across trials, it is possible that memories were represented in these areas on some recall trials. If the mPFC or mLTC deep layers was sharing memory-specific information with another brain area then I would expect the trial-by-trial representation strengths of the brain areas to covary (Coutanche & Thompson-Schill, 2013). I examined this by correlating the trial-by-trial representation strengths of all remote memory trials between the mPFC deep layers and each of the other 26 other ROIs under investigation and then repeating this for the mLTC. Previous studies have interpreted correlated representation strengths as an indication of information sharing between areas (Aitken & Kok, 2022; Koster et al., 2018; Shao et al., 2023).

I found 2 areas with correlated trial-by-trial representation strengths with the mPFC deep layers - the retrosplenial cortex deep layers ($p = 0.04$; Figure 31) and the occipital cortex deep layers ($p = 0.0172$; Figure 31; Table 12 for the results from other brain areas). There were 3 areas with correlated trial-by-trial representation strengths with the mLTC deep layers - the retrosplenial cortex deep layers ($p = 0.024$; Figure 31), the parahippocampal cortex deep layers ($p = 0.0328$; Figure 31) and hippocampal CA2/3 ($p = 0.0107$; Figure 31; see Table 13 for the results from other brain areas).

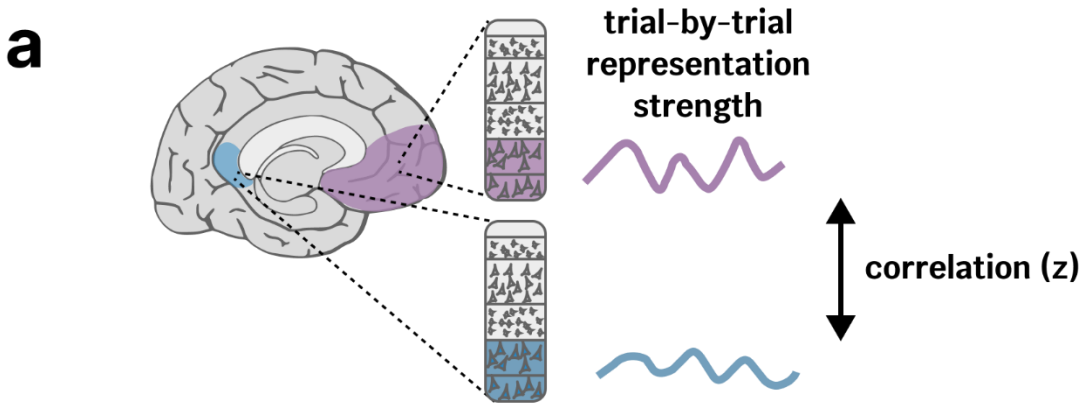
a
trial-by-trial
representation
strength
correlation (z)

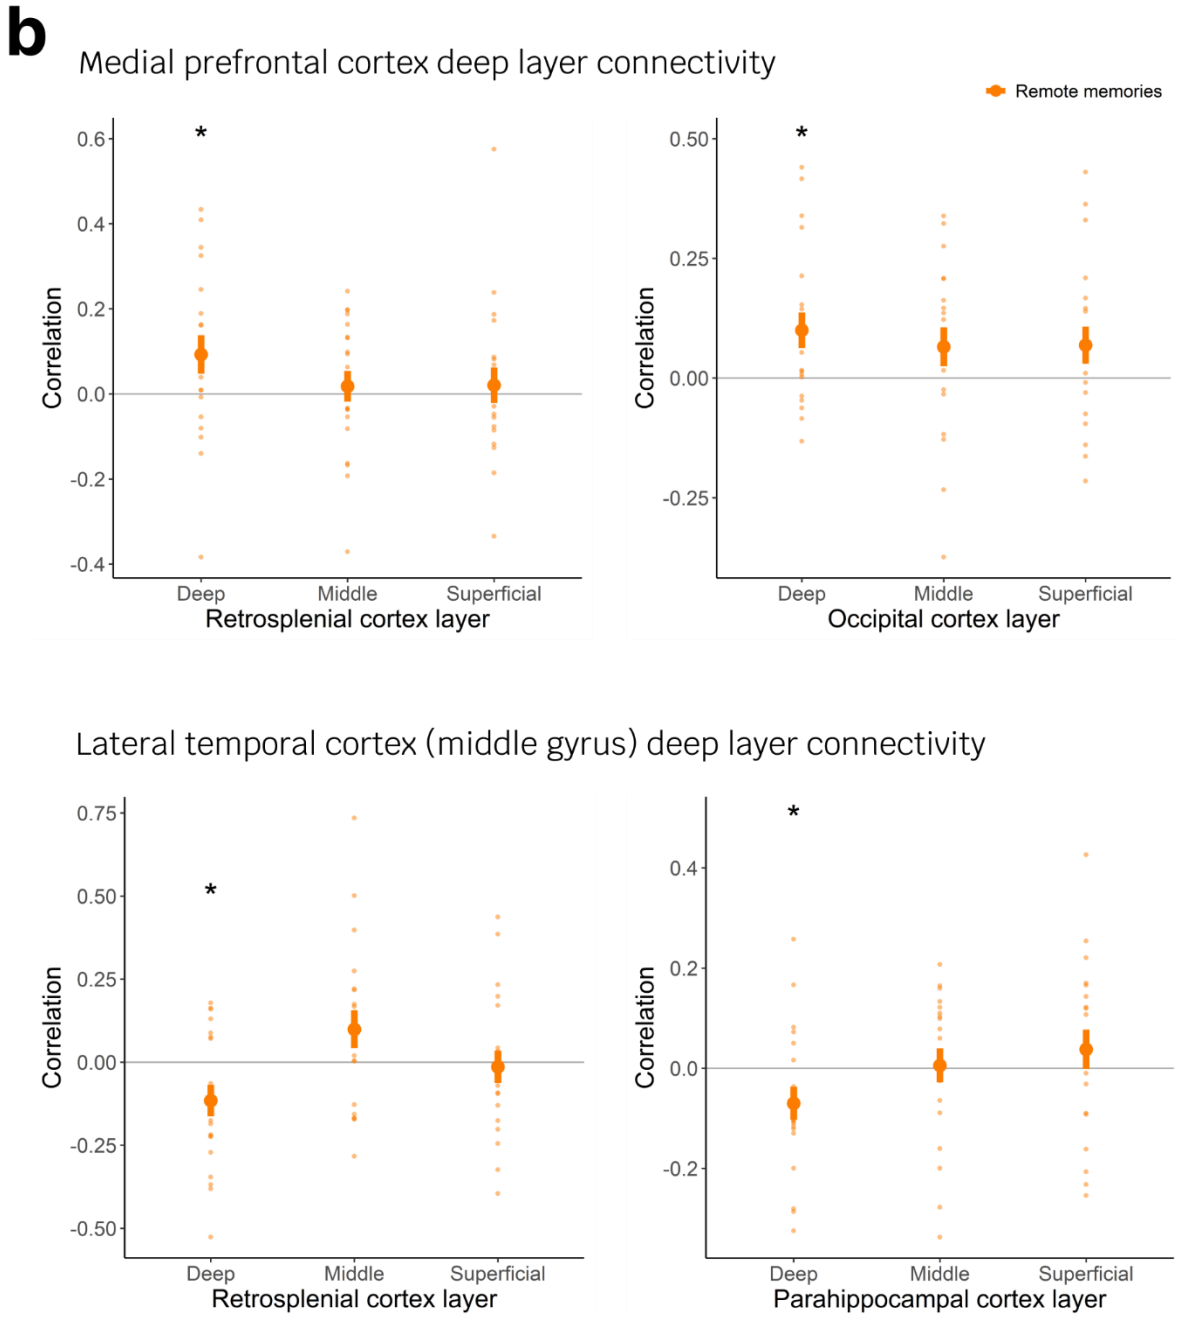
b
Medial prefrontal cortex deep layer connectivity
Remote memories
Correlation
Deep
Middle
Superficial
Retrosplenial cortex layer
Occipital cortex layer
Lateral temporal cortex (middle gyrus) deep layer connectivity
Retrosplenial cortex layer
Parahippocampal cortex layer

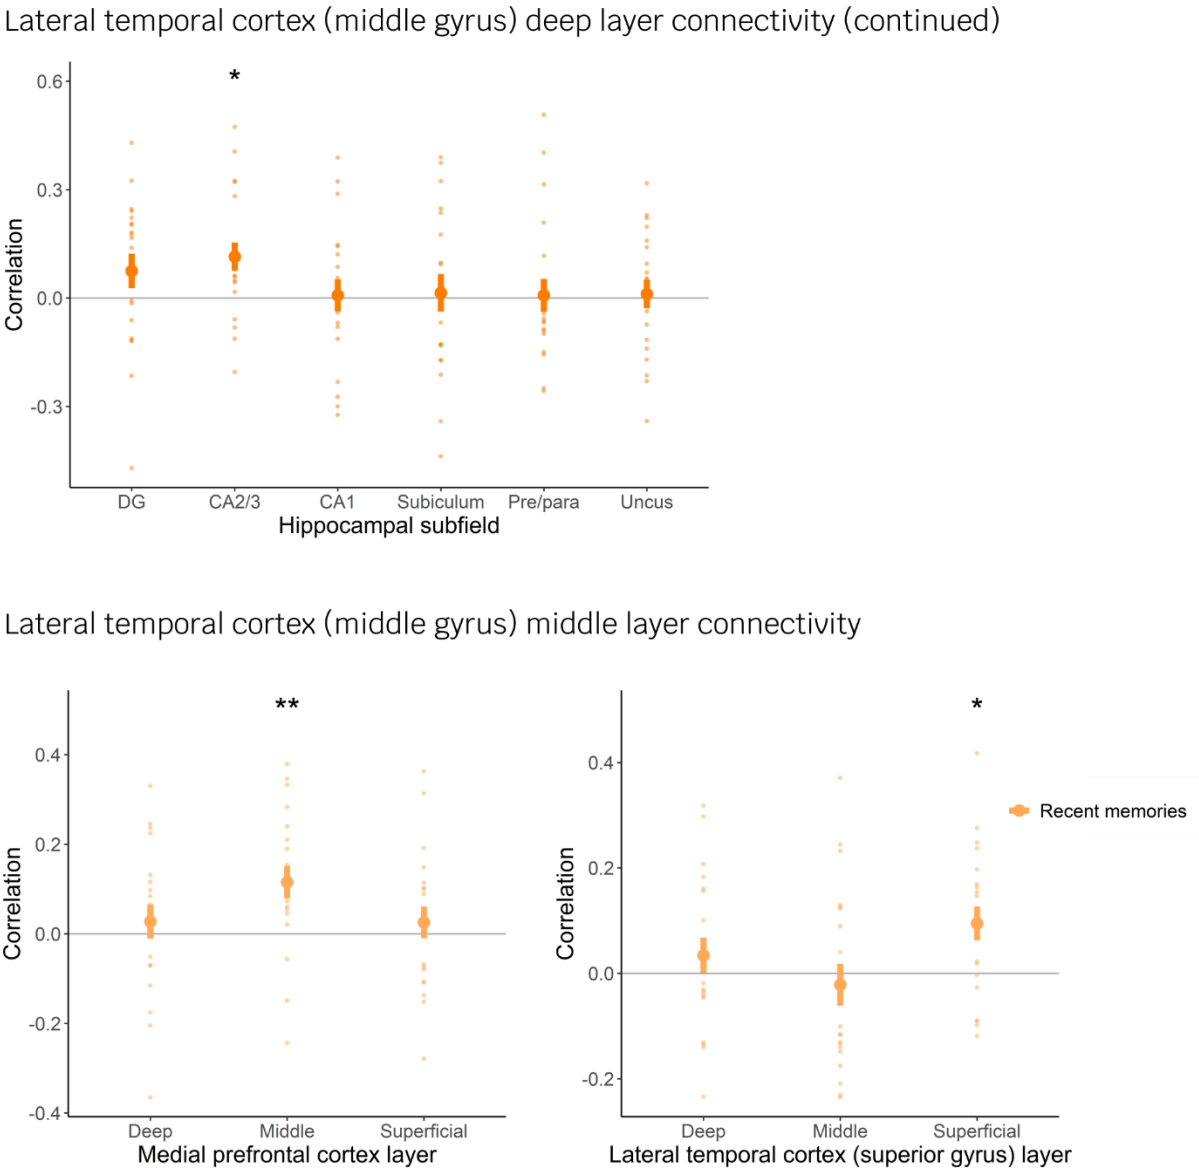

Figure 31. Informational connectivity during remote and recent autobiographical memory retrieval. (a) the trial-by-trial representation strengths of the mPFC deep layer, mLTC deep layer and mLTC middle layer were correlated with those in all other areas in the core network. A correlation that was different from zero suggested informational connectivity, sharing of memory content, between the brain areas. (b) Remote memory representation strengths were correlated between the mPFC deep layer and retrosplenial cortex deep layer and occipital cortex deep layer. The remote memory representation strengths were correlated between the mLTC deep layer and the retrosplenial cortex deep layer, parahippocampal cortex deep layer and hippocampal CA2/3. The recent memory representation strengths were correlated between the mLTC middle layer and the mPFC middle layer and sLTC superficial layer.

Table 12. Informational connectivity of the mPFC deep layer during recall of remote autobiographical memories.

ROI	Layer	Correlation		W-statistic	P-value
		mean	SD		
Lateral prefrontal cortex	Deep	0.0409	0.1837	132	0.33
	Middle	0.0283	0.2323	110	0.869
	Superficial	0.0344	0.2167	130	0.368
Lateral temporal cortex (superior gyrus)	Deep	-0.0338	0.1972	82	0.409
	Middle	-0.0151	0.1971	99	0.841
	Superficial	0.0805	0.1700	151	0.0897
Lateral temporal cortex (middle gyrus)	Deep	0.0158	0.1777	115	0.729
	Middle	-0.0242	0.1913	101	0.898
	Superficial	-0.0390	0.2261	87	0.522
Retrosplenial cortex	Deep	0.0927	0.2020	160	**0.04 ***
	Middle	0.0179	0.1609	126	0.452
	Superficial	0.0206	0.1864	109	0.898
Parahippocampal cortex	Deep	-0.0029	0.2131	115	0.729
	Middle	0.0225	0.1561	132	0.33
	Superficial	0.0378	0.1095	146	0.133
Occipital cortex	Deep	0.0999	0.1665	168	**0.0172 ***
	Middle	0.0651	0.1813	151	0.0897
	Superficial	0.0687	0.1737	146	0.133
DG	N/A	0.0701	0.1591	151	0.0897
CA2/3	N/A	0.0360	0.2257	125	0.475
CA1	N/A	0.0048	0.1526	109	0.898
Subiculum	N/A	-0.0235	0.2003	85	0.475
Pre/parasubiculum	N/A	-0.0386	0.1857	81	0.388
Uncus	N/A	0.0457	0.1424	130	0.368

Table 13. Informational connectivity of the lateral temporal cortex (middle gyrus) deep layer during remote autobiographical memory retrieval.

ROI	Layer	Correlation		W-statistic	P-value
		mean	SD		
Medial prefrontal cortex	Deep	0.0158	0.1777	115	0.729
	Middle	0.0574	0.2162	134	0.294
	Superficial	-0.0404	0.1842	75	0.277
Lateral prefrontal cortex	Deep	-0.0149	0.2433	95	0.729
	Middle	-0.0138	0.1904	96	0.756
	Superficial	0.0276	0.1813	120	0.596
Lateral temporal cortex (superior gyrus)	Deep	0.0266	0.1703	137	0.245
	Middle	0.0021	0.1915	124	0.498
	Superficial	-0.0052	0.1792	103	0.956
Retrosplenial cortex	Deep	-0.1156	0.2102	45	**0.024 ***
	Middle	0.0989	0.2549	146	0.133
	Superficial	-0.0139	0.2174	90	0.596
Parahippocampal cortex	Deep	-0.0699	0.1480	48	**0.0328 ***
	Middle	0.0057	0.1522	121	0.571
	Superficial	0.0377	0.1766	131	0.349
Occipital cortex	Deep	-0.0094	0.1578	89	0.571
	Middle	0.0423	0.2103	128	0.409
	Superficial	0.0114	0.1854	117	0.674
DG	N/A	0.0747	0.2126	152	0.0826
CA2/3	N/A	0.1143	0.1747	172	**0.0107 ***
CA1	N/A	0.0077	0.1993	111	0.841
Subiculum	N/A	0.0144	0.2332	111	0.841
Pre/parasubiculum	N/A	0.0077	0.2039	84	0.452
Uncus	N/A	0.0110	0.1758	114	0.756

5.4.6. Recent memory representations are correlated amongst cortical middle layers

I next investigated the informational connectivity of the mLTC middle layer during retrieval of recent autobiographical memories, using the same method as described above. I found just two ROIs with correlated trial-by-trial representation strengths: the sLTC superficial layer ($p = 0.0136$) and the mPFC middle layer ($p = 0.00422$; Figure 31; Table 14 for the results from the other brain areas).

Table 14. Informational connectivity of the lateral temporal cortex (middle gyrus) middle layer during recent autobiographical memory retrieval.

ROI	Layer	Correlation		W-statistic	P-value
		mean	SD		
Medial prefrontal cortex	Deep	0.0273	0.1696	127	0.43
	Middle	0.1153	0.1593	179	**0.00422 ****
	Superficial	0.0257	0.1589	121	0.571
Lateral prefrontal cortex (middle gyrus)	Deep	0.0041	0.1629	114	0.756
	Middle	0.0226	0.2105	121	0.571
	Superficial	-0.0044	0.2359	99	0.841
Lateral temporal cortex (superior gyrus)	Deep	0.0341	0.1498	127	0.43
	Middle	-0.0215	0.1775	84	0.452
	Superficial	0.0948	0.1441	170	**0.0136 ***
Retrosplenial cortex	Deep	-0.0552	0.2436	81	0.388
	Middle	0.0321	0.1342	135	0.277
	Superficial	-0.0628	0.1962	77	0.312
Parahippocampal cortex	Deep	0.0205	0.1588	111	0.841
	Middle	0.0296	0.1932	124	0.498
	Superficial	-0.0293	0.1900	89	0.571
Occipital cortex	Deep	-0.0114	0.1577	95	0.729
	Middle	0.0680	0.1869	144	0.154
	Superficial	0.0516	0.1912	135	0.277
DG	N/A	0.0500	0.2262	129	0.388
CA2/3	N/A	0.0536	0.2459	141	0.189
CA1	N/A	0.0009	0.1560	114	0.756
Subiculum	N/A	0.0620	0.1443	148	0.114
Pre/parasubiculum	N/A	-0.0364	0.1747	79	0.349
Uncus	N/A	0.0062	0.2295	107	0.956

5.5. Discussion

Supporting previous 3T fMRI studies, I identified a core network of brain areas that was associated with autobiographical memory retrieval, including the mPFC, lateral prefrontal cortex, mLTC, sLTC, parahippocampal cortex, retrosplenial cortex, occipital cortex, and the hippocampus (Buckner & Carroll, 2007; Hassabis & Maguire, 2007; Maguire, 2001; Schacter et al., 2007; Svoboda et al., 2006). Furthermore, I found increased activation during retrieval of recent compared to remote autobiographical memories in the DG and CA2/3 of the hippocampus. Next, I used RSA to investigate multivoxel representations of individual recent and remote memories at the level of the cortical layers and hippocampal subfields. Supporting my hypothesis, I detected multivoxel representations of rich and detailed remote memories

(from 2-5 years ago) in the mPFC deep layers and in the deep layers of a lower order cortical area, the mLTC, indicating the involvement of feedback processing pathways in remote autobiographical memory retrieval.

By contrast, recent autobiographical memories were represented in the middle layer of the mLTC, suggesting the involvement of feedforward signalling. To investigate these intermediate cortical areas further, I used a method called informational connectivity to examine whether the time course of layer-specific multivoxel representations in the mPFC or mLTC were correlated with one another or with any other brain region in the core network. Correlated representation time courses would indicate that there may be some sharing of representation content between the areas (Coutanche & Thompson-Schill, 2013). I found that the time course of remote memory representations in the mPFC and mLTC deep cortical layers was correlated with those in the deep layers of several other cortical areas and with hippocampal CA2/3, supporting my hypothesis of feedback signalling from the mPFC to the hippocampus via other cortical areas. Whereas, the recent memory representations in the mLTC middle layer were correlated with the sLTC superficial layer and the mPFC middle layer, indicating that there may be some feedforward transfer of information during retrieval of recent memories.

5.5.1. The hippocampal microcircuitry supporting memory retrieval over time

The finding of increased activation in the hippocampal DG and CA3 during recent compared to remote autobiographical memory retrieval suggest that the trisynaptic loop, consisting of projections from the entorhinal cortex to the DG, then onwards to CA3, and finally synapsing in CA1, becomes less involved in memory retrieval over time.

Supporting this view, molecular imaging and optogenetic studies in rodents found that, while CA3 is necessary for the retrieval of recently acquired memories, it is significantly less engaged when remote memories are retrieved (Denny et al., 2014; Lux et al., 2016). Retrieval of these older memories was, instead, found to solely rely on CA1 (Lux et al., 2016). These findings inspired the hypothesis that CA3's role in memory retrieval decreases over time, possibly because of a failure to successfully pattern complete remote memory traces, which may have degraded significantly in the hippocampus (Kesner & Rolls, 2015; Lux et al., 2016).

However, a study into human patients with damage to CA3 contradicts this view, finding deficits in the retrieval of autobiographical memories of all ages (Miller et al., 2020).

Furthermore, the strength of CA3 memory representations has actually been found to increase as memories age (Bonnici et al., 2013). Unlike Bonnici et al. (2013), I was unable to detect representations of individual remote (nor recent) memories in CA3 (nor in any of the hippocampal subfields). However, I did find some evidence for the involvement of CA3 in remote memory retrieval. Specifically, I found informational connectivity between the neocortex and CA3 during remote memory retrieval, indicating that memory content may still be passing through CA3 during retrieval of remote memories. These findings contradict the proposition that the trisynaptic loop is less involved in the retrieval of remote memories.

Differences may exist between the 3T fMRI findings and my 7T fMRI findings because of the different spatial resolution achieved by each of these field strengths (Bonnici et al., 2013). In my experiment, we used a 7T fMRI voxel size of 0.8 x 0.8 x 0.8 mm^3, whereas the 3T fMRI study used a voxel size of 1.5 x 1.5 x 1.5 mm^3. The higher resolution afforded by 7T fMRI may enhance the precision of subfield-specific effects.

In sum, my finding that the DG and CA3 are more engaged during recent compared to remote autobiographical memory retrieval supports research in rodents suggesting that the trisynaptic loop may be more engaged for retrieval of recent memories, whereas the monosynaptic pathway is engaged during remote memory retrieval. However, my finding that CA2/3 is informationally connected during retrieval of remote memories, as well as findings from previous human studies of the hippocampal subfields, indicate that the trisyanaptic loop still processes remote memory content during retrieval. These discrepancies highlight the need for further human high-resolution fMRI studies of the hippocampal subfields to examine the degree to which each input pathway to the hippocampus supports recent and remote autobiographical memory retrieval.

Laminar fMRI could aid with the further distinction of these pathways. The trisyaptic loop pathway extends from the entorhinal cortex 2 layer to the inner layers of CA1, via the DG and CA2/3. Whereas the monosynaptic pathway extends from the entorhinal cortex 3 layer to the outer layers of CA1. Unfortunately, significant geometric distortion and signal drop out around the entorhinal cortex in our, and indeed most, fMRI data precluded the investigation of memory representations in this brain area. Future studies could explore whether the entorhinal cortex superficial layers are more involved in recent memory retrieval and middle

layers in remote memory retrieval, which would distinguish between the trisynaptic loop and the monosynaptic pathways. Furthermore, the distinction of BOLD activity in the layers of the hippocampal subfields has recently been proven possible. Therefore, this approach could also be applied to distinguish between trisynaptic loop and monosynaptic pathways, based on the engagement of CA1 layers (Pfaffenrot et al., 2024).

5.5.2. Feedback signalling from medial prefrontal cortex during remote autobiographical memory retrieval

In line with my hypothesis, I found that remote memory representations were present in the deep layers of the mPFC, and connectivity between these deep layers and the deep layers of lower-order cortical areas. Various laminar schemes of feedforward and feedback processing have been proposed (Felleman & Van Essen, 1991; Markov et al., 2014; Shipp, 2023). While these models differ in some aspects, they all agree that long-range feedback connections between distant cortical areas travel in the deep layers. Consequently, the solely deep layer involvement in remote memory retrieval found in this experiment strongly indicates the involvement of feedback processing.

It should be noted that the studies investigating the anatomical pathways of feedforward and feedback signalling have primarily focused on the primary sensory cortices of non-human primates, and their generalisability to higher-order areas like the mPFC is debated (Barbas, 2015; Finn et al., 2021; Godlove et al., 2014; Rockland, 2019). However, recent evidence from a non-human primate study finds that the electrophysiological properties of the cortical layers is consistent throughout the neocortex (Mendoza-Halliday et al., 2024). It therefore seems likely that the laminar connectivity patterns that give rise to these electrophysiological properties is also ubiquitous.

From the informational connectivity results, a potential functional pathway for the top-down influence of the mPFC deep layers on hippocampal CA3 during the retrieval of remote memories can be reasoned (Figure 32). This feedback pathway would originate in the deep layers of the mPFC, travel through the deep layers of the retrosplenial cortex and the mLTC, and ultimately reach hippocampal CA3. Tracing studies in non-human primates have identified direct anatomical connections between each of these cortical areas, except between the LTC and hippocampal CA2/3 (Kobayashi & Amaral, 2003, 2007). Therefore, an additional

intermediate cortical area must exist between the mLTC and CA2/3. The entorhinal cortex is a prime candidate region, as it is known to serve as a gateway for information flow between the hippocampus and the neocortex (Seltzer & Pandya, 1991).

It is worth noting that the informational connectivity between the retrosplenial cortex and the mLTC was negative in sign. That is, there was a negative correlation between representation strengths in each area. When there were stronger memory representations in retrosplenial cortex, there were weaker memory representations in mLTC, and vice versa. This may indicate the existence of inhibitory signalling during remote autobiographical retrieval.

5.5.3. Feedforward signalling to the medial prefrontal cortex during recent autobiographical memory retrieval

I detected representations of recent autobiographical memories in the mLTC middle layer and connectivity between the mLTC middle layer and the sLTC superficial layer and mPFC middle layer (Figure 32). Supporting this, tracer studies in non-human primates have found anatomical connections between the mLTC, sLTC and the mPFC (Bachevalier et al., 1997). The middle and superficial layer involvement in retrieval of recent memories indicates feedforward signalling from the lower-order mLTC and sLTC to the higher-order mPFC. Long-range feedforward signalling pathways originate in the superficial layers of lower-order cortical areas and target the middle layers of higher-order cortical areas. Feedforward signalling of this kind may play a role in neocortical consolidation of recent autobiographical memories during retrieval.

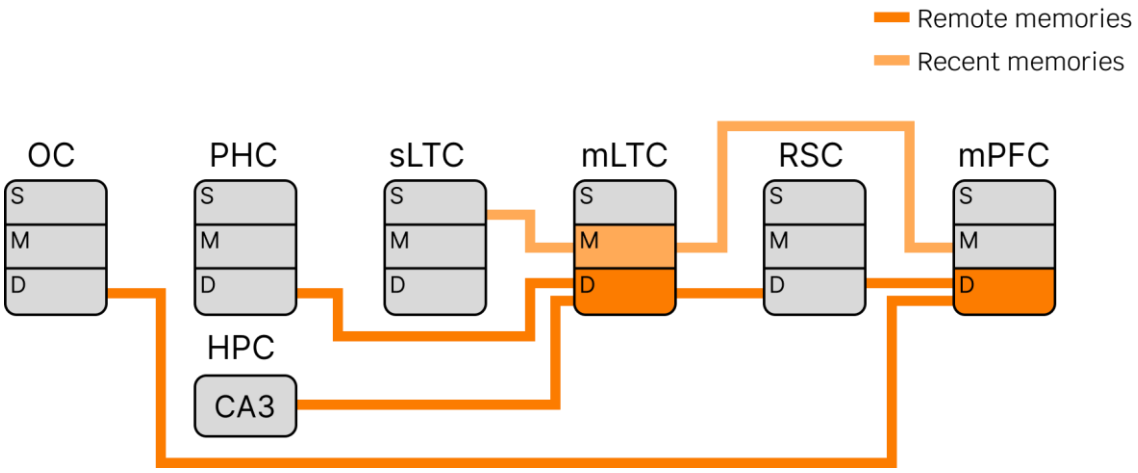

Figure 32: Suggested informational connectivity schematic for remote and recent autobiographical memory retrieval. Representations of remote memories were detected in the mPFC deep layers (dark orange) and mLTC deep layers (dark orange). These areas were informationally connected to several other cortical areas (indicated by dark orange lines). Representations of recent memories were detected in the mLTC middle layer (light orange). This area is informationally connected to other areas (indicated by the light orange lines). OC = occipital cortex, PHC = parahippocampal cortex, HPC = hippocampus, CA = Cornu Ammonis, sLTC = lateral temporal cortex (superior gyrus), mLTC = lateral temporal cortex (middle gyrus), RSC = retrosplenial cortex, mPFC = medial prefrontal cortex.

5.5.4. The medial prefrontal cortex and event schema

Previous studies have identified an important role for the mPFC in storing and processing event schemas (Audrain & McAndrews, 2022; Guo & Yang, 2020; Preston & Eichenbaum, 2013; Reagh & Ranganath, 2023; van Kesteren et al., 2020). Event schemas are general knowledge, derived from multiple prior experiences, about the typical sequence of events that would unfold in a certain context (e.g., eating out at a restaurant). Many studies have highlighted a role of event schemas in memory retrieval (Audrain & McAndrews, 2022; Maguire et al., 1999; Preston & Eichenbaum, 2013; Reagh & Ranganath, 2023; Spalding et al., 2015; van Kesteren et al., 2020); they may provide a structured way of cueing memories, by stepping through the various stages of the event in sequence (e.g., sit down at table, receive menu, order drinks etc.; Mandler, 1984). These cues could be used to guide selection of the specific memory details, which may be located in lower order cortical areas. This role of schemas may be particularly necessary for retrieval of remote memories where the specific event details are less accessible.

An outcome of using event schemas to cue retrieval of remote memories is that, when the event was non-typical, the specific details may be misremembered. Instead, the content of

typical events may be retrieved (e.g., thinking you ordered ice cream for dessert as it is your favourite). Indeed, people tend to confabulate information that is schema-congruent. For example, when tasked with recalling words from list of thematically related words (e.g., hill, valley and range), healthy controls tend to erroneously recall schema-congruent words that were not included in the list (e.g., mountain; Roediger & Mcdermott, 1995). Patients with mPFC damage make significantly fewer errors on this task (Ciaramelli et al., 2006; Warren et al., 2014). Although, when recalling real-life autobiographical memories, mPFC damaged patients tend to confabulate highly unrealistic stories, perhaps because their recollections are unconstrained by event schema (Gilboa, 2010).

5.5.6. Memory representations in other neocortical areas

Both remote and recent autobiographical memories were detectable in the mLTC, although in different layers. The lateral temporal cortex plays an important role in retrieving semantic aspects of autobiographical memories (Maguire et al., 2000). These aspects include both personal or generic information about the specific memory, such as people, places or objects. These semantic memory details may be selected by schema representations in the mPFC, via feedback connections, during remote autobiographical memory retrieval and consolidated into schema, via feedforward connections, during recent autobiographical memory retrieval.

Several other cortical areas, the retrosplenial cortex, occipital cortex and parahippocampal cortex had a time course of representation strengths that was correlated with that in the mPFC or the mLTC. These lower-order cortical areas may contain the perceptual and further semantic details that are specific to the autobiographical event being remembered. Additionally, representations in mLTC deep layer directly correlated with hippocampal CA3. Semantic details represented in the mLTC may be transmitted to CA3 for pattern completion of the full memory.

5.5.7. Qualitative differences between remote and recent memories

It is important to note that qualitative differences between recent and remote memories may contribute to the representational differences observed in this study. For events to be retained in memory for many years with a high degree of vividness and detail, like the remote events sampled in this study, they are likely to have a high degree of personal and emotional

significance. By contrast, recent memories may not contain the same degree of significance. Additionally, remote memories may evoke different emotional experiences to recent memories during their retrieval, such as the feeling of nostalgia. These potential qualitative differences between remote and recent memories could explain their representational differences in the mPFC. Indeed, previous research has shown that the mPFC processes self-referential, personally relevant and emotionally charged information (Gusnard et al., 2001; Abraham, 2013; Kim & Johnson, 2013). A direct examination of how these variables mediate the strength of memory representations in the mPFC could be an interesting direction for future research.

5.5.8. Conclusions

In summary, my findings demonstrate a role for the mPFC deep layers, and their connectivity with the deep layers of lower order cortical areas and hippocampal CA3, in retrieval of remote memories. The deep layer involvement supports the idea that information is transmitted from the mPFC to the hippocampus via feedback signalling pathways. Furthermore, my findings indicate a role of the mLTC middle layer and its connectivity with the sLTC superficial layer and mPFC middle layer in the retrieval of recent memories. This middle and superficial layer involvement suggest a role of feedforward signalling from the lateral temporal cortex to the mPFC. An interesting future direction would be to explore the precise nature of the memory information that is shared between the mPFC and wider network of brain areas, which I hypothesise may be related to event schemas.

6. Experiment 4: The neural microcircuitry underpinning the perception of scenes and events

6.1. Introduction

In this experiment, I sought to examine the specific role of the mPFC and hippocampus in processing static scenes and dynamic events that were not autobiographical and over which I had much more experimental control.

Participants watched short visual animations during fMRI scanning. The animations either displayed sequences of unrelated real-world scenes (Unlinked Scenes) or real-world scenes that were linked to form dynamic events (Linked Scenes). Control animations showed sequences of static abstract patterns (Unlinked Non-scenes) or dynamic evolving patterns (Linked Non-scenes). A control counting task was also included as a baseline condition.

First, I used a whole-brain mass-univariate analysis to identify a set of brain areas that responded during scene and event perception, which included the mPFC and hippocampus. Subsequently, I examined whether real-world, scene-based animations could be decoded from non-scene animations in any of these brain regions based on their multivoxel activation patterns, allowing the identification of scene-selective brain regions. I found that the DG, CA2/3 and CA1 hippocampal subfields decoded both Linked Scenes from Linked Non-scenes and Unlinked Scenes from Unlinked Non-scenes. This supports previous studies finding that the hippocampus has a preference for processing scene imagery over other types of stimuli. By contrast, the mPFC deep, middle and superficial layers decoded Linked Scenes from Linked Non-scenes, but not Unlinked Scenes from Unlinked Non-scenes. This suggests that the mPFC may have a preference for processing scene imagery, only when in the form of dynamic, temporally extended events. Finally, I investigated the nature of communication between the mPFC, hippocampus and wider system of brain regions using the method called informational connectivity (Coutanche & Thompson-Schill, 2013), finding extensive information sharing between the mPFC layers, hippocampal subfields and the layers of all other cortical regions in the core network during scene and event processing.

6.2. Background

A core network of brain regions, including the mPFC, hippocampus, lateral prefrontal, lateral temporal, retrosplenial, parahippocampal and visual cortices, are consistently engaged in any task that involves the visualisation of scenes and events, including retrieving past, planning future, and perceiving present moment events (Buckner & Carroll, 2007; Hassabis, Kumaran, & Maguire, 2007; Hassabis & Maguire, 2007; Hasson et al., 2008; Maguire, 2001; Robin et al., 2018; Schacter et al., 2007; Summerfield et al., 2010; Svoboda et al., 2006; Zeidman, Mullally, et al., 2015).

In Chapter 5, I investigated the microcircuits formed by these brain regions in support of remote and recent autobiographical memory retrieval. In this next experiment, I sought to understand the roles of the mPFC layers and hippocampal subfields in the temporal processing of events. I define a scene as a 3D environment that could be encountered in the world (e.g., a beach, kitchen or garden) and an event as a dynamic sequence of actions unfolding within a scene (e.g., a person swimming, preparing a meal, or sowing seeds) to create a story-like narrative.

Focal, bilateral damage to the hippocampus results in severe impairments in any task requiring the visualisation of scenes (Andelman et al., 2010; Scoville & Milner, 1957; St-Laurent et al., 2009; Steinvorth et al., 2005; reviewed in McCormick, Ciaramelli, et al., 2018). Conversely, damage to the mPFC does not prevent the vivid recollection of static scenes when specific cues are used (Kurczek et al., 2015). However, it does impair the visualisation of events that unfold over time (Bertossi & Ciaramelli, 2016; reviewed in McCormick, Ciaramelli, et al., 2018). Therefore, the hippocampus appears crucial for the visualisation of scenes, while the mPFC may integrate scene information across time to produce events that unfold over many seconds, minutes or hours. In the real-world, scenes are rarely static. They are often comprised of moving people and objects. Furthermore, our perspective of a given scene changes as we shift our gaze and position within it. The mPFC has therefore been proposed to provide a framework to the hippocampus into which scene imagery is continuously updated (Monk, Barry, et al., 2021).

A previous MEG study found higher synchronisation of mPFC and hippocampal theta oscillations during the visualisation of dynamic events when compared to static objects or

scenes (Kaplan et al., 2017). Furthermore, the mPFC response precedes and drives the hippocampal response during the reconstruction of past autobiographical events (McCormick et al., 2020; Nawa & Ando, 2019, 2020) and during imagination of scenes (Barry et al., 2019; Monk, Dalton, et al., 2021). Furthermore, in Experiment 3 (Chapter 5), I found that remote autobiographical memory retrieval involves the mPFC deep layers and feedback signalling to the deep layers of lower order cortical regions and the hippocampus, supporting the idea that the mPFC guides event processing in a top-down fashion.

In this experiment, I leveraged the high spatial resolution of 7T MRI to further investigate scene and event processing at the level of the cortical layers and hippocampal subfields, to reveal the microcircuitry underpinning these functions.

Specifically, I asked:

5. Which brain regions are active during the perception of scenes and events?
6. Are real-world scenes and/or events distinguished from non-scene imagery in specific mPFC layers or hippocampal subfields, or the layers of other cortical areas in the core network?
7. What is the nature of connectivity between the mPFC layers, hippocampal subfields and other cortical areas in the core network during scene and event processing?

In the task, participants watched short visual animations composed of individual image frames. These images were either naturalistic scenes containing people and relevant objects (Scenes) or abstract, patterns (Non-scenes). This design allowed for a direct comparison of multivoxel response patterns associated with processing Scenes versus Non-scenes to identify scene-selective brain regions. Additionally, the sequences of images were presented in two formats: linked image sequences that formed an event (Linked) and unlinked, random sequences of images (Unlinked). This manipulation allowed me to examine the differential involvement of brain regions in processing static scenes and dynamic events.

I hypothesised that:

5. A core network of brain areas would respond during scene and event perception, including the mPFC, lateral prefrontal, lateral temporal, parahippocampal, retrosplenial, occipital cortices and the hippocampus, in accordance with previous

research (e.g., Buckner & Carroll, 2007; Hasson et al., 2008; Zeidman, Mullally, et al., 2015) and my results in Experiment 3 (Chapter 5).

6. The hippocampus would show scene-selectivity (i.e., it would distinguish scene imagery from non-scene imagery) irrespective of the temporal condition.
7. The mPFC would distinguish scene imagery from non-scene imagery when in the form of dynamic events (Linked Scenes), but not when in the form of static scenes (Unlinked Scenes). Given that the mPFC has been found to drive activity in the hippocampus during the processing of real-world events, and the deep cortical layers are the source of top-down feedback signals, I expected this effect to be localised to the mPFC deep layers.
8. There would be functional connectivity between the mPFC deep layer and the deep layers of lower-order cortical areas (e.g., the lateral temporal, parahippocampal, retrosplenial and/or occipital cortex) and/or the hippocampus, again based on the perspective that the mPFC controls the construction of scenes and events.

This study will provide, for the first time in humans, an exploration of the differential processing of individual scenes and extended events in the cortical layers and hippocampal subfields.

6.3. Methods

6.3.1 Participants

I recruited and scanned 28 healthy young participants. Five participants moved their head considerably (> 1.6 mm movement) during fMRI scanning and were therefore excluded from the analysis. Three participants did not sufficiently attend to the fMRI task. This manifested as answering more than 2 out of 8 of the attention checking task questions incorrectly (see Section 6.3.2.1). A final participant was excluded for systematically linking together the Unlinked Scene animations such that they were perceived as single events unfolding over time. The remaining 19 participants were 13 females and 6 males aged between 18 and 35 years old (mean = 23.26, SD = 4.45) with no history of neurological or psychiatric conditions.

6.3.2. Experimental procedure

Participants attended for a single visit in which they underwent fMRI and structural MRI scanning. Following training, participants went into the MRI scanner to perform the scene and event perception task, where they watched short visual animations. After the fMRI session, they underwent a structural MRI session. Finally, they undertook post-scan testing and were debriefed outside of the scanner. The visit lasted approximately 3 hours in total, including task training, set up in the 7T MRI scanner, the experimental task, structural MRI scanning, and the post-scan testing and debrief.

6.3.2.1. Stimuli

The stimuli used in this experiment were bespoke short visual animations that had been produced for a previous MEG study (Monk, Barry, et al., 2021). The animations were built from individual image frames depicting either Scenes or Non-scenes (Figure 33a). Scenes were defined as spatially-coherent 3D recognisable environments, populated with a person, and relevant objects. All Scene images contained a main stick-figure character who was positioned relatively centrally in the image. Non-scenes were used as a control condition. They comprised of abstract patterns in a 2D space and contained a main shape near the centre of the image. Each image was formed of dark grey straight lines and circles on a light grey pixelated background. Greyscale images were used to ensure a low luminance contrast.

The image frames were arranged into sequences to form animations. The image frames could either be temporally linked, such that they produced an evolving event (Linked), or they could be unrelated, such that no event narrative could be conceived (Unlinked). In the Linked condition, each image frame displayed a small change relative to the previous image frame, allowing a flip-book-style event to evolve. In the Linked Scene animations, the main stick-figure character interacted with a main object to perform an action (e.g., skateboarding). The background comprised of other objects (e.g., a table) and space defining features (e.g., walls of a room; Mullally & Maguire, 2013) to give a sense of continuity between the image frames. Half of the Linked Scene animations depicted indoor events while the other half were outdoor events. In the Linked Non-scene animations, the main shape, gradually underwent several transformations such as rotation, expansion or shearing. In the Unlinked condition, each image frame was entirely unrelated to all of the other image frames. The central main

character or main shape was different on each successive image frame as were the objects and the backgrounds. This distinction between Linked and Unlinked animations allowed us to investigate the effect of temporal linking on scene perception.

For both the Linked and Unlinked animations, the Scene and Non-scene animations were matched in terms of low-level stimulus features. That is, each Linked Scene animation was matched to a Linked Non-scene animation and each Unlinked Scene animation was matched to an Unlinked Non-scene animation. The matched animations contained central items (main characters and main objects) formed from the same number of pixels. Furthermore, the backgrounds (everything except the central items) of the Non-scene image frames were scrambled versions of the backgrounds of matched Scene image frames.

There were 16 image frames per animation type. Each image frame was displayed for 700 ms and was followed by a "gap" frame of the same duration. Gap frames were blank images comprised of only the pixilated grey background. They were particularly important for the Unlinked condition as they ensured that each scene or non-scene was perceived as separate from the previous.

6.3.2.1. Scene perception fMRI task

During each trial of the fMRI task, a participant was first shown a condition cue for 3 seconds, informing them of the type of animation they were about to watch (Figure 33a). The terminology used for each condition was simplified such that Linked Scenes were called Pictures Linked, Linked Non-scenes were called Patterns Linked, Unlinked Scenes were called Pictures Unlinked, and Unlinked Non-scenes were called Patterns Unlinked. A chain or broken chain symbol was also used to signify Linked or Unlinked animations, respectively. After the cue, the animation was presented for a total duration of 22.4 seconds. A fixation cross was presented for 3 seconds before moving onto the next trial.

In 8 trials (2 trials per condition) a surprise probe question was presented after the animation (Figure 33b). The probe question was "Having watched that clip, would this next image fit well with that clip?". A subsequent novel image frame was then presented. For each condition, there was one trial in which a congruent image frame was shown (i.e., the correct answer was "yes") and one in which an incongruent image frame was shown (i.e., the correct answer was

“no”). The purpose of these questions was to check if the participant was attending to the animations. They were relatively evenly spread throughout the session.

Counting trials were also included, which were used as a baseline condition during the univariate fMRI analysis. In these trials, participants were shown a cue for 3 seconds, informing them that they were going to perform a counting trial. They were then shown a number on the screen and they were tasked with counting up in 3’s from the number for 22.4 seconds with their eyes open, followed by a 3 second fixation cross.

Each animation and each number were presented once in a random order across runs. There was a total of 50 trials (10 per condition, including the Counting condition) that were split into two runs of 25 trials each. Each run was approximately 12 minutes in length. The fMRI task was produced using the MATLAB 2021b in conjunction with the Psychophysics Toolbox Version 3 (PTB-3).

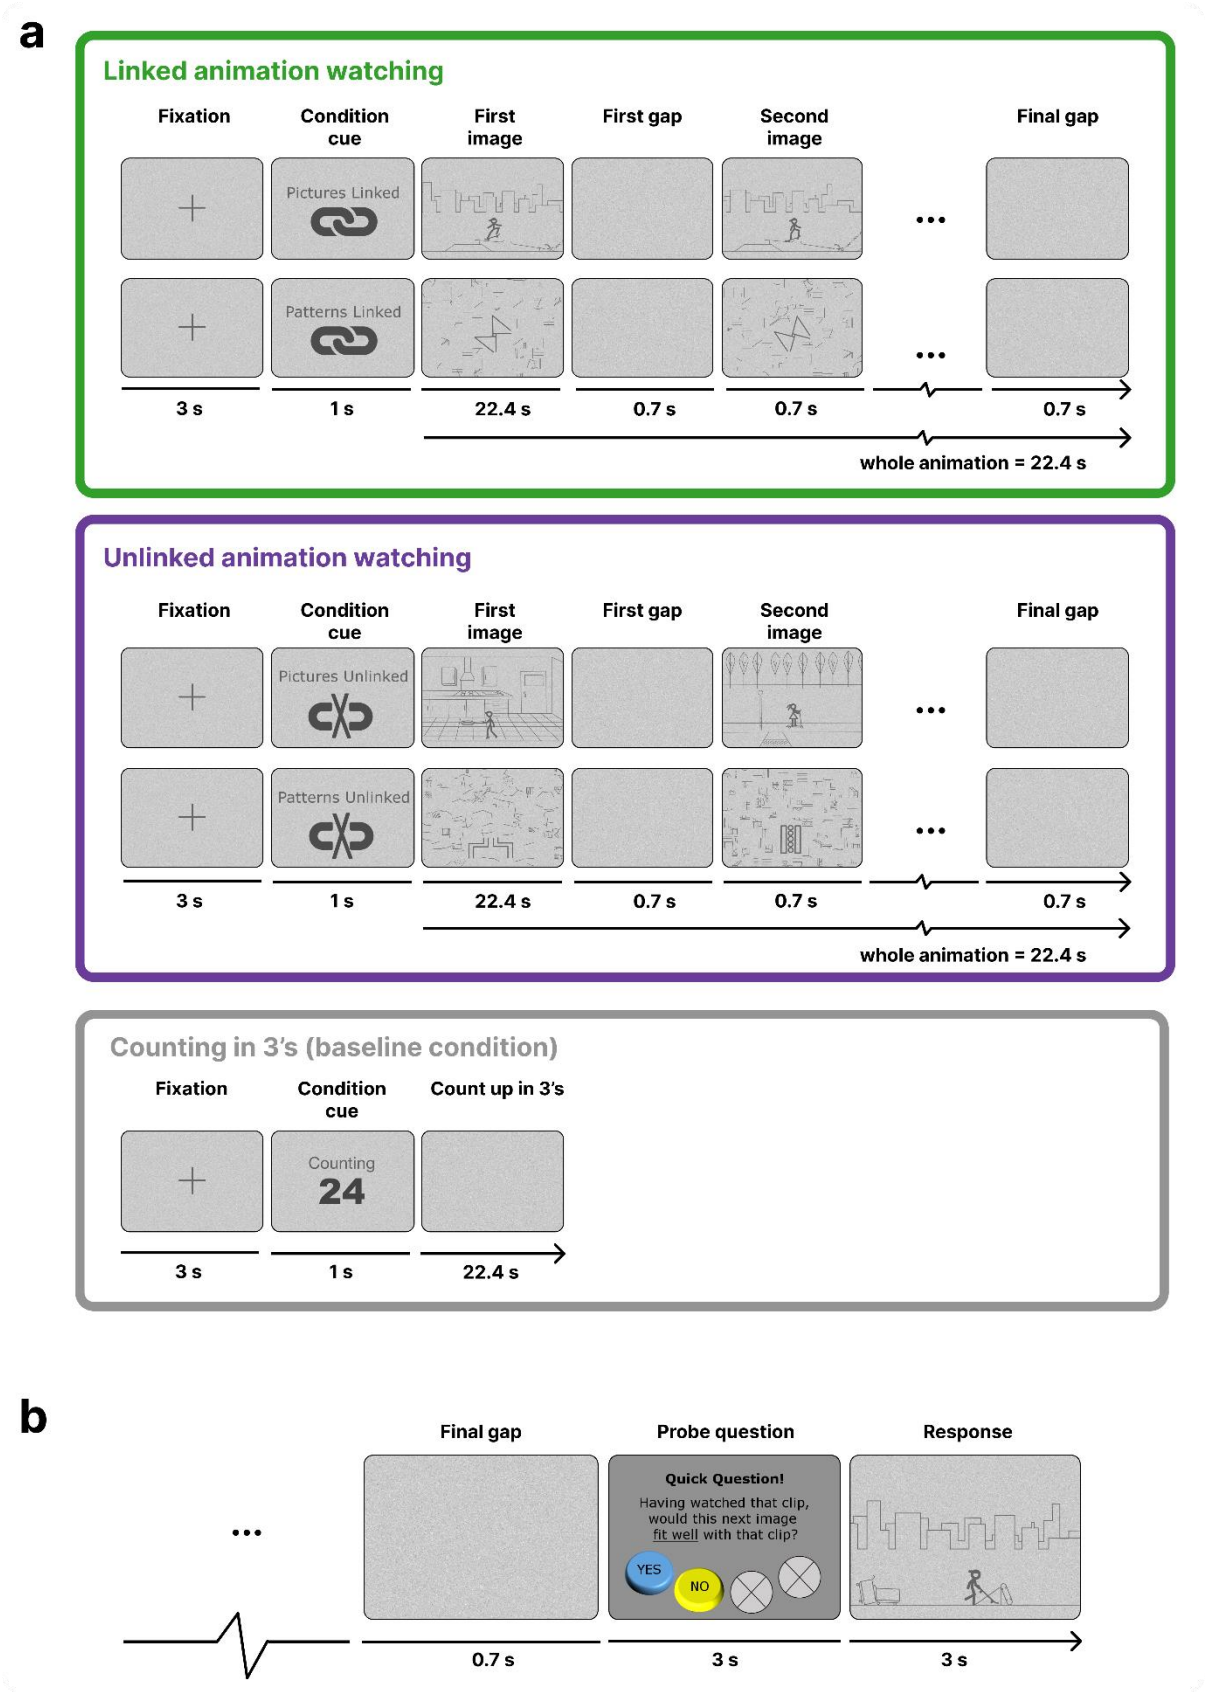

Figure 33. Scene and event perception task. (a) Linked (green box) and Unlinked (purple box) animation watching trials consisted of a fixation cross (3 s), then a cue informing them of the upcoming animation type (1 s), followed by the animation (22.4 s). The animation was built from 16 images each presented for 0.7 s and separated by a blank gap image (also presented for 0.7 s). Linked images (green

box) featured the same central character or shape and background throughout and, when presented in succession, formed a flip-book style movie. Unlinked images were unrelated to one another. Counting trials (grey box) were included as a baseline condition. They started with a fixation cross (3 s), then a number cue (1 s) and finally a period of counting up in 3's from the given number (22.4s) (b) In 20% of animation watching trials a surprise probe question was presented after the animation to check that the participant was attending to the animations.

6.3.2.1. Post-scan testing

Following the fMRI task, participants completed a post-scan test. In this session, they watched each of the animations again. After each animation, they were asked to answer 2 questions about their experience watching this clip during the fMRI task. They were asked to answer the questions with respect to the beginning (first 5 image frames), middle (next 6 image frames) and end (last 5 image frames) of the animation.

The questions were as follows:

1. "To what extent were you linking the images in this clip together?" They could answer on a 5 point scale where 1 = "I did not link the images at all", 2 = "I only linked two images together in my mind but on the whole the images felt unrelated", 3 = "I linked more than two images together in my mind", 4 = "I linked most of the images", and 5 = "I linked all of the images"
2. "To what extent did you feel you were predicting what might happen next?" They could answer on a 5 point scale, where 1 = "I did not predict at all", 2 = "I predicted what might happen next only rarely", 3 = "I predicted what might happen next quite often", 4 = "I predicted what might happen next most of the time", and 5 = "I predicted what might happen next all of the time".

The ratings for linking were compared, using paired t-tests, between conditions to check that the participants were experiencing the Linked animations as unfolding events and the Unlinked animations as many unrelated images. One participant was excluded from further analysis as they responded with a rating of 4 to all of the Unlinked Scene animations, indicating that they were systematically linked these animations to perceive them as single evolving events.

6.3.3. MRI sequences

MRI data were acquired on a Siemens Magnetom Terra 7T MRI system (see Chapter 2, Section 2.6.1).

5.3.4.1. Functional MRI

Partial brain fMRI images with submillimetre resolution (voxel size = 0.8 × 0.8 × 0.8 mm^3) were acquired with a T2*-weighted 3D GE-EPI sequence. The sequence had the following imaging parameters: volume acquisition time was 3872 ms, TR = 44 ms, TE = 18.7 ms, flip angle = 14°, field of view = 192 × 192 × 70.4 mm, slab angle = 15°, direction of phase encoding = anterior-to-posterior, parallel imaging acceleration in both the direction of the partition (factor = 2) and the direction of phase encoding (factor = 4), partial Fourier 6/8 in the phase-encoded direction of the EPI readout, and a binomial (1331) water-selective excitation, and advanced B0 shimming (WIP 1441). Two images were acquired at the start of the sequence with the same parameters except that the opposite phase encoding direction was used (posterior-to-anterior). These images were used for distortion correction of the fMRI data.

5.3.4.2. Structural MRI

A whole brain MT-weighted EPI image was acquired using a T2*-weighted 3D GE-EPI sequence. Many of the parameters (e.g., echo spacing and acceleration factor) used to acquire this image were matched to the fMRI sequence. This meant that the nature of distortion in the MT-weighted image and fMRI images was the same. Some parameters did differ between the acquisitions: MT-based contrast = on, volume acquisition time = 3 m 45 s, TE = 16.97 ms, flip angle = 8°, field of view = 192 × 192 × 128 mm^3, parallel imaging (factor = 4). The MT weighting resulted in higher contrast between grey and white matter, making the MT-weighted image superior for coregistering cortical surfaces to the fMRI data.

A whole brain T1 weighted image with submillimetre resolution (voxel size = 0.65 × 0.65 × 0.65 mm) was acquired using an MP2RAGE sequence (Marques et al., 2010). This MP2RAGE sequence had the following imaging parameters: a volume acquisition time = 8 m 42 s, TR = 5000 ms, TE = 2.54 ms, TI = 900 ms and 2,750 ms, 5° and 3° flip angles, field of view 208 × 208 × 156 mm^3, and in-plane GRAPPA acceleration (factor = 3). The MP2RAGE sequence produces 4 images, including a T1-weighted image and an inverted T1-weighted image. These two

images are used for defining the boundaries between the grey and white matter and the grey matter and CSF.

Two partial brain T2-weighted images with submillimetre resolution (voxel size = 0.52 × 0.52 × 0.5 mm^3) were acquired using a T2-weighted 3D SPACE sequence (Mugler et al., 2000). This T2-weighted sequence used the following parameters: a volume acquisition time = 10 m 25 s, TR = 3,500 ms, TE = 229 ms, flip angle = 8°, field of view = 200 × 169 mm x 56 mm, partition oversampling = 14.3%, and in-plane GRAPPA acceleration (factor = 2). These two images were used for segmenting the hippocampal subfields. I acquired 2 images so that a single image with improved SNR could be used by coregistering, denoising and averaging the images.

6.3.4. Preprocessing

6.3.4.1. Defining cortical surface

As explained in previous chapters, the first preprocessing step was to define the pial and white matter surfaces. The pial surface is the boundary between the grey matter and CSF and the white matter surface is the boundary between grey matter and white matter. Later, two additional surfaces were be produced between the pial and white matter surface that define the boundaries between the cortical layers.

For each participant, a "spliced" whole brain T2-weighted image was produced by coregistering, denoising and averaging the MP2RAGE inverted T1-weighted image and the 2 high resolution partial brain T2-weighted images. This was then inputted to the FreeSurfer (version 7.3.2) recon-all pipeline along with the MP2RAGE T1 weighted image. The recon-all pipeline uses the MP2RAGE T1-weighted image to define the pial and white matter surfaces. It then uses the "spliced" whole brain T2-weighted image to re-define the pial surface availing of the superior grey matter to CSF contrast in this image. I visually inspected the cortical surfaces to ensure that they were accurate to the anatomical boundaries.

6.3.4.2. Motion and distortion correcting fMRI data

I used a map of the B0 field to correct for geometric distortion in the functional images. The FSL (version 6.0) topup tool was used to estimate the B0 field map from the 2 opposite phase encoded images and the first 2 functional images of the fMRI data. A voxel displacement map, which estimated the translation of each voxel due to the B0 field, was then calculated from

the field map using SPM12. This voxel displacement map was then inputted into SPM12 Realign and Unwarp along with the fMRI data. SPM Realign and Unwarp corrects for motion and distortion in the fMRI data, taking into consideration the susceptibility-motion interactions by estimating changes to the voxel displacement map with head motion. Importantly, the fMRI data were only interpolated and resampled once, minimising the loss of effective resolution that can result from interpolation. Finally, a mean fMRI image was calculated from the corrected fMRI data.

The whole brain MT-weighted EPI image was also corrected for geometric distortion using FSL Topup. This time, I used just one image that had been acquired with the opposite phase encoding direction to the fMRI data and the single MT-weighted image with the same phase encoding direction. This image would be used for coregistration of the cortical surfaces to the fMRI data. Therefore, first it was brought into alignment with the fMRI data by coregistering it to the mean fMRI image using SPM.

Five participants moved their head more than 2 voxels (1.6 mm) during the scanning session. Excessive head motion can cause image artefacts such as “ghosting” (repeating of the brain or parts of the brain) and susceptibility-motion interactions, where the susceptibility field changes between functional images altering the nature of distortions (Andersson et al., 2001). Therefore, I decided to exclude these participants from the experiment. The remaining 20 participants moved their head by, on average, 0.85 mm (SD = 0.49 mm) during the course of the fMRI task (although a further participant was excluded for systematically linking together the Unlinked Scene animations such that they were perceived as single events unfolding over time).

6.3.4.3. Aligning cortical surfaces to fMRI data

Three coregistrations were performed to align the pial and white matter surfaces to the fMRI data. First, a NMI coregistration, which maximises the shared information between the two images, was applied using FSL Flirt to move the MP2RAGE T1-weighted EPI image into alignment with the MT-weighted. EPI image. Following this, a BBR, which maximises the signal intensity gradients across the pial and white matter surfaces, was applied using FreeSurfer BBRegister to more the surfaces into alignment with the MT-weighted EPI image. Finally, a recursive BBR, in which a surface is divided into iteratively smaller subsections BBR is

performed on each subsection separately, was applied using the OpenFmriAnalysis toolbox to improve the surface coregistration. During this recursive BBR step, 6 iterations of surface divisions were applied. Furthermore, scaling was allowed in the phase-encoding direction to correct for geometric distortion, which is most extreme along the phase encoding axis (Jezzard & Clare, 1999).

6.3.5. Whole brain mass univariate analysis

Spatial normalisation was performed using SPM12 to warp the fMRI data into MNI space. Smoothing (FWHM = 6 mm) was then performed, also using SPM12. For each participant, a GLM analysis was then performed using SPM12. The 22.4 s animation watching periods of the Linked-Scenes, Linked-Non-scenes, Unlinked-Scenes and Unlinked-Non-scenes, as well as the 22.4 s counting period of the Counting trials, were used as regressors of interest in the GLM. The six motion parameters, estimated during realignment, were also included as nuisance regressors. The 4 animation watching conditions were contrasted with the baseline counting condition to yield a t-statistic map of animation watching. Two additional contrasts were made, one between the Linked animation conditions and the counting baseline and another between the Unlinked animation conditions and the counting baseline. Group level analyses were performed using SPM12 to yield group level t-statistic maps for each contrast.

6.3.6. Delineating regions of interest

6.3.6.1. Cortical layer ROIs

Next, I defined the ROIs that would be used in the ROI-based MVPA analysis. The pial surface was segmented into cortical areas, according to the Desikan-Killiany and Destrieux atlases (Desikan et al., 2006; Destrieux et al., 2010), using FreeSurfer. Cortical areas were selected by consulting the areas activated, at the group level, by animation watching. The cortical area segmentation was coregistered to the fMRI data by applying the same transformation that was applied to the pial surface. Each cortical area, which were in the form of surface masks, was projected across the surface normal, thereby projecting it into volume space. Gaps often existed in the volume masks, which were filled using FSL Maths Edge, Dilate and Erode. Finally, the animation watching group level t-statistic map was consulted again to identify broad areas of activation. Cortical area masks were combined to create masks that corresponded to these

broad areas. Where necessary, some manual editing was performed to bring the anatomical masks into alignment with the group level functional activations.

Next, the cortical area masks were split into 3 equivolume layers (deep, middle and superficial) using the OpenFmriAnalysis toolbox. First, the position of two additional surfaces between (and running parallel to) the pial and white matter surfaces were calculated using the level set method. Then for each voxel, the distance between it and each of the 5 compartments (white matter, CSF, and the superior, middle and deep cortical layers) was calculated (Waehnert et al., 2014). Based on these distances, the percentage of the voxel volume in each of the 5 compartments was estimated. Voxels were assigned to a layer ROI if at least 80% of their volume was within the layer.

To correct for the superficial layer bias effect (see Chapter 2 and Chapter 3 for a detailed explanation), voxels with tSNR below the 65th percentile were excluded (Jia et al., 2021). This is because low tSNR can be an indication of the presence of large veins, which cause a large variation in signal.

To correct for any remaining geometric distortions in the data, Jacobian values were calculated using HySCO in the ACID MATLAB toolbox and voxels with extreme values, below the 5th percentile or above 95th percentile, were excluded (Ruthotto et al., 2012). This is because the Jacobian is a measure of signal displacement of a voxel and is therefore a useful quantitative measure of the degree of geometric distortion in the voxel (Jezzard & Balaban, 1995).

6.3.6.2. Hippocampal subfield ROIs

The hippocampal subfields were manually segmented on the partial brain T2-weighted images. First the two partial brain T2-weighted images were realigned, denoised and averaged to produce a single image with superior SNR. Then, using the ITK-SNAP version 3.2.0 software, 6 hippocampal subfields, the DG/CA4, CA3/2, CA1, subiculum, pre/parasubiculum and uncus, were manually defined according to a detailed segmentation protocol (see Chapter 4). To assess the reliability of manual hippocampal segmentations, inter-rater reliability analysis was performed. A quarter of hippocampi (5 of 19 participants) were segmented by a second, independent experimenter. The similarity in segmentations performed by each experimenter was evaluated using the Dice similarity coefficient, ranging from 0 (no overlap) to 1 (identical; Table 20 for interrater reliability results).

6.3.7. Multi-voxel pattern analysis

In order to perform MVPA analysis, first the multi-voxel pattern of activation caused by each animation trial was estimated. For each animation trial, a GLM was fit to the fMRI data that modelled the 22.4 second animation watching period of the trial and the animation watching periods of all other trials. Motion parameters and drift terms were used as nuisance regressors. The beta estimate of the single animation watching trial represented the neural activation that was unique to that trial. Noisier voxels were down-weighted by dividing the beta estimate of each voxel by the square root of the residuals (Walther et al., 2016).

MVPA was used to attempt to decode the stimulus category (Scene or Non-scene) from the multi-voxel patterns of activation. Specifically, for each participant and each cortical layer ROI, I used an SVM to classify animations as Scenes or Non-scenes based on their multi-voxel pattern of noise normalised beta estimates. This scene decoding was performed separately for Linked and Unlinked animations.

The data was first split into a training set and a testing set. There was a total of 18 animation trials (9 Scene trials and 9 Non-scene trials) used as training instances in the training set. The SVM learned to classify each trial as a Scene or Non-scene based a set of features, in this case the noise normalised beta estimates of each voxel in the ROI. During training, the SVM plots each trial in high-dimensional feature space, where each dimension corresponds to the noise normalised beta estimate of a given voxel in the ROI. It then learns the optimal hyperplane that separates Scene from Non-scene trials, known as a decision boundary. It does this by maximising the margin, the distance between the closest points (called support vectors) of each trial condition to the decision boundary. Importantly, SVM's permit some data points to be within the margin or on the wrong side of the decision boundary, making them less sensitive to noise and more generalisable to new data. I used a linear kernel, meaning that no transformations were performed to the features when plotting them in feature space, and regularisation of C = 1, which strikes a balance between maximising the margin and minimising misclassification. During testing, the new data points are plotted in feature space and classified as Scenes or Non-scenes based on the side of the decision boundary that they fall on.

I used a cross-validation approach to test performance of the classifier. This means that, for each participant and each ROI, multiple SVM's were trained and tested. On each iteration, a different pair of Scene and Non-scene trials were allocated to the testing set and the remaining trials were allocated to the training set. I iterated through all possible pair combinations of Scene and Non-scene trials, yielding a total of 10 tests per trial. For each trial, I calculated its classification accuracy by calculating the proportion of correctly classified tests for that trial.

The trial-by-trial accuracy values were then averaged within Linked trials and Unlinked trials. At the group level, if the SVM achieved above chance (i.e., above 50%) classification accuracy, evaluated using a one-sample t test, for a given cortical layer ROI then I inferred that the cortical layer ROI in some way represented scene imagery or was involved in scene processing. Differences in the accuracy of scene decoding were compared, using Bonferroni corrected paired t tests, between scenes that unfold into events (i.e., Linked) and sequences of scenes that are unrelated (i.e., Unlinked), indicating the involvement of the cortical layer ROI in Linked compared to Unlinked scene processing.

6.3.8 Informational connectivity analysis

I next sought to understand the nature of connectivity between the mPFC layers and the layers of other cortical areas and the hippocampal subfields. To investigate this, I performed informational connectivity analysis, which assesses the similarity of representation time courses between brain areas (Coutanche & Thompson-Schill, 2013). Specifically, the Fisher Z-transformed Pearson product-moment correlation was calculated between the trial-by-trial decoding accuracy values in each mPFC layer and each other ROI under investigation. At the group level, I used a one sample t test to test for a correlation that was reliably different from zero. A positive correlation means that on trials where the mPFC layer was strongly representing scenes, so was the other brain region. Whereas a negative correlation means that stronger representations in the mPFC layer coincided with weaker scene representations in the other brain region.

6.4. Results

6.4.1. Participants maintained attention throughout the task

On average, participants answered 94.74% (SD = 7.59%) of the probe trials correctly showing that they were attending to the animations throughout the session.

6.4.2. Participants linked and predicted successive images during the linked animations

During the debrief session, participants rated, on a scale from 1 to 5, each of the animations in terms of the degree to which they were linking the image frames to perceive them as unfolding events. They provided these linking ratings with respect the beginning, middle and end of the animations. On average, the 19 participants linked all of the image frames in Linked animations across all phases of the experiment, whereas, they did not link the image frames at all in the Unlinked animations (see Table 15). The linking ratings for Linked animations were significantly higher than the linking ratings for Unlinked animations (p = 2.79e-174), thereby validating the Linked versus Unlinked experimental manipulation.

Table 15. Linking ratings of animations.

Condition	Phase	Linking rating	
		mean	SD
Linked Scenes	Beginning	5	0
	Middle	5	0
	End	5	0
Linked Non-scenes	Beginning	4.979	0.092
	Middle	4.979	0.092
	End	4.958	0.126
Unlinked Scenes	Beginning	1.032	0.1
	Middle	1.032	0.1
	End	1.053	0.187
Unlinked Non-scenes	Beginning	1	0
	Middle	1	0
	End	1	0

Participants also rated, on a scale from 1 to 5, each of the animations in terms of the degree to which they were predicting what would happen next at the beginning, middle and end of the animations. Unlinked animations were consistently rated as completely unpredictable throughout the animations, whereas, when watching Linked animations participants reported

predicting to some degree what would happen next throughout the animation (see Table 16). Linked animations were significantly more predictable than Unlinked animations ($p = 9.72e-32$), which was assessed using a one-sample t-test ($\mu = 1$) due to the lack of variation in responses to Unlinked animations.

Table 16. Predicting ratings of animations.

Condition	Animation phase	Predicting rating	
		mean	SD
Linked Scenes	Beginning	2.348	1.055
	Middle	3.103	1.105
	End	2.78	1.142
Linked Non-scenes	Beginning	2.885	1.463
	Middle	3.335	1.326
	End	3.04	1.258
Unlinked Scenes	Beginning	1	0
	Middle	1	0
	End	1	0
Unlinked Non-scenes	Beginning	1	0
	Middle	1	0
	End	1	0

6.4.3. Core network of brain areas are activated by scene and event perception

Next, I identified the brain areas that were activated during scene and event perception by comparing all of the animation watching trials with all of the counting trials. I found significant ($p > 0.01$, uncorrected) activity in several brain areas (Figure 34a), including the mPFC (peak voxel at left: x = -8, y = 51, z = 17; right: x = 15, y = 48, z = 5), retrosplenial cortex (left: x = -17, y = -51, z = 13; right: x = 2, y = -44, z = 8), lateral prefrontal cortex (left: x = -54, y = 21, z = 20; right: x = 34, y = 27, z = -15), mLTC (left: x = -69, y = -17, z = -12; right: x = 48, y = -11, z = -15), sLTC (left: x = -57, y = -2, z = -12; right: x = 45, y = 11, z = -16), parahippocampal cortex (left: x = -27, y = -27, z = -23; right: x = 26, y = -18, z = -19), occipital cortex (left: x = -38, y = -79, z = -10; right: x = 32, y = -41, z = -22), and hippocampus (right: 21, -28, -4; Table 17 for all activated areas). This same set of brain areas was active when Linked and Unlinked recall trials were separately compared with counting trials (Figure 34b and 34c).

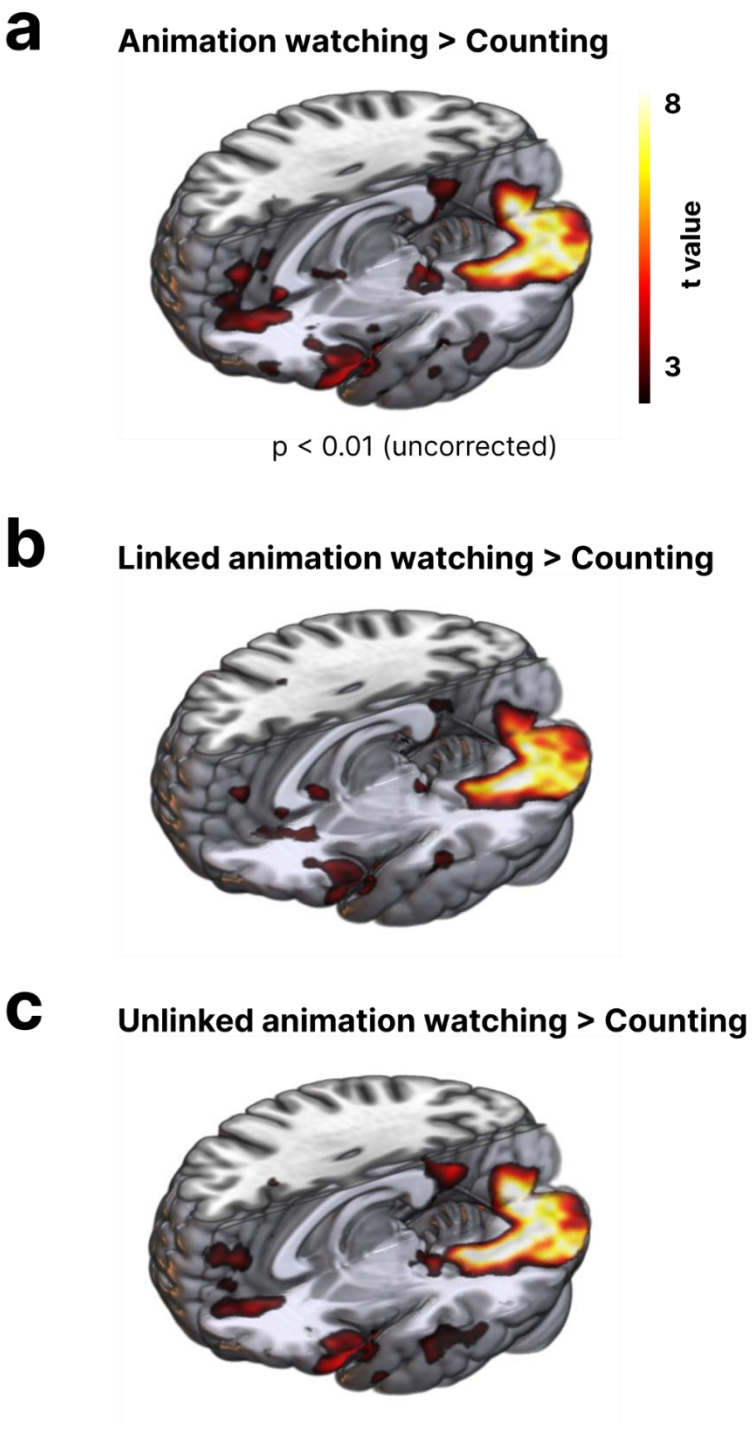

Figure 34. Core network of brain areas active during animation watching. A core network of brain areas was activated during (a) animation watching, (b) Linked animation watching and (c) Unlinked animation watching. T-statistics are thresholded at $p < 0.01$ (uncorrected).

Table 17. Activation peaks for the animation watching > counting baseline contrast.

Brain region	Cluster size (voxels)	Z score	Coordinates		
			x	y	z
Occipital cortex, fusiform gyrus (R)	145670	6.57	32	-41	-22
Occipital cortex, inferior occipital gyrus (L)		6.08	-38	-79	-10
Parahippocampal cortex (L)		2.64	-27	-27	-23
Amygdala (R)	5730	3.95	22	-5	-14
Lateral prefrontal cortex, pars orbitalis (R)		3.72	34	27	-15
Medial prefrontal cortex, medial superior frontal gyrus (L)	8576	3.98	-8	51	17
Medial prefrontal cortex, medial orbital superior frontal gyrus (L)		3.88	-8	47	-9
Medial prefrontal cortex, anterior cingulate (L)		3.62	-11	36	-6
Hippocampus (R)	2174	3.72	21	-28	-4
Lateral temporal cortex, middle gyrus (R)	1361	3.62	48	-11	-15
Lateral prefrontal cortex, pars triangularis (L)	1512	3.28	-54	21	20
Lateral temporal cortex, middle gyrus (L)	584	3.24	-69	-17	-12
Lateral temporal cortex, temporal pole superior gyrus (R)	106	3.24	45	11	-16
Lateral temporal cortex, superior temporal gyrus (L)	433	2.42	-57	-2	-12
Lateral prefrontal cortex, pars orbitalis (L)	157	2.92	-35	56	-8
Lateral prefrontal cortex, pars triangularis (R)	37	2.88	64	38	16
Lateral temporal cortex, temporal pole middle gyrus (R)	54	2.79	-54	16	-26
Retrosplenial cortex (L)	52	2.72	-17	-51	13
Lateral prefrontal cortex, dorsolateral (L)	155	2.72	-23	58	10
Parahippocampal cortex (R)	25	2.69	26	-18	-19
Lateral prefrontal cortex, pars triangularis (R)	148	2.66	48	33	9
Lateral prefrontal corte, pars orbitalis (R)	15	2.62	12	58	-13
Lateral prefrontal cortex, pars triangularis (L)	3	2.60	-62	37	15
Posterior cingulate gyrus (R)	13	2.59	5	-40	10
Lateral temporal cortex, superior temporal gyrus (R)	13	2.58	67	2	2
Lateral temporal cortex, temporal pole superior gyrus (L)	12	2.55	-53	7	-11
Lateral prefrontal cortex, dorsolateral (R)	11	2.50	14	59	22
Thalamus (L)	1	2.41	-15	-20	0
Retrosplenial cortex (R)	1	2.40	2	-44	8
Medial prefrontal cortex, medial superior frontal gyrus (R)	15	2.40	15	48	5
Caudate (L)	5	2.39	-14	25	-5
Cerebellum (R)	8	2.39	10	-63	-24
Lateral temporal cortex, inferior temporal gyrus (L)	1	2.35	-58	-15	-24
Medial prefrontal cortex, medial orbital superior frontal gyrus (R)	1	2.34	11	55	-9

Note: (R) refers to right hemisphere and (L) refers to left hemisphere.

6.4.4. Linked scenes are represented in all layers of the mPFC

Next, I sought to understand whether the perception of scene imagery could be decoded from the multivoxel activity patterns in any of the mPFC layers or hippocampal subfields. Additionally, in an exploratory analysis, I investigated scene decoding in the layers of the other 6 cortical areas that were engaged during animation watching. I divided each of the 7 cortical areas (the mPFC, lateral prefrontal cortex, mLTC, sLTC, retrosplenial cortex, parahippocampal cortex and occipital cortex) into 3 cortical layers: superficial, middle and deep. This yielded a total of 27 cortical layer and hippocampal subfield ROIs.

For each ROI, I used a SVM to classify Scene from Non-scene stimuli based on their multivoxel response patterns. This scene decoding was performed separately for Linked and Unlinked animations. Above chance (50%) classification accuracy, assessed using one-sample t-tests, indicated that real-world scene imagery was in some way represented or processed by the brain area. Differences in scene decoding accuracy between Linked and Unlinked animations were assessed using paired t-tests. Higher accuracy of scene decoding for Linked compared to Unlinked animations would indicate that the brain region preferentially represents or processes real-world dynamic events compared to single scene snapshots.

I found above chance decoding of Linked Scenes in all layers of the mPFC (deep: $p = 0.00603$, middle: $p = 0.0329$, superficial: $p = 0.0181$; Figure 35a). Whereas, Unlinked Scenes were not decoded above chance in any of the mPFC layers (deep: $p = 0.206$, middle: $p = 0.385$, superficial: $p = 0.318$). Despite decoding Linked Scenes but not Unlinked Scenes in the mPFC layers, there were no pairwise differences in decoding accuracy between the Linked and Unlinked animations in any of the layers (deep: 0.22, middle: 0.233, superficial: 0.176).

In the hippocampal subfields, both Unlinked and Linked Scenes were decoded above chance in the DG (Unlinked: $p = 0.0383$, Linked: $p = 0.0174$), CA2/3 (Unlinked: $p = 0.0253$, Linked: $p = 0.00151$) and CA1 (Unlinked: $p = 0.0211$, Linked: $p = 0.0259$; Figure 35a). While only Linked Scenes were decoded above chance in the pre/parasubiculum (Linked: $p = 0.015$, Unlinked: $p = 0.138$). Neither Unlinked nor Linked Scenes were decoded above chance in the subiculum (Unlinked: $p = 0.1$, Linked: $p = 0.175$) and uncus (Unlinked: $p = 0.16$, Linked: $p = 0.142$). As in the mPFC, there were no pairwise differences in scene decoding accuracy between the Linked

and Unlinked animations in any of the hippocampal subfields (DG: $p = 0.989$, CA2/3: $p = 0.716$, CA1: $p = 0.447$, subiculum: $p = 0.543$, pre/parasubiculum: $p = 0.354$, uncus: $p = 0.794$).

In exploratory analyses, I investigated representations of Linked and Unlinked scenes in the other cortical areas in the core network, finding above chance decoding accuracy in many areas (Table 18 for a full report of the results). However, just two of the cortical areas displayed differences in scene decoding accuracy between Linked and Unlinked animations, the retrosplenial cortex deep layer ($p = 0.015$) and the parahippocampal cortex superficial layer ($p = 0.04$; Figure 35b; Table 19 for pairwise comparisons of the other ROIs). In both of these regions, there was more accurate decoding of Unlinked Scenes compared to Linked Scenes, indicating a preference of these brain regions towards processing multiple unrelated scene images compared to extended events.

a

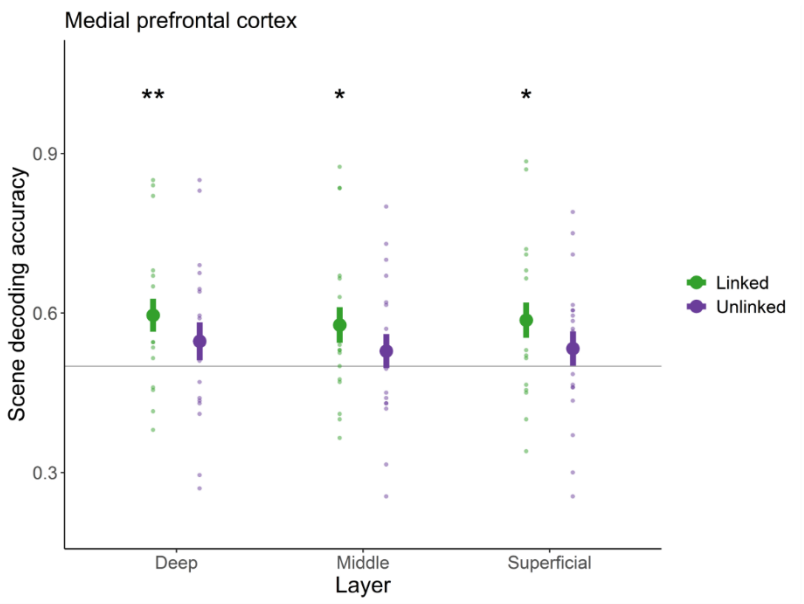
Medial prefrontal cortex
**
*
*
0.9
0.6
0.3
Scene decoding accuracy
Linked
Unlinked
Deep
Middle
Superficial
Layer

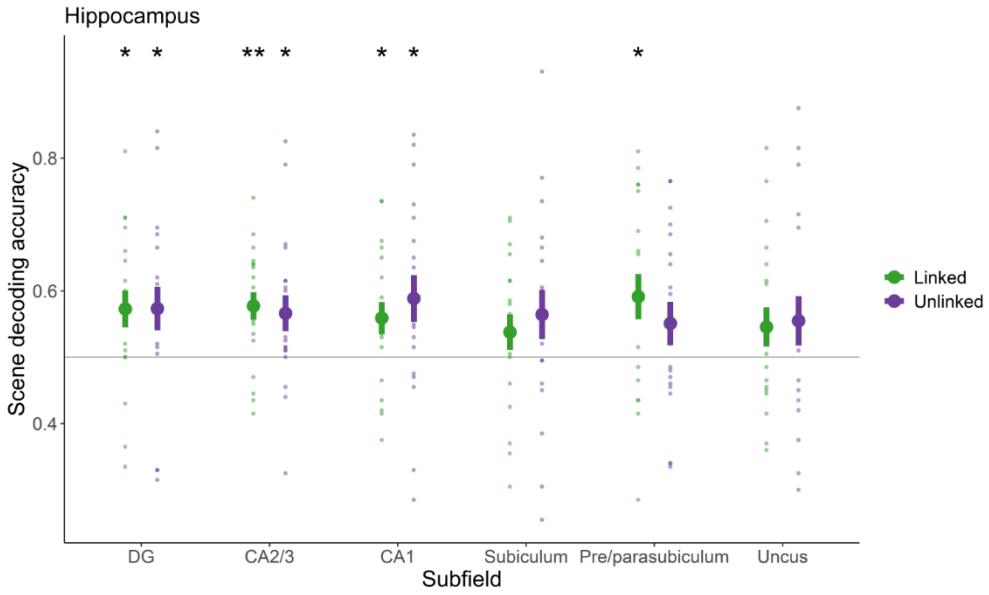
Hippocampus
* *
** *
* *
*
0.8
0.6
0.4
Scene decoding accuracy
Linked
Unlinked
DG
CA2/3
CA1
Subiculum
Pre/parasubiculum
Uncus
Subfield

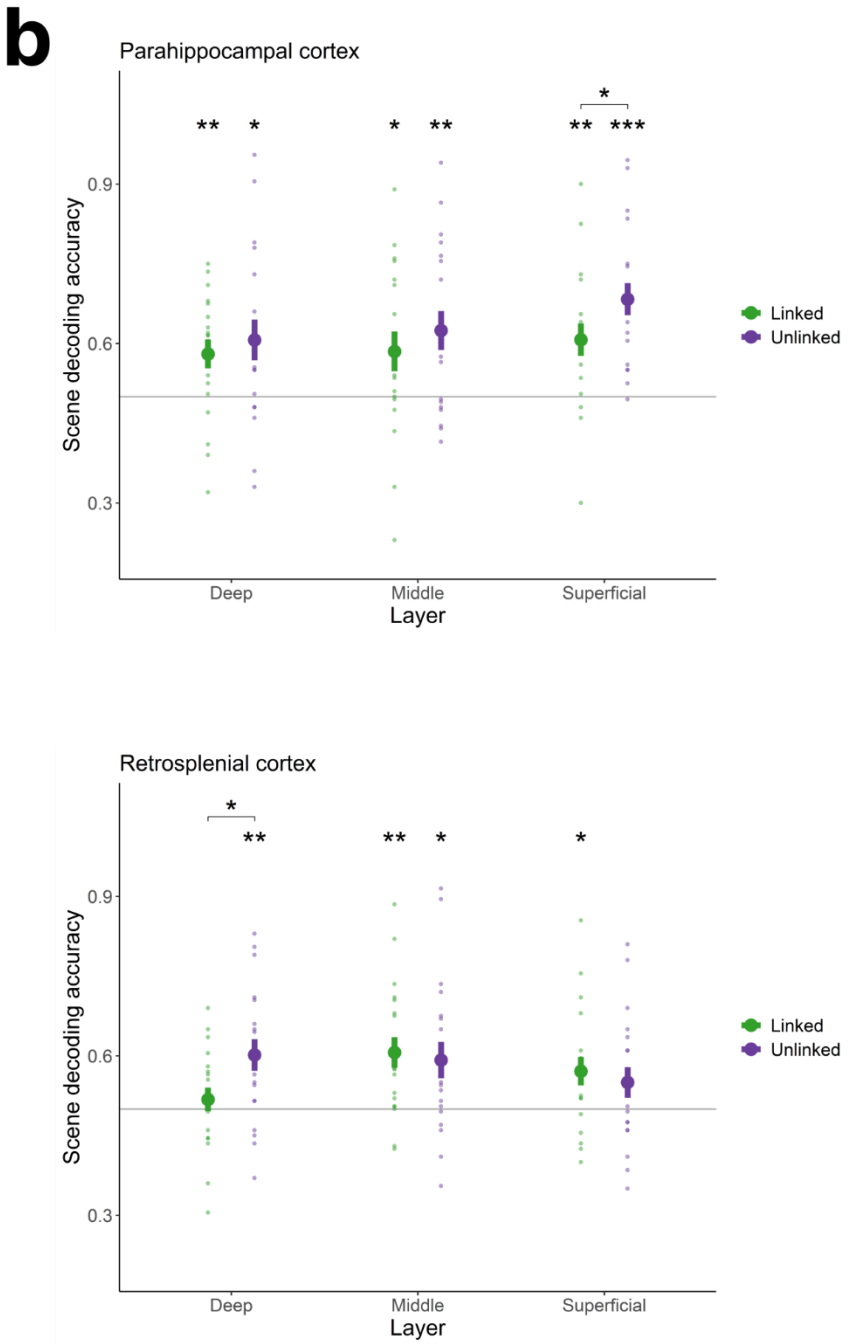

Figure 35. Scene decoding of Linked and Unlinked animations. The accuracies of Linked (green) and Unlinked (purple) scene decoding were compared to one another and to 0.5 (i.e., chance level decoding). (a) In the mPFC layers, there was above chance scene decoding of only Linked animations. In hippocampal DG, CA2/3 and CA1 there was above chance scene decoding of Linked and Unlinked animations and in the pre/parasubiculum there was above chance scene decoding of only Linked animations. (b) Two cortical areas contained layers with scene decoding that was significantly different between Linked and Unlinked animations. These were the retrosplenial cortex deep layer and parahippocampal cortex superficial layer.

Table 18. Scene decoding of Linked and Unlinked animations.

Brain area	Cortical layer	Condition	Scene decoding accuracy		t-statistic	p-value
			mean	SD		
Medial prefrontal cortex	Deep	Linked	0.596	0.134	3.11	**0.00603** **
		Unlinked	0.547	0.156	1.31	0.206
	Middle	Linked	0.577	0.146	2.31	**0.0329** *
		Unlinked	0.528	0.139	0.89	0.385
	Superficial	Linked	0.587	0.145	2.60	**0.0181** *
		Unlinked	0.533	0.141	1.03	0.318
Lateral prefrontal cortex	Deep	Linked	0.569	0.131	2.30	**0.0338** *
		Unlinked	0.574	0.152	2.12	**0.0481** *
	Middle	Linked	0.575	0.141	2.32	**0.0324** *
		Unlinked	0.573	0.141	2.27	**0.036** *
	Superficial	Linked	0.582	0.133	2.69	**0.015** *
		Unlinked	0.580	0.138	2.54	**0.0204** *
Lateral temporal cortex (superior gyrus)	Deep	Linked	0.572	0.112	2.81	**0.0117** *
		Unlinked	0.570	0.142	2.15	**0.045** *
	Middle	Linked	0.579	0.118	2.94	**0.00884** **
		Unlinked	0.562	0.150	1.81	0.0866
	Superficial	Linked	0.594	0.113	3.61	**0.00199** **
		Unlinked	0.572	0.149	2.09	0.0511
Lateral temporal cortex (middle gyrus)	Deep	Linked	0.587	0.103	3.68	**0.00172** **
		Unlinked	0.596	0.145	2.89	**0.00968** **
	Middle	Linked	0.611	0.111	4.33	**0.000403** ***
		Unlinked	0.588	0.141	2.72	**0.0141** *
	Superficial	Linked	0.645	0.102	6.16	**0.00000811** ***
		Unlinked	0.602	0.166	2.68	**0.0153** *
Parahippocampal cortex	Deep	Linked	0.580	0.119	2.93	**0.00886** **
		Unlinked	0.607	0.167	2.78	**0.0125** *
	Middle	Linked	0.585	0.164	2.26	**0.0363** *
		Unlinked	0.624	0.160	3.40	**0.0032** **
	Superficial	Linked	0.607	0.132	3.53	**0.00241** **
		Unlinked	0.683	0.132	6.07	**0.00000978** ***
Retrosplenial cortex	Deep	Linked	0.518	0.097	0.80	0.434
		Unlinked	0.602	0.130	3.39	**0.00323** **
	Middle	Linked	0.606	0.126	3.69	**0.00167** **
		Unlinked	0.592	0.149	2.69	**0.0151** *
	Superficial	Linked	0.571	0.117	2.65	**0.0161** *
		Unlinked	0.550	0.126	1.73	0.1
Occipital cortex	Deep	Linked	0.611	0.122	3.96	**0.000918** ***
		Unlinked	0.598	0.160	2.68	**0.0153** *
	Middle	Linked	0.637	0.131	4.55	**0.000247** ***
		Unlinked	0.627	0.154	3.60	**0.00207** **

	Superficial	Linked	0.677	0.126	6.16	**0.00000816 *** **
		Unlinked	0.651	0.158	4.14	**0.000609 *** **
DG	N/A	Linked	0.573	0.121	2.62	**0.0174 ***
		Unlinked	0.573	0.143	2.24	**0.0383 ***
CA2/3	N/A	Linked	0.577	0.090	3.74	**0.00151 ****
		Unlinked	0.566	0.118	2.44	**0.0253 ***
CA1	N/A	Linked	0.559	0.106	2.43	**0.0259 ***
		Unlinked	0.588	0.153	2.53	**0.0211 ***
Subiculum	N/A	Linked	0.538	0.116	1.41	0.175
		Unlinked	0.564	0.161	1.73	0.1
Pre/parasubiculum	N/A	Linked	0.591	0.148	2.69	**0.015 ***
		Unlinked	0.551	0.143	1.55	0.138
Uncus	N/A	Linked	0.546	0.129	1.54	0.142
		Unlinked	0.555	0.163	1.47	0.16

Table 19. Pairwise comparisons between scene decoding accuracy of Linked and Unlinked accuracy.

Brain area	Cortical layer	t-statistic	p-value
Medial prefrontal cortex	Deep	1.270	0.22
	Middle	1.234	0.233
	Superficial	1.409	0.176
Lateral prefrontal cortex	Deep	-0.140	0.89
	Middle	0.055	0.956
	Superficial	0.059	0.954
Lateral temporal cortex (superior gyrus)	Deep	0.062	0.951
	Middle	0.479	0.638
	Superficial	0.663	0.516
Lateral temporal cortex (middle gyrus)	Deep	-0.306	0.763
	Middle	0.697	0.495
	Superficial	1.250	0.227
Retrosplenial cortex	Deep	-2.698	**0.015 ***
	Middle	0.442	0.663
	Superficial	0.622	0.542
Parahippocampal cortex	Deep	-0.701	0.492
	Middle	-0.847	0.408
	Superficial	-2.213	**0.04 ***
Occipital cortex	Deep	0.334	0.742
	Middle	0.282	0.781
	Superficial	0.687	0.501
DG	N/A	-0.015	0.989
CA2/3	N/A	0.370	0.716
CA1	N/A	-0.778	0.447
Subiculum	N/A	-0.620	0.543
Pre/para	N/A	0.952	0.354
Uncus	N/A	-0.264	0.794

Table 20. Interrater reliability results for hippocampal subfield segmentation. Dice similarity coefficients for the Two researchers independently segmented 4 of 19 participants. All subfields achieved Dice coefficients that were higher than the pre-established thresholds from previous work (Barry et al., 2021; Berron et al., 2017; Bonnici et al., 2013; Chadwick et al., 2014; Clark et al., 2023; Dalton et al., 2019; Lee et al., 2014; Palombo et al., 2013; Yeung et al., 2019).

Subfield	Dice similarity coefficient (mean ± SD)	
	Left	Right
DG	0.87 ± 0.02	0.86 ± 0.01
CA2/3	0.71 ± 0.02	0.66 ± 0.03
CA1	0.79 ± 0.02	0.79 ± 0.02
Subiculum	0.81 ± 0.03	0.79 ± 0.03
Pre/parasubiculum	0.70 ± 0.05	0.71 ± 0.02
Uncus	0.87 ± 0.03	0.84 ± 0.05

6.4.5. Linked scene representations are correlated between mPFC layers and all other cortical areas and hippocampal subfields

I next explored the idea that the scene representations in the mPFC layers and hippocampal subfields are communicated with other brain areas in the core network. If the mPFC or hippocampus was sharing scene-related information with another brain area then I would expect the trial-by-trial scene decoding accuracies of the brain areas to covary (Coutanche & Thompson-Schill, 2013). I examined this by correlating the trial-by-trial scene decoding accuracy of Linked animation trials between each of the mPFC layers and each of the other 26 other ROIs under investigation. I repeated this for the trial-by-trial Linked and Unlinked scene representations of each of the hippocampal subfields. Previous studies have interpreted correlated representation strengths as an indication of information sharing between areas (Aitken & Kok, 2022; Koster et al., 2018; Shao et al., 2023).

I found that all 26 areas had correlated trial-by-trial scene decoding accuracies with all layers of the mPFC (Table 21 for an example of the mPFC deep layer). Furthermore, the trial-by-trial decoding accuracies of all hippocampal subfields were correlated with those of all other brain regions under investigation.

Table 21. Informational connectivity of the medial prefrontal cortex deep layer during Linked animation watching.

Brain area	Cortical layer	Correlation (z)		t-statistic	p-value	Effect size (Cohen's d)
		mean	SD			
Lateral prefrontal cortex	Deep	1.218	0.532	19.95	**1.03E-31** ***	2.29
	Middle	1.169	0.56	18.19	**3.23E-29** ***	2.09
	Superficial	1.039	0.373	24.30	**2.74E-37** ***	2.79
Lateral temporal cortex (middle gyrus)	Deep	0.943	0.37	22.23	**9.65E-35** ***	2.55
	Middle	0.933	0.304	26.80	**3.76E-40** ***	3.07
	Superficial	0.91	0.245	32.41	**7.37E-46** ***	3.72
Lateral temporal cortex (superior gyrus)	Deep	0.976	0.261	30.33	**7.61E-44** ***	3.48
	Middle	1.01	0.217	20.53	**1.67E-32** ***	2.35
	Superficial	0.934	0.323	25.79	**5.03E-39** ***	2.96
Retrosplenial cortex	Deep	0.68	0.33	25.77	**5.28E-39** ***	2.96
	Middle	0.654	0.279	24.85	**6.07E-38** ***	2.85
	Superficial	0.613	0.289	14.30	**3.97E-23** ***	1.64
Parahippocampal cortex	Deep	0.64	0.23	16.89	**2.89E-27** ***	1.94
	Middle	0.583	0.23	18.22	**2.89E-29** ***	2.09
	Superficial	0.562	0.374	16.96	**2.25E-27** ***	1.94
Occipital cortex	Deep	0.907	0.281	15.13	**0.0000000000111** ***	3.47
	Middle	0.846	0.429	17.01	**0.00000000000155** ***	3.90
	Superficial	0.823	0.316	11.10	**0.00000000176** ***	2.55
Dentate gyrus	N/A	0.867	0.305	12.38	**0.000000000307** ***	2.84
CA2/3	N/A	0.729	0.251	12.64	**0.000000000217** ***	2.90
CA1	N/A	0.892	0.367	10.59	**0.00000000366** ***	2.43
Subiculum	N/A	0.894	0.504	7.72	**0.000000404** ***	1.77
Pre/para	N/A	0.661	0.22	13.11	**0.00000000012** ***	3.01
Uncus	N/A	0.744	0.31	10.47	**0.00000000438** ***	2.40

6.5. Discussion

In this experiment, participants watched short movie-like animations depicting either continuous events (Linked Scenes) or static scenes (Unlinked Scenes). Control animations showed either dynamic abstract patterns that evolved over time (Linked Non-scenes) or static abstract patterns (Unlinked Non-scenes). First, I compared the BOLD signal while watching these clips to a baseline counting task to identify a core network of brain regions involved in scene and event processing. This core network included the mPFC, lateral prefrontal cortex, mLTC, sLTC, parahippocampal cortex, retrosplenial cortex, occipital cortex and the hippocampus.

Next, I used MVPA to investigate multivoxel representations of scenes and events in the cortical layers of these brain regions and in the hippocampal subfields. I trained SVM's to decode Scenes from Non-scenes based on the multivoxel response patterns. This scene decoding was performed separately for Linked and Unlinked animations, further allowing me to identify brain regions that differentially processed dynamic events compared to single scenes. Supporting my hypothesis, the mPFC distinguished between Linked Scenes and Linked Non-scenes. This was the case in the deep, middle and superficial mPFC layers, indicating the involvement of both feedforward and feedback signalling to and from the mPFC during the processing of dynamic unfolding events. In the hippocampal DG, CA2/3 and CA1, both Linked and Unlinked Scenes were decodable from Non-scene imagery, suggesting that the hippocampus processes both still scene imagery and unfolding events. In two brain regions in the core network, the parahippocampal cortex superficial layer and the retrosplenial cortex deep layer, I found higher decoding accuracy of Unlinked Scenes compared to Linked Scenes, indicating that these regions are preferentially involved in processing multiple individual scene images compared to a single unfolding event.

I next used the informational connectivity method to examine whether the time course of layer- and subfield-specific multivoxel scene representations in the mPFC or hippocampus were correlated with one another or with any other brain region in the core network, indicating sharing of representation content between the regions (Coutanche & Thompson-Schill, 2013). I found that the time course of event decoding accuracy in the mPFC layers and hippocampal subfields were correlated with those in the layers of all other cortical areas in the network and the hippocampal subfields, indicating network-wide connectivity.

6.5.1. The core network of brain regions supporting scene and event perception

The same core network of brain regions that was active during the retrieval of past autobiographical events (see Experiment 3, Chapter 5) was also active during the real-time perception of present scenes and events. However, the degree of activation within these regions differed between perception and memory. While autobiographical memory retrieval produced large responses in higher-order brain areas such as the retrosplenial, parahippocampal and prefrontal cortices, scene and event perception primarily activated the lower-order visual cortices. This difference may be explained by the source of the initiating signal in each cognitive process. Memory retrieval is endogenously initiated and may

therefore engage, more strongly, these higher-order brain areas to initiate and drive memory retrieval. On the contrary, perception is initiated by exogenous signals, manifesting first in lower-order sensory cortices and causing the strongest responses in these brain regions. The differential activation of sensory cortices during memory and perception might also explain the phenomenological difference between the highly vivid imagery experienced during perception and the weaker visual imagery that is typically experienced during memory recall.

6.5.2. The medial prefrontal cortex and event schema

I found that the mPFC distinguished dynamic events from dynamic non-scene stimuli. Whereas, it failed to distinguish static scenes from static non-scene stimuli. However, it is worth noting that there were no significant differences in the decoding accuracy of dynamic events and static scenes an any of the mPFC layers. Therefore, based on these findings, it would be incorrect to conclude that the mPFC is preferentially involved in processing dynamic events. Rather, my findings support its role in dynamic event processing, but fail to support its involvement in individual scene processing.

As discussed in Chapter 5, the mPFC plays an important role in processing event schema, high-level temporal information about events, such as the expected sequence of actions when doing grocery shopping (e.g., pick up a basket, collect groceries, walk to the checkout till etc.; Audrain & McAndrews, 2022; Guo & Yang, 2020; Preston & Eichenbaum, 2013; Reagh & Ranganath, 2023; van Kesteren et al., 2020). The role of the mPFC in processing dynamic unfolding events may be one of activating these event schemas. These schemas may then provide a template into which each individual image frame can be anticipated, understood and linked.

Despite being unable to decode static scenes from static non-scenes in the mPFC, I found significant univariate activation in the mPFC during processing of unlinked stimuli when compared to the baseline counting task, indicating that the mPFC is in some way involved in processing static images. This accords with several previous MEG studies that found mPFC engagement during the imagination of scene imagery (Barry et al., 2019; Kaplan et al., 2017; Monk, Dalton, et al., 2021). These static scenes may activate other types of schematic information in the mPFC, such as context schema, the typical people, objects and landmarks that might be found within a context. If this were true, each Unlinked Scene animation would

activate many different context schemas, resulting in strong univariate engagement of the mPFC, as observed in this experiment.

6.5.3. Feedforward and feedback signalling in the medial prefrontal cortex

Recall that long-range feedback signalling pathways originate and terminate in the deep cortical layers, whereas feedforward signalling pathways originate in the superficial layers of lower-order cortical areas and target the middle layers of higher-order cortical areas (Felleman & Van Essen, 1991; Markov et al., 2014; Shipp, 2023). In this experiment, Linked Scene animations were decoded in all layers of the mPFC, indicating that perception of events involves both feedforward and feedback signalling. This differs from studies using DCM that found that the higher-order mPFC drives activity in the lower-order hippocampus during scene and event visualisation (see Experiment 3, Chapter 5). These previous DCM studies used paradigms that involved the endogenous initiation of mental imagery, such as imagination and memory retrieval, which involve considerable feedback signalling (Bergmann et al., 2024; Carricarte et al., 2024). By contrast, the online perception of events is predominantly driven by exogenous signals and therefore also involves considerable bottom-up processing via feedforward signalling pathways as well as a feedback component (Bergmann et al., 2024; Carricarte et al., 2024; Hardstone et al., 2021; Kok et al., 2016; Lamme et al., 1998). For example, an external stimulus will cause a first wave of feedforward signalling (Semedo et al., 2022). This feedforward signal will pass upwards through the brain's processing hierarchy until it reaches the higher order brain areas such as the mPFC. Here, event schema representations may be activated, which may then be fed back through the processing hierarchy via feedback connections. These feedback signals may play a role in contextualising or predicting each successive scene as it is processed, allowing faster comprehension of the event. Indeed, we found that participants reported predicting what would happen next during Linked animations.

6.5.4. The hippocampus and scene-based cognition

Several hippocampal subfields, the DG, CA2/3 and CA1, represented both static scenes (Unlinked Scenes) and dynamic events (Linked Scenes). The pre/parasubiculum also represented dynamic events (Linked Scenes), although no single scene representations (Unlinked Scenes) were detectable. Previous research found strong hippocampal activation

during the processing of scene compared to non-scene stimuli (Graham et al., 2010; Hassabis, Kumaran, & Maguire, 2007; Hodgetts et al., 2016; Monk, Dalton, et al., 2021; Zeidman, Mullally, et al., 2015). Furthermore, a previous MEG study using a similar paradigm to that used in this experiment revealed that a distributed set of brain regions preferentially responded to Linked Scene compared to Linked Non-scene animations, with the greatest change in broadband (1–30 Hz) power in the hippocampus (Monk, Barry, et al., 2021). Interestingly, this hippocampal scene-selective effect was only found at the first image frame, supporting the hippocampus's primary involvement in processing individual static scenes.

My findings support previous work that has found that the pre/parasubiculum is particularly involved in scene-based cognition (reviewed in Dalton & Maguire, 2017). It is thought to combine spatial, object, and contextual information to form a cohesive scene (Dalton & Maguire, 2017; Zeidman, Lutti, et al., 2015; Zeidman & Maguire, 2016). This area is part of a wider hippocampal circuit that supports the spatial processing of information. Indeed scene-selective effects have also been found in CA1 (Read et al., 2024) and place cells and spatial view cells have been identified in CA1, CA2, CA3 and the DG (Leutgeb et al., 2007; Oliva et al., 2016; Robertson et al., 1998; Rolls, 2023). The more complex 3D spatial environment present in Scene compared to Non-scene imagery in our paradigm may explain the presence of a scene-selective effect in all of the above mentioned hippocampal subfields.

6.5.5. Scene and event processing in other cortical areas in the core network

In two brain areas, the retrosplenial cortex deep layer and the parahippocampal cortex superficial layer, there was more accurate scene decoding of unrelated individual scenes (Unlinked Scenes) compared to unfolding events (Linked Scenes), indicating that these two regions may have a preference towards processing individual scene imagery. The retrosplenial cortex and parahippocampal cortex are functionally connected to one another and to the hippocampus, and are consistently implicated in in spatial navigation and scene-based cognition (Baldassano et al., 2016; Epstein, 2008; Epstein & Baker, 2019; Park & Chun, 2009; Summerfield et al., 2010; Vann et al., 2009). Previous human and non-human primate studies have identified a specific portion of the posterior parahippocampal cortex, known as the parahippocampal place area, that responds selectively to scenes compared to objects or faces (Bastin et al., 2013; Cukur et al., 2016; Epstein & Kanwisher, 1998; Mormann et al., 2017). Specifically, the parahippocampal place area has been found to represent the perspective-

dependent visuospatial organisation of a scene, including the relationships between objects and landmarks, as well as the more general scene context (Aminoff et al., 2013; Park & Chun, 2009; Vass & Epstein, 2013). Whereas, the retrosplenial cortex has been found to represent the higher-level global properties of a scene, such as relationships between objects that are independent of the perspective as well as its location within the wider extended environment (Kim & Maguire, 2018; Park & Chun, 2009; Vass & Epstein, 2013).

My results are consistent with the roles of the parahippocampal cortex and retrosplenial cortex in scene cognition. Processing of unrelated individual scenes would likely involve higher engagement of the retrosplenial cortex and parahippocampal cortex, due to the high quantity of novel scenes, each with unique contexts, objects and characters, relative to a single unfolding event in which a single context, main character and object were present throughout the entire animation.

The preference for processing Unlinked Scenes was localised to the superficial layers of the parahippocampal cortex, which forms part of the feedforward circuitry and is a primary input to the hippocampus. These feedforward pathways would receive a high exogenous influx of novel scene information relative to the Linked animations. In the retrosplenial cortex, the effect was localised to the deep layer, which suggests a top-down feedback role, though its functionality is currently not understood.

6.5.6. Network-wide informational connectivity

I found network-wide informational connectivity between all mPFC layers, hippocampal subfields and cortical layers of all other brain areas in the core network. This differed from my findings in Experiment 3 (Chapter 5) in which specific brain regions were informationally connected during autobiographical memory retrieval. This difference may be explained by the subtlety of the information being decoded in each experiment. In this experiment, I decoded real-world scenes by differentiating multivoxel response patterns to animations with scene imagery and patterns caused by animations with non-scene abstract patterns. Whereas, in Experiment 3 (Chapter 5), specific autobiographical memories were decoded by comparing multivoxel responses to retrieval of a specific memory with multivoxel responses to retrieval of all other memories. The scene and event decoding investigated in the current experiment reflects a broader processing of real-world scene imagery compared abstract non-scene

imagery Experiment 3 (Chapter 5). This coarse scene-related information likely corresponds to scene and event processing in general. Whereas, the event-specific content investigated in Experiment 3 (Chapter 5) may be considerably subtler to detect, and sharing of this event-specific information between areas may be more limited.

6.5.7. Methodological limitations

There are several key differences between linked and unlinked animations worth considering. First, the temporal contrast differs between the two, with linked animations exhibiting higher similarity in low-level visual features between successive image frames compared to unlinked animations. This could result in greater habituation of visual responses during the watching of linked animations. To mitigate these potential habituation effects, a blank "gap" image was inserted between frames. However, full recovery after visual adaption can last for several minutes (e.g., Rose & Lowe, 1982), so the gap image was unlikely to entirely eliminate these effects. This difference in temporal contrast between linked and unlinked images may therefore account for the differences in visual cortex responses to each experimental condition.

Another potential difference between the linked and unlinked conditions is the cognitive effort required for the task. Before each animation started, participants were informed about the type of animation they were about to view. They were instructed to perceive the images in unlinked animations as entirely unrelated, while they were encouraged to perceive successive images of linked animations as related. In the real world, the brain's natural processing of temporally successive scenes involves linking them as a means of making sense of the world. Consequently, unlinked animations may require more effort to process, as participants must inhibit this innate linking function. I aimed to limit this difference in required effort by conducting extensive training with participants to train them to perceive the images in unlinked animations as unrelated. However, a potential difference in the degree of effort required may remain, which may explain observed differences between linked and unlinked conditions.

Finally, as discussed above, the brain may naturally link temporally successive stimuli to derive meaning from their relationships, such as cause and effect. Consequently, the increased involvement of the mPFC during the perception of linked versus unlinked animations may

reflect the greater engagement of causal processing when interpreting linked scenes. This ability to extract meaning from temporally extended stimuli aligns with the mPFC's role in event schematics — representations of temporal sequences of actions and outcomes that are causally connected. Thus, temporal processing may be just one dimension of the mPFC's broader role in processing event schema, specifically its role in learning and predicting meaningful relationships between actions and outcomes (Alexander & Brown, 2011).

6.5.8. Conclusion

In summary, I found support for a role of the mPFC in processing dynamic unfolding events. Whereas, I found evidence for the involvement of the hippocampus in processing both static scenes and temporally extended events. This supports a proposed hierarchical organisation of the hippocampus and mPFC in event processing, with the hippocampus representing single scenes and the mPFC playing a higher-level role in integrating this scene information over time to represent dynamic events (McCormick, Ciaramelli, et al., 2018). The engagement of all of the deep, middle and superficial layers of the mPFC indicates the involvement of bidirectional (i.e., both feedforward and feedback) signalling during event perception.

7. General Discussion

7.1. Introduction

In this final chapter, I will focus on the two event processing 7T fMRI experiments, discussing the findings in the context of the existing literature. I will refer to the experiment in which I compared remote and recent autobiographical memory retrieval (Chapter 5) as "Event Retrieval". While I will refer to the experiment where I compared the perception of dynamic and static, scene and non-scene animations (Chapter 6) as "Event Perception". First, I will briefly recap the motivation behind the work and the main findings. Following this, I will discuss these findings in terms of hierarchical information processing in the mPFC, hippocampus and wider system of brain regions during event processing. Finally, I will outline some possible future directions for research within this domain.

7.2. Summary of the main results

In my PhD research I aimed to characterise the neural microcircuitry underpinning the retrieval and perception of life events. This research was motivated by evidence to suggest that the mPFC and hippocampus are arranged in a functional hierarchy, with the lower-order hippocampus processing more basic event features such as specific scenes, which are then integrated and abstracted by the higher-order mPFC to produce dynamic events and schema (i.e., abstract prior knowledge structures that have been extracted over multiple repeated events). The anatomical projections between hierarchically organised brain regions have been systematically mapped and, interestingly, they follow a consistent laminar pattern (Felleman & Van Essen, 1991; Markov et al., 2014). Long range feedforward projections originate in the superficial layers (layers 2/3) of lower order areas and target the middle layer (layer 4) of higher order areas. Long range feedback projections originate in the deep layers (layers 5/6) of higher order areas and target the deep layers (layers 5/6) and superficial layers (layers 2/3) of lower order areas. Within the hippocampus, feedforward inputs from the superficial layer of the lower-order entorhinal cortex arrive into a specific set of subfields (the DG, CA3, CA1 and the subiculum) and outputs feeding back to the deep layer of the entorhinal cortex and forward to the middle layer of the mPFC originate in a different set of subfields (CA1, subiculum; Aggleton et al., 2015; Carmichael & Price, 1995; Witter & Amaral, 1991, 2020). The

inputs from the entorhinal cortex to the hippocampus are divided into two separate pathways (Amaral & Witter, 1989; Andersen et al., 1971; Lorente de Nò, 1934; Witter & Amaral, 1991, 2020). The trisynaptic pathway is a series of 3 projections: the entorhinal cortex to DG, then onwards to CA3, and finally synapsing in CA1. Whereas, the monosynaptic pathway, comprises direct connections from layer 3 of the entorhinal cortex to CA1 and the subiculum.

The distinct laminar sites of feedforward and feedback processing pathways between the hippocampus and neocortical areas mean that, by examining the laminar sites of multivoxel representations, event retrieval and perception can be understood in terms of the relative involvement of feedforward and feedback signalling, allowing inferences about the direction of connectivity between brain areas. To this end, I designed two experiments, examining the cortical and hippocampal microcircuitry underpinning event retrieval and perception, and thereby allowing the characterisation of each process in terms of feedforward and feedback processing. Each experiment was designed to probe a potential difference in roles of the mPFC and hippocampus in event processing, which I hypothesised may be evident in the microcircuitry.

The first experiment, Event Retrieval, was motivated by previous research that reported differences in the involvement of the hippocampus and mPFC in the retrieval of autobiographical memories of different ages, with the mPFC more involved as memories become more remote retrieval (Barry et al., 2018; Bonnici et al., 2012; Bonnici & Maguire, 2018). An examination of recent and remote retrieval at the level of cortical layers and hippocampal subfields revealed differential involvement of feedforward and feedback signalling pathways. During retrieval of remote memories, I detected multivoxel representations of individual memories in the deep layers of the mPFC and another cortical area, the mLTC, as well as connectivity with the deep layers of other cortical areas and hippocampal CA2/3. These findings indicate the involvement of feedback signalling from the mPFC to the lower order cortical areas and the hippocampus during retrieval of remote autobiographical memories. By contrast, during retrieval of recent memories, I detected individual memory representations in the lateral temporal cortex middle and superficial layers and connectivity with the middle layer of the mPFC. These findings suggest the involvement of feedforward signalling from the lateral temporal cortex to the mPFC in retrieval of recent memories. Therefore, the direction of information flow between the mPFC, hippocampus and

other areas in the autobiographical memory network during retrieval may shift as memories age.

In the second experiment, Event Perception, I probed the specific roles of the mPFC and hippocampus in event processing. I was inspired by patient and neuroimaging research that suggested the hippocampus may be involved in the construction of single scenes, whereas the mPFC may process the temporal dimension of events to produce extended events that unfold over time (Kaplan et al., 2017; Kurczek et al., 2015). I found that the hippocampal DG, CA2/3 and CA1 distinguished scene from non-scene imagery, whether the stimuli were in the form of static images or temporally extended events, indicating that hippocampus has a preference for processing scenes, irrespective of the time dimension. By contrast, the mPFC layers distinguished scene from non-scene imagery only when stimuli were in the form of temporally extended events, indicating the involvement of the mPFC in processing the temporal dimension of realistic events. The involvement of all mPFC layers indicate the involvement of both feedforward and feedback signalling during the processing of dynamic events.

In this general discussion, I will consider the two main questions that I sought to answer with this work: What are the specific roles of the mPFC and hippocampus in the retrieval and perception of events? How do hippocampal and cortical microcircuits support these roles? I will examine these questions in the context of hierarchical information processing in the mPFC and hippocampus.

7.3. Functional hierarchies of event processing

Scene and event information has been proposed to be processed hierarchically along three dimensions: time, executive control and specificity. First, there is a proposed temporal hierarchy, where the lower-level hippocampus processes single instances in time (i.e., single scenes), which are integrated and linked by the higher-level mPFC to produce events that unfold over longer time scales (McCormick, Ciaramelli, et al., 2018). Second, there is a proposed hierarchy of executive control, where the higher-level mPFC initiates and coordinates the processing of scenes and events in the lower-level hippocampus (McCormick, Ciaramelli, et al., 2018). Finally, there is an abstraction hierarchy, where the lower-level hippocampus processes information pertaining to specific autobiographical events and the higher-level mPFC extracts more general and abstract knowledge across multiple events (i.e.,

event schema; Gilboa & Marlatte, 2017; McClelland et al., 1995; McCormick, Ciaramelli, et al., 2018; Preston & Eichenbaum, 2013; Robin & Moscovitch, 2017).

These three hierarchical relationships have been combined into a single model of scene and event construction (Figure 7; McCormick, Ciaramelli, et al., 2018). In the model, the mPFC internally initiates event construction. The mPFC selects the appropriate scene elements and inhibits irrelevant elements by activating the representations of these elements (e.g., people, objects and landmarks) in neocortical areas (Ghosh et al., 2014; Gilboa & Marlatte, 2017; Preston & Eichenbaum, 2013; van Kesteren et al., 2010; van Kesteren et al., 2012). The scene elements are then passed on to the hippocampus where a spatially-coherent scene is constructed (Hassabis & Maguire, 2007; Maguire & Mullally, 2013). To envisage an extended event, the constructed scene is then passed on to the mPFC and the process repeats for each successive scene of the event. As such, the mPFC is proposed to exert top-down control of the hippocampus and neocortex to initiate and coordinate the construction and linking together of the series of scenes that comprise a full event.

In Sections 7.3.1 to 7.3.3, I will discuss evidence, including from my PhD research, for each of these three proposed functional hierarchies. I will discuss how two of the proposed hierarchies, the temporal and executive hierarchies, do not fully fit with my findings, but the third proposed hierarchy, the abstraction hierarchy, does.

7.3.1. A temporal hierarchy

At the lowest level of a temporal hierarchy are single moments in time, such as a single scene (i.e., a snapshot). At subsequent levels, fast, transient motions may be processed, such as the flicker of a leaf in the wind, which may occur over milliseconds or seconds. At higher levels still are slower events, such as the playing of a tennis match, which may last for several hours.

Neuroimaging studies have found that the hippocampus is particularly engaged when processing single scene snapshots, when compared to other types of stimuli (Brandman & Peelen, 2017; Graham et al., 2010; Hodgetts et al., 2016; Hodgetts et al., 2017; McCormick et al., 2021; Read et al., 2024; Summerfield et al., 2010; Zeidman, Mullally, et al., 2015; but also see Barense et al., 2010). Furthermore, damage to the hippocampus causes impairments in any task involving the representation of scenes, including those involving the visualisation of extended events, without affecting the visualisation of other stimuli, such as objects, faces,

words, or concepts (reviewed in Graham et al., 2010; Lee et al., 2012; McCormick, Ciaramelli, et al., 2018). Unlike the hippocampus, damage to the mPFC does not affect the visualisation of individual scenes, provided specific cues are given (Kurczek et al., 2015), but does severely impair tasks requiring the visualisation of temporally extended events (Bertossi et al., 2017; Bertossi et al., 2016; Della Sala et al., 1993; Kopelman et al., 1999). Furthermore, in healthy individuals, there was a significantly greater coupling in theta phase between the mPFC and hippocampus during the visualisation of dynamic events when compared to static objects or scenes (Kaplan et al., 2017). Together, these findings suggest that the connectivity between the mPFC and hippocampus may facilitate the temporal integration of single scenes into dynamic events (McCormick, Ciaramelli, et al., 2018).

My findings (in the Event Perception experiment) support this work. The hippocampus distinguished scene from non-scene imagery, whether the stimuli were in the form of static images or temporally extended events, indicating that hippocampus has a preference for processing scenes, irrespective of the time dimension. By contrast, the mPFC distinguished scene from non-scene imagery only when stimuli were in the form of temporally extended events, indicating the involvement of the mPFC in processing the temporal dimension of realistic events.

Despite not decoding static scenes from non-scenes in the mPFC using multivariate analyses, univariate analyses of my Event Perception data revealed that the mPFC was still engaged during the viewing of static images. This finding aligns with previous neuroimaging studies that show that the mPFC activates during the imagination of meaningful static scenes (Barry et al., 2019; McCormick et al., 2021; Monk, Dalton, et al., 2021). Furthermore, previous studies using MEG found that neural activity in the hippocampus and the mPFC synchronise during imagination of static scenes, with the mPFC driving activity in the hippocampus (Barry et al., 2019; Kaplan et al., 2017; Monk, Dalton, et al., 2021; but also see Campbell et al., 2018). These findings complicate the notion that the mPFC only acts to temporally bind scene imagery, instead indicating that it may also play a role in atemporal processing.

The proposed temporal hierarchy is further obfuscated by evidence from patient and neuroimaging studies suggesting that, as well as processing atemporal scenes, the hippocampus also plays an important role in the processing of temporal sequences (reviewed in Barry & Maguire, 2024; Davachi & DuBrow, 2015). For instance, damage to the

hippocampus prevented participants from recognising the correct temporal order of previously presented sequences of stimuli, although they could remember the stimuli themselves (Mayes et al., 2001). Similar impairments were observed when hippocampal activity was transiently disrupted using direct electrical stimulation in epileptic patients with surgically implanted electrodes (Goyal et al., 2018). These participants were unable to recall the correct order of previously presented word sequences. A similar impairment is also observed when processing naturalistic event stimuli. Unlike healthy controls, patients with hippocampal damage were unable to recall of the correct order of events that were experienced on a 25 minute walk around a university campus (Dede et al., 2016). Furthermore, in an fMRI study, the hippocampus was active during retrieval of the temporal order of events within a previously watched movie and its degree of engagement correlated with the accuracy of sequence recall (Lehn et al., 2009). Together, these findings indicate an important role for the hippocampus in encoding sequential associations between information that was experienced within the same temporal context, occurring close together in time.

However, the hippocampus appears to be less involved in the processing of information over longer time scales, that span days, months and even years, while the mPFC does appear to be important for this function (Barry & Maguire, 2024). For instance, damage to the mPFC impairs the ability to place autobiographical memories on a timeline that spans many years, while this ability remains intact following hippocampal damage (Tranel & Jones, 2006). Furthermore, mPFC damaged patients display strong temporal discounting, preferring smaller immediate rewards to larger delayed rewards (Peters & D'Esposito, 2016; Sellitto et al., 2010). By contrast, hippocampal damaged patients display temporal discounting in line with healthy controls (Kwan et al., 2013). Therefore, perhaps the mPFC, but not the hippocampus, functions to integrate information over longer periods of time.

Together, these findings suggest that the mPFC and hippocampus may indeed be organised into a temporal hierarchy, with the mPFC processing information that spans across multiple distinct events, perhaps occurring on different days, weeks or years, and the hippocampus processing temporal sequence information, perhaps occurring within a shorter time period such as within a single event.

7.3.2. An executive control hierarchy

A second hierarchical relationship that has been proposed to exist between the mPFC and hippocampus is one of executive control, with the mPFC situated higher than the hippocampus and undertaking a role of initiating and supervising the (re)construction of scenes, and possibly sequences of subevents, by the hippocampus (McCormick, Ciaramelli, et al., 2018).

Studies into patients with bilateral hippocampal or mPFC damage revealed a role for the mPFC in exerting top-down control of the hippocampus. Hippocampal damaged patients seem to have a fundamental issue with any task requiring the visualisation of scenes or events (Bird et al., 2008; Cipolotti et al., 2001; Corkin, 2002; Hassabis, Kumaran, Vann, et al., 2007; Miller et al., 2020; Mullally et al., 2012; Nadel & Moscovitch, 1997; Rosenbaum et al., 2008; Steinvorth et al., 2005; Taylor et al., 2007; Viskontas et al., 2000). Whereas, mPFC damaged patients are able to imagine or recall scenes and events when highly detailed external cues are provided, but they are unable to *spontaneously* visualise scenes and events (Bertossi et al., 2017; Bertossi & Ciaramelli, 2016; Bertossi et al., 2016; Della Sala et al., 1993; Kopelman et al., 1999), indicating a potential issue with the endogenous initiation of event visualisation. Furthermore, unlike hippocampal damaged patients, mPFC damaged patients engage less frequently in day-dreaming, which is characterised by the internal direction of attention towards internal thoughts and visualisations (Bechara, 2004; Bertossi & Ciaramelli, 2016; Leopold et al., 2012; McCormick, Rosenthal, et al., 2018). This led to the idea that the mPFC may play a role in the endogenous initiation of internal experiences (McCormick, Ciaramelli, et al., 2018).

One approach to investigating the influence of one brain area on another in healthy individuals is using effective connectivity analyses, such as DCM, of neuroimaging data. DCM is used to test biologically plausible network models of the effective (i.e., causal) connectivity between regions. DCM analysis of MEG and fMRI data found that activity in the mPFC drives activity in the hippocampus, rather than vice versa, during the recollection of past autobiographical memories, regardless of their age (McCormick et al., 2020; Nawa & Ando, 2019, 2020), as well as during imagination of static scene images (Barry et al., 2019; Monk, Dalton, et al., 2021; but also see Campbell et al., 2018). These findings suggest a top-down influence, presumably via feedback signalling pathways, of the mPFC on the hippocampus during the internal visualisation of scenes and events.

Despite finding that the mPFC drives activity in the hippocampus throughout retrieval, McCormick et al. (2020) found that the mPFC responded earlier than the hippocampus specifically during the retrieval of remote memories (aged between 4 months to 5 years old). By contrast, no difference in response time between the regions was observed for recent memories (less than 1 month old), suggesting that the mPFC may initiate remote but not recent event retrieval. Based on this evidence, it was proposed that the mPFC always drives activity in the hippocampus, and, except perhaps during retrieval of very recent memories where a hippocampal memory trace may still exist, the mPFC initiates the internal visualisation of scenes and events.

In my Event Perception study, I observed the involvement of all layers of the mPFC during the perception of dynamic events, indicating the involvement of both feedforward and feedback signalling. This may be explained by the idea that, unlike internally-driven scene and event visualisation, externally-driven scene and event visualisation may not require the mPFC for initiation and control. Supporting this view, two recent laminar fMRI studies compared the internal imagination and external perception of events in a single study. In the first study, they found that internally-generated imagery activated the deep layers of the parahippocampal and visual cortices, whereas externally-generated imagery produced equal activation across the layers (Carricarte et al., 2024). Similarly, the second study found that imagination activated the deep layers of the primary visual cortex, while perception activated all layers (Bergmann et al., 2024). These findings suggest that, while internally generated imagery may rely, predominantly, on feedback signalling, perception may be supported by both feedforward and feedback signalling between neocortical areas.

Based on this hypothesis, in the Event Retrieval study, I might have expected memory retrieval to engage only the feedback layers, irrespective of the age of the memory being retrieved. I did indeed find evidence for top-down control of the hippocampus by the mPFC during retrieval of remote memories. Specifically, I found evidence for the sharing of memory content from the mPFC to hippocampal CA3 (via other neocortical areas) via feedback anatomical pathways. However, during retrieval of recent memories, I found no evidence of this top-down influence. Instead, I found evidence for the sharing of information from the mLTC to the mPFC via feedforward anatomical pathways.

An important distinction between the DCM studies and my laminar fMRI studies is in the nature of the connectivity being examined. DCM estimates the connectivity between regions based on the time courses of univariate power changes or BOLD signal changes, typically taken from a peak voxel within each region. These univariate measurements correspond to the estimated amplitude of neural activation during a given cognitive process. Whereas, I measured the layer-specific multivariate representations within brain regions and the sharing of this information content between regions based on the time course of these multivoxel representations. These multivoxel representations correspond to distributed patterns of neuronal activity representing specific event content (i.e., individual memories). As such, while DCM may more sensitive to detecting the overall influence of the activity in one area on the activity in another during a given cognitive process, my layer-specific multivariate analyses may be more sensitive to detecting the transfer of informational content between areas. These two processes may be orthogonal. For example, during recent autobiographical memory retrieval, there may be bidirectional connectivity between the hippocampus, mPFC and other cortical areas. This may sum to an overall influence of the mPFC on the hippocampus, perhaps due to its role in initiating and controlling the visualisation of events by the hippocampus. However, the predominant passing of memory-specific content may occur feedforward, from the lateral temporal cortex to the mPFC. Indeed, despite not detecting stable hippocampal and mPFC representations of recent memories, I observed univariate activation of both the mPFC and the hippocampus during recent memory retrieval, indicating the involvement of these regions beyond the passing of memory-specific content.

In sum, while there is evidence from previous DCM studies for a hierarchy of control between the mPFC and hippocampus, it, like the temporal hierarchy, appears to be insufficient for explaining my findings fully.

7.3.3. An abstraction hierarchy

A third hierarchical relationship that has been observed between the hippocampus and mPFC is based on the degree of abstraction of processed information, where the lower order hippocampus has been proposed to process specific autobiographical events, and the higher order mPFC to process generalised information that has been abstracted from multiple individual events (Gilboa & Marlatte, 2017; McClelland et al., 1995; McCormick, Ciaramelli, et al., 2018; Preston & Eichenbaum, 2013; Robin & Moscovitch, 2017).

The complementary learning systems theory suggests that the hippocampus rapidly learns and maintains unique representations of specific events by forming associations between event elements (Kumaran et al., 2016; McClelland et al., 1995). The neocortex, on the other hand, gradually learns the statistical structure of the world, in the form of generalised representations that have been abstracted across multiple events. Perhaps the most abstract, generalised knowledge structures that have been studied are so-called schemas. Many different types of schema have been described, including but not limited to context and event schema (Ghosh & Gilboa, 2014). Context schemas are knowledge structures of associations between objects, people, landmarks that would typically be found within a given spatial context. For example, a restaurant might typically contain tables, chairs and service staff. Event schemas, on the other hand, are knowledge structures of temporal sequences that might occur within a given context. For instance, the typical sequence of subevents when dining in a restaurant includes being escorted to a table, sitting down, reading the menu, selecting a meal, and placing an order. As demonstrated by this example, context and event schema are not distinct from one another and may be linked under a broader "restaurant" schema.

There is considerable evidence supporting a role for the hippocampus in representing specific event content and of the mPFC in representing of more generalised schema (Audrain & McAndrews, 2022; Chadwick et al., 2011; Reagh & Ranganath, 2023; Schlichting et al., 2015). For example, in a paradigm known as associative inference, participants learn associations between pairs of stimuli that share a common element (e.g., AB and BC). Immediately after encoding, multivoxel representations for A and C became more similar in the mPFC, despite never being presented together, indicating the integration of information across "events" (Schlichting et al., 2015). Whereas in the (posterior) hippocampus, neural representations for A and C remained distinct indicating the maintenance of episode-specific representations. In this study, arbitrary relationships between stimuli were learned during the experiment. However, because in the real-world schemas are learned over a lifetime of experiences, to achieve ecological validity, it is important to investigate schema in terms of real-world prior knowledge. A second study sought to do this (Audrain & McAndrews, 2022). While being fMRI scanned, participants retrieved memories of object-scene pairs that were either congruent or incongruent with existing schema. The multivoxel response patterns in the mPFC were similar for schema-congruent, but not incongruent, pairs that shared the same schematic context

(e.g., beach). In the (posterior) hippocampus, on the other hand, multivoxel response patterns remained distinct to the specific scene-object pair, irrespective of its schema-congruency. Together, these studies highlight a role of the mPFC in integrating information across events and the hippocampus in maintaining event-specific information.

Schemas are called upon in the service of two main functions: 1) the top-down provision of prior knowledge to guide the retrieval, imagination, prediction or interpretation of specific events, and 2) the bottom-up integration of novel or recent events into existing knowledge. For example, during the online perception of events, schemas may function in the top-down anticipation and interpretation of unfolding events. This may be particularly helpful when events are ambiguous. Event schemas in particular, which store information about the temporal structure of events, may guide event perception by giving rise to expectations about what will happen next (Ortiz-Tudela et al., 2024; Richmond et al., 2017). These expectations influence our interpretation of sensory information, leading us to perceive events in a way that aligns with our prior knowledge (Knill & Richards, 1996). During the retrieval of past events, schemas have been proposed to guide, in a top-down fashion, the rapid selection of the most plausible event details (Gilboa & Marlatte, 2017; Preston & Eichenbaum, 2013). This has been found to improve long-term memory of events, with subevents, such as individual actions, being remembered better if they are congruent with event schemas (e.g., Bartlett, 1932; Spalding et al., 2015; Tse et al., 2007). The bottom-up integration of newly and recently perceived events into existing schema ensures that schemas which accurately reflect our experiences can continue to perform their important roles in perception and memory retrieval.

The role of the mPFC in processing schema would explain its engagement in both autobiographical memory retrieval (as in my Event Retrieval study), where schema may be called upon to guide retrieval, and during scene and event perception (as in my Event Perception study), where schema may be called upon to guide interpretation and prediction of scenes and events as well as their integration into existing schema. A possible explanation for the differential involvement of feedforward and feedback signalling pathways in the processing of remote, recent and present moment events observed in my experiments may correspond to the differential involvement of the two major schema functions in each process: the top-down guidance of retrieval and perception, presumably operating via feedback

pathways, and the bottom up (re)integration of information into schema, presumably operating via feedforward pathways.

In the Event Perception experiment, realistic imagery was decodable from abstract pattern imagery in the mPFC when in the form of temporally extended events. Whereas, when static images were presented, the mPFC did not distinguish between meaningful and abstract imagery. One possible explanation for this is that real-world temporal sequences engage mPFC-based schema to a greater degree than real-world atemporal information. Indeed, static scene imagery would likely activate context schema in the mPFC. Whereas, dynamic meaningful events may additionally activate event schema.

Supporting this view, the mPFC has been found to be particularly involved in the representation of event schema when compared to context schema (Baldassano et al., 2018). In another study, participants watched movies or listened to audio narratives that took place in one of two schematic contexts (a restaurant or an airport). For half of the participants, the events were congruent with pre-existing event schema, unfolding in the typical temporal sequence (e.g., entering the restaurant, being seated, ordering, and eating food). For the other half of participants, the same events were incongruent with existing event schema, unfolding in a scrambled order. Multivoxel representations of the contexts were detected in the mPFC, but only when the events were shown in the correct order. This research indicates that the mPFC may be particularly involved in representing the typical temporal sequences that form event schema and less involved in the representation of the typical scene elements that comprise context schema. Not only would this explain my findings, but it would also explain why patients with mPFC lesions are impaired at recalling extended events but are able to recollect individual scenes (Bertossi et al., 2017; Bertossi et al., 2016; Della Sala et al., 1993; Kurczek et al., 2015), and why the mPFC-hippocampus coupling is stronger during dynamic compared to static scene imagination (Kaplan et al., 2017).

In terms of the laminar microcircuitry, in my Event Perception experiment, I found that dynamic real-world events were decodable from dynamic abstract patterns in all mPFC layers. This indicates that, dynamic events and dynamic abstract patterns may differ in terms of their reliance on schema for both top-down guidance, operating via feedback signalling pathways, and bottom-up integration, operating via feedforward signalling pathways. When viewing dynamic events and dynamic abstract patterns, participants reported predicting ahead to

what might happen next. This function of predicting ahead may rely heavily on top-down processing by existing event schema, particularly when real-world meaningful events are perceived. By contrast, the function of predicting ahead during the perception of dynamic abstract patterns may rely more on lower-order brain areas that process simple motions such as rotation. This predictive function would be accompanied by the bottom-up integration of information, perhaps in the form of prediction errors, into event schemas.

In my Event Retrieval study, specific remote, but not recent, memories were represented in the mPFC during their retrieval, corroborating the findings of previous 3T fMRI studies, and supporting the idea that the mPFC is a major neocortical consolidation site (Barry et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Bonnici & Maguire, 2018). It is widely believed that, rather than being directly transferred into the neocortex, new memories are integrated into pre-existing schemas (McClelland et al., 1995; McKenzie & Eichenbaum, 2011; Tse et al., 2007; Wang & Morris, 2010). Supporting this view, recent fMRI studies found that memory representations that were detected in the mPFC at 3 to 7 days post-encoding were not memory-specific but instead overlapped with other events with shared features, such as a shared schematic context (Audrain & McAndrews, 2022; Tompary & Davachi, 2017). This coincided with a decreased memory for the specific event details, indicating that the process of schema consolidation may have resulted in more generalised memories (Audrain & McAndrews, 2022).

My results may therefore be explained by the notion that retrieval of remote autobiographical events requires more top-down schema guidance, operating via feedback pathways, than retrieval of recent memories. This would explain why, during retrieval of remote but not recent memories, I found memory representations in the mPFC deep layers and sharing of these representations (ultimately) with hippocampal CA3 (although passing through the retrosplenial cortex deep layers and mLTC deep layers first). By contrast, retrieval of recent autobiographical events may require more feedforward integration into existing schema than retrieval of remote memories. Indeed, I found that recent episodic memories were not yet stably represented in the mPFC and may therefore be in the process of consolidation into mPFC schemas. This process of schema integration may involve the feedforward passing of event information into the mPFC and may explain why recent memory representations in the

lateral temporal cortex feedforward layers were shared with the mPFC feedforward layers during retrieval.

Supporting the idea that schema integration occurs during memory retrieval, systems-level consolidation is believed to occur, not just during rest and sleep, but also during memory retrieval (Antony et al., 2017). During rest and sleep, the same sequences of neuronal firing that took place during wakefulness have been found to be replayed in the hippocampus and prefrontal cortex (Colgin, 2011; Ji & Wilson, 2007; Siapas et al., 2005). Blocking these replay events impairs memory performance in rodents (Jadhav et al., 2012), highlighting their critical role in memory consolidation. In a similar fashion to offline replay, memory retrieval activates the same hippocampal and neocortical firing patterns that took place during encoding and also improves memory performance (i.e., testing effects; Rowland, 2014). This has led to the proposition that, like sleep and rest, memory retrieval also functions to consolidate memories.

While I detected remote representations in the mPFC, I did not detect representations of memories of any age in the hippocampus, challenging the idea that it represents specific events. However, I did find univariate activation of the hippocampus during retrieval of memories of all ages, supporting its long-term involvement in memory retrieval (Barry & Maguire, 2019; Moscovitch et al., 2016; Nadel et al., 2007). Scene Construction Theory posits that memories are stored in the hippocampus for a short period, specifically, until they are replaced by new memory traces. Following this, the hippocampus continues to be involved in memory retrieval, but, instead of activating a stored memory trace, reconstructs a coherent episodic memory from a combination of neocortical memory traces (Barry & Maguire, 2019). It may be that the recent (1 – 6 week old) memories examined in my study had already undergone this kind of degradation in the hippocampus. Indeed, enough time had passed for the recent memories to undergo some neocortical consolidation, as they were stably represented in the mLTC, suggesting that a hippocampal trace may no longer be required. Therefore, perhaps, as proposed by Barry and Maguire, 2019, in my Event Retrieval study, the hippocampus was functioning to flexibly reconstruct both recent and remote memories.

An alternative perspective, Trace Transformation Theory, posits that various types of memory representations can coexist, including detailed, episode-specific representations in the hippocampus (with coarse, gist-like representations in the anterior hippocampus and fine-grained, detailed representations in the posterior hippocampus) and more generalised,

schematic event representations in the mPFC (Robin & Moscovitch, 2017; Sekeres et al., 2018). According to this theory, hippocampal traces are in the form of indexes, which point to the relevant neocortical memory traces, including schematic, semantic and perceptual details. The involvement of each type of representation is thought to depend on task demands, such as the level of detail required during recollection and the nature of the memory cue provided. For example, this theory posits that hippocampal representations will only engage when episode-specific information is retrieved. The cue is also important as it is proposed to determine which brain region initiates the retrieval process. When generic cues are used, schematic representations in the mPFC will initiate and guide the retrieval process, whereas when episode-specific cues are given the hippocampus will control retrieval.

In my Event Retrieval study, I ensured that highly detailed memories were selected and that recent and remote memories were matched in terms of the number of episode-specific details. Furthermore, the task demands were the same for recent and remote retrieval. A reasonably coarse, but episode-specific, memory cue was provided and participants were tasked with recollecting the specific event with a high degree of detail and vividness. However, despite this, I did not detect stable hippocampal representations, nor did I find evidence for them initiating and driving the engagement of mPFC representations. Instead, I found evidence for the opposite effect during remote memory retrieval, that is, mPFC representations driving the hippocampus via feedback signalling pathways. Therefore, my findings do not support the predictions made by Trace Transformation Theory.

In summary, the arrangement of the mPFC and hippocampus into a functional hierarchy of abstraction, with the lower-order hippocampus processing specific events and the mPFC processing knowledge abstracted across multiple events, may explain why the mPFC is more involved in the processing of temporally unfolding events compared to static images (Bertossi et al., 2017; Bertossi et al., 2016; Della Sala et al., 1993; Kaplan et al., 2017; Kurczek et al., 2015). This is because the mPFC may be particularly involved in representing the typical temporal structure of commonly experienced events (event schema), which may act as a template under which specific events unfold. Furthermore, it may explain why I found remote, schema-consolidated, memories, were stably represented in the mPFC, while recent memories were not.

7.3.4. Other cortical areas

While the focus of this thesis is on the microcircuitry of the hippocampus and mPFC in support of event processing, these brain regions clearly do not work in isolation. Rather, there is a widely distributed set of brain areas involved (Beaty et al., 2018; Bonnici, Chadwick, Lutti, et al., 2012; Hassabis & Maguire, 2007; Maguire, 2001; McDermott et al., 2009; Schacter et al., 2007). In my Event Perception study, I found that, during perception, all of the areas in the core network displayed scene-selectivity during scene and event perception and were functionally connected to one another. This indicates the network-wide preference of the core network of brain regions for scene compared to non-scene imagery and its strong functional connectivity during their perception.

In my Event Retrieval study, one cortical area, the mLTC, contained recent and remote memory representations. Furthermore, several cortical areas were informationally connected to the mLTC or mPFC, which both contained memory representations. The specific informational connectivity between brain areas during memory retrieval led to the creation of a connectivity diagram that could explain an indirect pathway of connectivity between the mPFC and the hippocampus. It revealed two brain areas, the retrosplenial cortex and mLTC that memory-related information may pass through as the mPFC and hippocampus communicate.

Along with the mPFC, the mLTC has also been shown to engage during schema-guided retrieval (Webb et al., 2016). For example, in one study, participants were presented with scenes (bathroom and farm) containing schema-congruent and incongruent objects (Webb et al., 2016). They subsequently underwent a recognition memory test while being fMRI scanned. During the recognition memory test they were presented with the schema congruent and incongruent objects from the encoded scenes as well as schema-congruent lure objects that they had not seen before. The mPFC and lateral temporal cortices coactivated whenever a schema-congruent object was reported as being seen before, whether the memory was accurate or not. In contrast, the accurate recognition of memories coincided with increased activity in the hippocampus and posterior visual cortex.

A recent study probed the specific type of event information that is represented by the mLTC as well as other cortical regions and the hippocampus (Reagh & Ranganath, 2023). While being fMRI scanned, participants watched and then recollected short events that were either

overlapping in terms of context schema (café or grocery store), place (specific café or specific grocery store) or person (specific woman or specific man). Supporting previous studies, the study found that, during both encoding and retrieval, the mPFC represented schematic context, while the hippocampus represented specific events. The mLTC represented people but did not distinguish between locations. By contrast, the retrosplenial cortex represented the specific location of the event, but did not distinguish between people. It is unclear from this study if the people and place representations were semantic or perceptual in nature, although the lateral temporal cortex in particular is known to play an important role in the retrieval of the semantic gist of events (Buckner, 1996; Gabrieli et al., 1998; Noppeney et al., 2007; Price, 2000; Simons et al., 2005). What is clear is that these kinds of event details would be positioned between schematic and episode-specific representations in a hierarchy of abstraction. For example, beaches in general are an example of a context schema, perhaps represented in the mPFC. Three Cliffs Bay (a beach that I used to go to with my family) is an example of a place, perhaps represented in the retrosplenial cortex. A specific trip to Three Cliffs Bay when I was 12 years old is an example of a specific episode, perhaps represented by the hippocampus.

These people and place event details may be precisely the type of event details that schemas are proposed to select during memory retrieval, perhaps explaining the feedback passing of information from the mPFC to the mLTC and retrosplenial cortex during retrieval of remote memories (in my Event Retrieval study). They may also be the type of information that is consolidated into mPFC schema representations, perhaps explaining the feedforward signalling from the lateral temporal cortex to the mPFC during retrieval of recent memories.

Another prominent model of event processing, the “PMAT” framework (Ranganath & Ritchey, 2012), proposes a structural and functional division of the core network. According to this framework, the anterior temporal (AT) system – including the anterior lateral temporal cortex, lateral prefrontal cortex, and amygdala – processes semantic and conceptual event information and its affective significance. In contrast, the posterior medial (PM) system – including the parahippocampal cortex, retrosplenial cortex, and the pre/parasubiculum – processes the spatial, temporal, and social context of events. Hippocampal CA3 and the DG, along with the mPFC, are thought to mediate interactions between these networks for successful event retrieval (Ritchey et al., 2015).

My findings support the idea that the mPFC integrates information from both systems. During event retrieval, the mPFC was functionally connected to an area in the AT system, the mLTC, as well as areas in the PM system, the retrosplenial cortex and occipital cortex. However, I did not find evidence for the convergence of AT and PM systems in the hippocampus. Hippocampal CA3 was exclusively connected to the AT system, specifically the mLTC. Moreover, my findings did not support a strict division between the PM and AT systems. Instead, I observed direct functional connectivity between the mLTC (AT system) and the retrosplenial cortex and parahippocampal cortex (PM system) during memory retrieval.

In summary, there is considerable evidence to suggest that the hippocampus and mPFC are hierarchically organised during event processing. Furthermore, there is evidence to suggest that lower-order cortical areas, such as the mLTC and retrosplenial cortex may sit between the mPFC and hippocampus in a functional hierarchy of abstraction, representing the semantic and perceptual details of events, such as people and places.

7.3.5. The hippocampal microcircuitry

In my Event Retrieval study, there was evidence for sharing of information content between the neocortex and hippocampal CA2/3 during the retrieval of remote autobiographical memories. By contrast, I did not find evidence for the sharing of memory content with any hippocampal subfields during the retrieval of recent memories. This suggests that CA3 might be more involved in retrieval as memories age. However, my analysis of univariate activation in the hippocampal subfields revealed an opposite effect. Both the CA3 and the DG produced greater activation during the retrieval of recent compared to remote autobiographical memories. These univariate results instead suggest that the trisynaptic loop, consisting of projections from the entorhinal cortex to the DG, then to CA3, and finally to CA1, become less involved in retrieval as memories age.

The subfields of the trisynaptic loop, the DG and CA3, support two important computations, pattern separation and completion (Baker et al., 2016; Berron et al., 2016; Guzman et al., 2016; Leutgeb et al., 2007; Lisman, 1999; McNaughton & Morris, 1987; Rebola et al., 2017; Rolls & Treves, 1994; Treves & Rolls, 1994). Neocortical information arrives into the DG, where it is encoded by sparsely distributed ensembles of neurons (pattern separation). This allows the incoming information to be encoded as distinct from past overlapping experiences. The

distinct DG representations are projected onwards to CA3, which, given its dense recurrent collateral connections, is able to access previously encoded information that overlaps with the input (pattern completion). As such, complete memories can be retrieved based on partial cues. These two computations are thought to be essential for memory function.

Evidence from rodent studies suggests that, while retrieval of recent contextual fear memories is supported by both CA3 and CA1, CA3 is not necessary for the retrieval of remote memories (Denny et al., 2014; Lux et al., 2016). Retrieval of these older memories was, instead found to solely rely on CA1 (Lux et al., 2016). In line with Standard Consolidation Theory and Scene Construction Theory, it was proposed that, over time, hippocampal memory representations may degrade, making it difficult to successfully retrieve full memory representations via pattern completion (Kesner & Rolls, 2015; Lux et al., 2016). This could explain why, in my Event Retrieval study, the DG and CA3 were more engaged during the retrieval of recent compared to remote memories.

Speaking against this hypothesis, focal damage to the human CA3 results in impoverished episodic details of memories of all ages (Miller et al., 2020). An fMRI study has also found stronger representations of remote autobiographical memories in CA3 when compared to recent memories (Bonnici et al., 2013). Supporting these findings, an alternative proposition has been made that CA3's role in memory retrieval may instead increase as memories age, due to a greater need to pattern complete memories as cues arriving from the neocortex are likely to be highly vague or degraded (Bonnici et al., 2013). While I did not find evidence for representations in CA3, I did find sharing of information content with CA3 for remote and not recent memories, which could reflect a greater demand for pattern completion by the retrieval of these older memories.

Given that pattern completion is believed to access stored memory representations, it implies the existence of stable representations of the complete memory in CA3. However, I did not find such stable representations in any of the hippocampal subfields. This challenges the idea that the hippocampus pattern completes previously stored representations during retrieval, instead supporting the view that the hippocampus flexibly reconstructs memories by recruiting different neuronal populations on each retrieval. The reconstruction theory does not, however, specify how hippocampal computations might support the proposed reconstruction of events from individual details.

Clearly, there are conflicting views on the involvement of the trisynaptic loop, the DG and CA3 subfields, in retrieval over time. I hope that future studies using 7T fMRI to investigate human hippocampal subfields will help to resolve these debates. The variability in findings suggest that any shift in the reliance of memory retrieval on CA3 over time may be highly nuanced, rather than an all-or-nothing phenomenon. The conflicting findings also provide further evidence for the importance of conducting both multivariate and univariate investigations of cognitive processes. Each approach can provide insights about different types of information processing and together offer a more fine-tuned understanding of cognitive processing.

7.4. Future directions and conclusion

My PhD work provides novel insights into the cortical and hippocampal microcircuits supporting event processing, including memory of past autobiographical events and perception of present moment events. Relatively few studies have investigated the anatomical pathways of connectivity between the hippocampus, mPFC and other cortical areas that support these important functions. My PhD research aimed to address this knowledge gap. First, I developed, with the Department's Physics Team, MRI acquisition and analysis methods that would allow me to detect neuronal activity in the mPFC layers and hippocampal subfields. As part of this methods development, I produced a publicly available, open source tool for automatically delineating the hippocampal subfields in MRI scans. Subsequently, I applied these methods to distinguish the roles of the mPFC layers and hippocampal subfields in recent and remote autobiographical memory retrieval as well as their respective roles in the processing of new events.

I found that, during Event Retrieval, remote but not recent autobiographical memories were stably represented in the mPFC deep layers. The hippocampal subfields, on the other hand, while actively engaged during retrieval, did not produce stable representations of recent or remote autobiographical memories. One other cortical area, the mLTC, did contain stable representations of both remote and recent memories in the deep and superficial layers, respectively. During perception, several hippocampal subfields (the DG, CA3 and CA1) distinguished between scene and non-scene imagery irrespective of whether they were in the form of individual atemporal scenes or temporally extended events. The mPFC layers, on the

other hand, only distinguished between scene and non-scene imagery when in the form of dynamic events.

Overall, my results are consistent with the idea that the hippocampus and mPFC are organised into a temporal hierarchy, with the higher-level mPFC representing information extracted over long time scales and the lower-order hippocampus flexibly (re)constructing individual scenes, and perhaps short temporal sequences. While I did not formally test the specificity of information represented by each brain area, my results are also consistent with the idea that the hippocampus and mPFC may be organised into an abstraction hierarchy, with the higher-level mPFC representing schematic information abstracted over multiple episodes and the hippocampus flexibly (re)constructing specific episodes. These co-existing hierarchies can be unified under the rubric that the mPFC plays a particularly important role in representing event schema, the typical temporal sequence of actions or subevents that comprise an extended event, and the hippocampus in processing episode-specific scenes and short temporal sequences.

This work provides first insights into the cortical microcircuitry involved in naturalistic event processing. The involvement of different cortical layers tells us about the direction of connectivity, feedforward or feedback, therefore contributing to our knowledge about the connectivity between brain regions. I propose that remote, schema-consolidated, memory representations in the mPFC guide retrieval via feedback signalling pathways, perhaps by selecting relevant and suppressing irrelevant event details in lower order areas. Whereas, recent memory representations in the mLTC may undergo more feedforward consolidation into the mPFC during retrieval. Event perception may involve an equal balance of schema-guided perception, for example the prediction and interpretation of events as they are unfolding, and the assimilation of novel events into existing schema.

This work also adds to the knowledge of hippocampal microcircuitry involved in event processing. I found that several hippocampal circuits, including subfields of the input pathways to the hippocampus, the trisynaptic loop (the DG, CA2/3 and CA1), and the monosynaptic pathway (CA1), and subfields in the output pathways from the hippocampus (CA1 and the pre/parasubiculum), have a preference for scene over non-scene imagery. However, their involvement in remote compared to recent memory retrieval was unclear, with univariate activations implying greater engagement of the trisynaptic loop in recent compared

to remote retrieval, but multivariate analyses finding that remote memory content was passed into the trisynaptic loop. This shows how important it is to do studies of univariate activation and content-specific multivariate studies. Multivariate laminar fMRI combined with informational connectivity analysis is unique in providing a method that allows for inferences about the direction of the passing of information content between brain regions.

Clearly, numerous questions remain. For a start, what type of memory content is represented by the mPFC deep layers and the deep layers of lower order cortical areas during event retrieval? While I suspected their involvement, I did not formally test how schemas underpin the retrieval of autobiographical memories and the online perception of events. Future studies could include events that are overlapping or distinct in terms of context schemas (e.g., airports or grocery stores) or event schemas (e.g., the typical sequence of subevents when moving through an airport or grocery store). Representations of other event features, such as specific people, places, objects or actions, could also be probed by including events are overlapping or distinct along these dimensions.

It would also be interesting to track mPFC and hippocampal representations of the same events over time to understand if, at any stage, the hippocampus stably represents specific events and at what point they become stably represented in neocortical areas, including the mPFC. One study performed this kind of analysis of autobiographical memory representations in the mPFC, finding that they were not stably represented at 2 weeks after their initial encoding, but they were after 4 months (Barry et al., 2018). Future studies could investigate this further by measuring both hippocampal and mPFC event representations from the time of their encoding. MRI-compatible VR could be used to expose participants to naturalistic events while fMRI scanning. Events could then be retrieved at various different intervals following encoding, such as immediately after, 1 day later and several months later. The use of VR to present naturalistic events in a laboratory setting, as opposed to real-world autobiographical memories, would also allow the comparison of schema-congruent and schema-incongruent false details between remote and recent memories, as a ground truth would be known.

Related to this, there is the question of whether the hippocampus stores lasting representations of detailed memories, as proposed by Trace Transformation Theory, or if, after a certain amount of time, the hippocampus flexibly reconstructs memories from neocortical

representations, as proposed by Barry and Maguire (2019). My findings of hippocampal activation during retrieval, but the absence of multivoxel representations of memories, support the idea that the hippocampus flexibly reconstructs memories ranging from 1 week to 5 years old. However, as discussed, there are conflicting results in both the rodent and human literature on this topic. Therefore, further human studies of the hippocampal subfields are required to resolve this question with respect to humans. I hope that the ASHS atlas that I developed as part of my PhD will help to make these kinds of studies more accessible.

Another important future study to test the predictions made by Trace Transformation Theory would be to understand the impact of cue specificity on hippocampal-neocortical microcircuitry supporting retrieval. Trace Transformation Theory proposes that the mPFC drives retrieval when schematic cues are used, whereas the anterior and posterior hippocampus drive retrieval when gist-like or detailed episode specific cues are used, respectively. This could be tested by manipulating the specificity of memory cues used in an autobiographical memory retrieval study and observing the impact on hippocampal and mPFC representations.

There is evidence to suggest that the integration of new information into schematic representations can be achieved during its initial encoding when the new information is schema congruent (Wang & Morris, 2010). Therefore, an interesting approach to test the hypothesis that schema integration occurs feedforward and schema-guided retrieval occurs feedback could be an experiment in which participants are presented with schema-congruent and incongruent events. In a similar approach to Audrain and McAndrews (2022), but using laminar and subfield fMRI, the events could fall into two schematic contexts (e.g., beaches and offices). I would predict that the mPFC middle layers and/or superficial layers would represent the schematic context during encoding of schema-congruent but not incongruent events, indicating the feedforward passing of new schema-congruent information into existing schema. Whereas, during retrieval, I would expect the mPFC deep layers to represent schematic context, again, during schema-congruent but not incongruent retrieval, indicating the involvement of feedback pathways in schema-guided retrieval.

In terms of the hippocampal subfield microcircuitry, it would be interesting to further explore the idea that the trisynaptic loop may be more engaged during retrieval of recent memories. A recent study has shown that it is possible to perform laminar analysis of the hippocampal

subfields (Pfaffenrot et al., 2024). Therefore, using my same dataset, a laminar subfield analysis of CA1 could be performed that would distinguish between the two pathways. Alternatively, new fMRI data could be acquired that is optimised for imaging the entorhinal cortex to distinguish between these pathways based on involvement of the superficial (layer 2) or middle layer (layer 3), which form part of the trisynaptic loop and monosynaptic pathway, respectively.

Neurons in layer 2 of the entorhinal and transentorhinal cortices are the first neurons in the brain to be affected by Alzheimer's pathology, developing neurofibrillary tangles, which leads to neurodegeneration (Braak & Braak, 1991). To understand this pathology better, laminar fMRI could be deployed in patients with mild cognitive impairment, a condition that can lead to Alzheimer's disease, to see if there are differences in the engagement of EC layer 2 and the subfields of the trisynaptic loop during memory encoding and retrieval compared to control participants and whether these differences predict differences in memory ability. If functional changes did exist, they could be a potential early biomarker for Alzheimer's risk. It would also be interesting to investigate the specific effects of any functional differences on hippocampal computations like pattern separation and pattern completion, which depend on the trisynaptic pathway.

In conclusion, my PhD research provides new insights into how the hippocampal subfields, mPFC layers and layers of other cortical areas in the core autobiographical memory network cooperate to support the processing of events. My findings highlight how various forms of event processing can differentially engage feedforward and feedback signalling pathways. Clearly, there is still much to learn and I hope that my findings will inspire future studies. Any future studies could make use of the acquisition and analysis tools that I developed as part of this PhD. Particularly, I hope that the work I undertook to develop a tool for automatically defining the hippocampal subfields on MRI scans will remove any potential barriers for future researchers interested in the functions of the human hippocampal microcircuits.

References

Abraham A. (2013). The world according to me: personal relevance and the medial prefrontal cortex. *Frontiers in human neuroscience*, *7*, 341. https://doi.org/10.3389/fnhum.2013.00341

Addis, D. R., Moscovitch, M., & McAndrews, M. P. (2007). Consequences of hippocampal damage across the autobiographical memory network in left temporal lobe epilepsy. *Brain*, *130*(Pt 9), 2327-2342. https://doi.org/10.1093/brain/awm166

Addis, D. R., Pan, L., Vu, M. A., Laiser, N., & Schacter, D. L. (2009). Constructive episodic simulation of the future and the past: distinct subsystems of a core brain network mediate imagining and remembering. *Neuropsychologia*, *47*(11), 2222-2238. https://doi.org/10.1016/j.neuropsychologia.2008.10.026

Addis, D. R., & Schacter, D. L. (2011). The hippocampus and imagining the future: where do we stand? *Front Hum Neurosci*, *5*, 173. https://doi.org/10.3389/fnhum.2011.00173

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, *45*(7), 1363-1377. https://doi.org/10.1016/j.neuropsychologia.2006.10.016

Aggleton, J. P., Wright, N. F., Rosene, D. L., & Saunders, R. C. (2015). Complementary Patterns of Direct Amygdala and Hippocampal Projections to the Macaque Prefrontal Cortex. *Cereb Cortex*, *25*(11), 4351-4373. https://doi.org/10.1093/cercor/bhv019

Aitken, F., & Kok, P. (2022). Hippocampal representations switch from errors to predictions during acquisition of predictive associations. *Nat Commun*, *13*(1), 3294. https://doi.org/10.1038/s41467-022-31040-w

Aitken, F., Menelaou, G., Warrington, O., Koolschijn, R. S., Corbin, N., Callaghan, M. F., & Kok, P. (2020). Prior expectations evoke stimulus-specific activity in the deep layers of the primary visual cortex. *PLoS Biol*, *18*(12), e3001023. https://doi.org/10.1371/journal.pbio.3001023

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature neuroscience*, *14*(10), 1338–1344. https://doi.org/10.1038/nn.2921

Aly, M., Ranganath, C., & Yonelinas, A. P. (2013). Detecting changes in scenes: the hippocampus is critical for strength-based perception. *Neuron*, *78*(6), 1127-1137. https://doi.org/10.1016/j.neuron.2013.04.018

Amaral, D. G., & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience*, *31*(3), 571-591. https://doi.org/10.1016/0306-4522(89)90424-7

Aminoff, E. M., Kveraga, K., & Bar, M. (2013). The role of the parahippocampal cortex in cognition. *Trends Cogn Sci*, *17*(8), 379-390. https://doi.org/10.1016/j.tics.2013.06.009

Andersen, P., Bliss, T. V., Lomo, T., Olsen, L. I., & Skrede, K. K. (1971). Lamellar organization of hippocampal excitatory pathways. *Acta Physiol Scand*, *76*(1), 4A-5A. https://doi.org/10.1111/j.1748-1716.1969.tb04499.x

Andersson, J. L., Hutton, C., Ashburner, J., Turner, R., & Friston, K. (2001). Modeling geometric deformations in EPI time series. *Neuroimage*, *13*(5), 903-919. https://doi.org/10.1006/nimg.2001.0746

Andersson, J. L., Skare, S., & Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage*, *20*(2), 870-888. https://doi.org/10.1016/S1053-8119(03)00336-7

Antony, J. W., Ferreira, C. S., Norman, K. A., & Wimber, M. (2017). Retrieval as a Fast Route to Memory Consolidation. *Trends in cognitive sciences*, *21*(8). https://doi.org/10.1016/j.tics.2017.05.001

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *Neuroimage*, *38*(1), 95-113. https://doi.org/10.1016/j.neuroimage.2007.07.007

Attardo, A., Fitzgerald, J. E., & Schnitzer, M. J. (2015). Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature*, *523*(7562), 592-596. https://doi.org/10.1038/nature14467

Audrain, S., & McAndrews, M. P. (2022). Schemas provide a scaffold for neocortical integration of new memories over time. *Nat Commun*, *13*(1), 5795. https://doi.org/10.1038/s41467-022-33517-0

Augustinack, J. C., Huber, K. E., Stevens, A. A., Roy, M., Frosch, M. P., van der Kouwe, A. J., Wald, L. L., Van Leemput, K., McKee, A. C., Fischl, B., & Alzheimer's Disease Neuroimaging, I. (2013). Predicting the location of human perirhinal cortex, Brodmann's area 35, from MRI. *Neuroimage*, *64*, 32-42. https://doi.org/10.1016/j.neuroimage.2012.08.071

Averbeck, B. B., & Seo, M. (2008). The statistical neuroanatomy of frontal networks in the macaque. *PLoS Comput Biol*, *4*(4), e1000050. https://doi.org/10.1371/journal.pcbi.1000050

Bachevalier, J., Meunier, M., Lu, M. X., & Ungerleider, L. G. (1997). Thalamic and temporal cortex input to medial prefrontal cortex in rhesus monkeys. *Exp Brain Res*, *115*(3), 430-444. https://doi.org/10.1007/pl00005713

Baker, S., Vieweg, P., Gao, F., Gilboa, A., Wolbers, T., Black, S. E., & Rosenbaum, R. S. (2016). The Human Dentate Gyrus Plays a Necessary Role in Discriminating New Memories. *Curr Biol*, *26*(19), 2629-2634. https://doi.org/10.1016/j.cub.2016.07.081

Baldassano, C., Esteva, A., Fei-Fei, L., & Beck, D. M. (2016). Two Distinct Scene-Processing Networks Connecting Vision and Memory. *eNeuro*, *3*(5). https://doi.org/10.1523/ENEURO.0178-16.2016

Baldassano, C., Hasson, U., & Norman, K. A. (2018). Representation of Real-World Event Schemas during Narrative Perception. *J Neurosci*, *38*(45), 9689-9699. https://doi.org/10.1523/JNEUROSCI.0251-18.2018

Barbas, H. (2015). General cortical and special prefrontal connections: principles from structure to function. *Annu Rev Neurosci*, *38*, 269-289. https://doi.org/10.1146/annurev-neuro-071714-033936

Barense, M. D., Henson, R. N., Lee, A. C., & Graham, K. S. (2010). Medial temporal lobe activity during complex discrimination of faces, objects, and scenes: Effects of viewpoint. *Hippocampus*, *20*(3), 389-401. https://doi.org/10.1002/hipo.20641

Barry, D. N., Barnes, G. R., Clark, I. A., & Maguire, E. A. (2019). The Neural Dynamics of Novel Scene Imagery. *J Neurosci*, *39*(22), 4375-4386. https://doi.org/10.1523/JNEUROSCI.2497-18.2019

Barry, D. N., Chadwick, M. J., & Maguire, E. A. (2018). Nonmonotonic recruitment of ventromedial prefrontal cortex during remote memory recall. *PLoS Biol*, *16*(7), e2005479. https://doi.org/10.1371/journal.pbio.2005479

Barry, D. N., Clark, I. A., & Maguire, E. A. (2021). The relationship between hippocampal subfield volumes and autobiographical memory persistence. *Hippocampus*, *31*(4). https://doi.org/10.1002/hipo.23293

Barry, D. N., & Maguire, E. A. (2019). Remote Memory and the Hippocampus: A Constructive Critique. *Trends Cogn Sci*, *23*(2), 128-142. https://doi.org/10.1016/j.tics.2018.11.005

Barry, D. N., & Maguire, E. A. (2024). Functions of the human hippocampus. In R. Morris, D. G. Amaral, T. Bliss, K. Duff, & J. O'Keefe (Eds.), *The Hippocampus Book*. Oxford University Press.

Bartlett, F. C. (1932). Remembering: a study in experimental and social psychology. *Macmillan*.

Bartsch, T., Dohring, 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

Bastin, J., Vidal, J. R., Bouvier, S., Perrone-Bertolotti, M., Benis, D., Kahane, P., David, O., Lachaux, J. P., & Epstein, R. A. (2013). Temporal components in the parahippocampal place area revealed by human intracerebral recordings. *J Neurosci*, *33*(24), 10123-10131. https://doi.org/10.1523/JNEUROSCI.4646-12.2013

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, *76*(4), 695-711. https://doi.org/10.1016/j.neuron.2012.10.038

Beaty, R. E., Thakral, P. P., Madore, K. P., Benedek, M., & Schacter, D. L. (2018). Core Network Contributions to Remembering the Past, Imagining the Future, and Thinking Creatively. *J Cogn Neurosci*, *30*(12), 1939-1951. https://doi.org/10.1162/jocn_a_01327

Bechara, A. (2004). The role of emotion in decision-making: evidence from neurological patients with orbitofrontal damage. *Brain Cogn*, *55*(1), 30-40. https://doi.org/10.1016/j.bandc.2003.04.001

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, *100*(2). https://doi.org/10.1016/j.neuron.2018.10.002

Benoit, R. G., Szpunar, K. K., & Schacter, D. L. (2014). Ventromedial prefrontal cortex supports affective future simulation by integrating distributed knowledge. *Proc Natl Acad Sci U S A*, *111*(46), 16550-16555. https://doi.org/10.1073/pnas.1419274111

Bergmann, J., Petro, L. S., Abbatecola, C., Li, M. S., Morgan, A. T., & Muckli, L. (2024). Cortical depth profiles in primary visual cortex for illusory and imaginary experiences. *Nat Commun*, *15*(1), 1002. https://doi.org/10.1038/s41467-024-45065-w

Berron, D., Schutze, H., Maass, A., Cardenas-Blanco, A., Kuijf, H. J., Kumaran, D., & Duzel, E. (2016). Strong Evidence for Pattern Separation in Human Dentate Gyrus. *J Neurosci*, *36*(29), 7569-7579. https://doi.org/10.1523/JNEUROSCI.0518-16.2016

Berron, D., Vieweg, P., Hochkeppler, A., Pluta, J. B., Ding, S. L., Maass, A., Luther, A., Xie, L., Das, S. R., Wolk, D. A., Wolbers, T., Yushkevich, P. A., Duzel, E., & Wisse, L. E. M. (2017). A protocol for manual segmentation of medial temporal lobe subregions in 7 Tesla MRI. *Neuroimage Clin*, *15*, 466-482. https://doi.org/10.1016/j.nicl.2017.05.022

Bertossi, E., Candela, V., De Luca, F., & Ciaramelli, E. (2017). Episodic future thinking following vmPFC damage: Impaired event construction, maintenance, or narration? *Neuropsychology*, *31*(3), 337-348. https://doi.org/10.1037/neu0000345

Bertossi, E., & Ciaramelli, E. (2016). Ventromedial prefrontal damage reduces mind-wandering and biases its temporal focus. *Soc Cogn Affect Neurosci*, *11*(11), 1783-1791. https://doi.org/10.1093/scan/nsw099

Bertossi, E., Tesini, C., Cappelli, A., & Ciaramelli, E. (2016). Ventromedial prefrontal damage causes a pervasive impairment of episodic memory and future thinking. *Neuropsychologia*, *90*, 12-24. https://doi.org/10.1016/j.neuropsychologia.2016.01.034

Bird, C. M., Vargha-Khadem, F., & Burgess, N. (2008). Impaired memory for scenes but not faces in developmental hippocampal amnesia: a case study. *Neuropsychologia*, *46*(4), 1050-1059. https://doi.org/10.1016/j.neuropsychologia.2007.11.007

Blamire, A. M., Ogawa, S., Ugurbil, K., Rothman, D., McCarthy, G., Ellermann, J. M., Hyder, F., Rattner, Z., & Shulman, R. G. (1992). Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging. *Proc Natl Acad Sci U S A*, *89*(22), 11069-11073. https://doi.org/10.1073/pnas.89.22.11069

Bohbot, V. D., Iaria, G., & Petrides, M. (2004). Hippocampal function and spatial memory: evidence from functional neuroimaging in healthy participants and performance of patients with medial temporal lobe resections. *Neuropsychology*, *18*(3), 418-425. https://doi.org/10.1037/0894-4105.18.3.418

Bonnici, H. M., Chadwick, M. J., Kumaran, D., Hassabis, D., Weiskopf, N., & Maguire, E. A. (2012). Multi-voxel pattern analysis in human hippocampal subfields. *Front Hum Neurosci*, *6*, 290. https://doi.org/10.3389/fnhum.2012.00290

Bonnici, H. M., Chadwick, M. J., Lutti, A., Hassabis, D., Weiskopf, N., & Maguire, E. A. (2012). Detecting representations of recent and remote autobiographical memories in vmPFC and hippocampus. *J Neurosci*, *32*(47), 16982-16991. https://doi.org/10.1523/JNEUROSCI.2475-12.2012

Bonnici, H. M., Chadwick, M. J., & Maguire, E. A. (2013). Representations of recent and remote autobiographical memories in hippocampal subfields. *Hippocampus*, *23*(10), 849-854. https://doi.org/10.1002/hipo.22155

Bonnici, H. M., Kumaran, D., Chadwick, M. J., Weiskopf, N., Hassabis, D., & Maguire, E. A. (2012). Decoding representations of scenes in the medial temporal lobes. *Hippocampus*, *22*(5), 1143-1153. https://doi.org/10.1002/hipo.20960

Bonnici, H. M., & Maguire, E. A. (2018). Two years later - Revisiting autobiographical memory representations in vmPFC and hippocampus. *Neuropsychologia*, *110*, 159-169. https://doi.org/10.1016/j.neuropsychologia.2017.05.014

Boxerman, J. L., Hamberg, L. M., Rosen, B. R., & Weisskoff, R. M. (1995). MR contrast due to intravascular magnetic susceptibility perturbations. *Magnetic resonance in medicine*, *34*(4). https://doi.org/10.1002/mrm.1910340412

Braak, H., & Braak, E. (1991). Neuropathological stageing of Alzheimer-related changes. *Acta Neuropathol*, *82*(4), 239-259. https://doi.org/10.1007/BF00308809

Brandman, T., & Peelen, M. V. (2017). Interaction between Scene and Object Processing Revealed by Human fMRI and MEG Decoding. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *37*(32). https://doi.org/10.1523/JNEUROSCI.0582-17.2017

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*.

Buckner, R. L. (1996). Beyond HERA: Contributions of specific prefrontal brain areas to long-term memory retrieval. *Psychon Bull Rev*, *3*(2), 149-158. https://doi.org/10.3758/BF03212413

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

Callaghan, M. F., Josephs, O., Herbst, M., Zaitsev, M., Todd, N., & Weiskopf, N. (2015). An evaluation of prospective motion correction (PMC) for high resolution quantitative MRI. *Front Neurosci*, *9*, 97. https://doi.org/10.3389/fnins.2015.00097

Callaghan, M. F., Lutti, A., Ashburner, J., Balteau, E., Corbin, N., Draganski, B., Helms, G., Kherif, F., Leutritz, T., Mohammadi, S., Phillips, C., Reimer, E., Ruthotto, L., Seif, M., Tabelow, K., Ziegler, G., & Weiskopf, N. (2019). Example dataset for the hMRI toolbox. *Data Brief*, *25*, 104132. https://doi.org/10.1016/j.dib.2019.104132

Campbell, K. L., Madore, K. P., Benoit, R. G., Thakral, P. P., & Schacter, D. L. (2018). Increased hippocampus to ventromedial prefrontal connectivity during the construction of

episodic future events. *Hippocampus*, *28*(2), 76-80. https://doi.org/10.1002/hipo.22812

Carmichael, S. T., & Price, J. L. (1995). Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. *J Comp Neurol*, *363*(4), 615-641. https://doi.org/10.1002/cne.903630408

Carricarte, T., Iamshchinina, P., Trampel, R., Chaimow, D., Weiskopf, N., & Cichy, R. M. (2024). Laminar dissociation of feedforward and feedback in high-level ventral visual cortex during imagery and perception. *iScience*, *27*(7), 110229. https://doi.org/10.1016/j.isci.2024.110229

Catani, M., Dell'acqua, F., Vergani, F., Malik, F., Hodge, H., Roy, P., Valabregue, R., & Thiebaut de Schotten, M. (2012). Short frontal lobe connections of the human brain. *Cortex*, *48*(2), 273-291. https://doi.org/10.1016/j.cortex.2011.12.001

Chadwick, M. J. (2012). What information is represented in the human hippocampus? Doctoral thesis, UCL. https://discovery.ucl.ac.uk/id/eprint/1363077/1/Chadwick_Thesis.pdf

Chadwick, M. J., Bonnici, H. M., & Maguire, E. A. (2014). CA3 size predicts the precision of memory recall. *Proc Natl Acad Sci U S A*, *111*(29), 10720-10725. https://doi.org/10.1073/pnas.1319641111

Chadwick, M. J., Hassabis, D., & Maguire, E. A. (2011). Decoding overlapping memories in the medial temporal lobes using high-resolution fMRI. *Learn Mem*, *18*(12), 742-746. https://doi.org/10.1101/lm.023671.111

Chadwick, M. J., Hassabis, D., Weiskopf, N., & Maguire, E. A. (2010). Decoding individual episodic memory traces in the human hippocampus. *Curr Biol*, *20*(6), 544-547. https://doi.org/10.1016/j.cub.2010.01.053

Chang, H., & Fitzpatrick, J. M. (1992). A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE Trans Med Imaging*, *11*(3), 319-329. https://doi.org/10.1109/42.158935

Chang, W., Langella, S., Seo, M. S., Huynh, K., Yap, P., & Lin, W. (2022). Cross-layer Balance of Visuo-hippocampal Functional Connectivity Is Associated With Episodic Memory Recognition Accuracy. *Research Square*. https://doi.org/10.21203/rs.3.rs-1789565/v1

Chen, Y., Beech, P., Yin, Z., Jia, S., Zhang, J., Yu, Z., & Liu, J. K. (2024). Decoding dynamic visual scenes across the brain hierarchy. *PLoS Comput Biol*, *20*(8), e1012297. https://doi.org/10.1371/journal.pcbi.1012297

Ciaramelli, E., De Luca, F., Monk, A. M., McCormick, C., & Maguire, E. A. (2019). What "wins" in VMPFC: Scenes, situations, or schema? *Neuroscience and biobehavioral reviews*, *100*. https://doi.org/10.1016/j.neubiorev.2019.03.001

Ciaramelli, E., Ghetti, S., Frattarelli, M., & Ladavas, E. (2006). When true memory availability promotes false memory: evidence from confabulating patients. *Neuropsychologia*, *44*(10), 1866-1877. https://doi.org/10.1016/j.neuropsychologia.2006.02.008

Cipolotti, L., Shallice, T., Chan, D., Fox, N., Scahill, R., Harrison, G., Stevens, J., & Rudge, P. (2001). Long-term retrograde amnesia...the crucial role of the hippocampus. *Neuropsychologia*, *39*(2), 151-172. https://doi.org/10.1016/s0028-3932(00)00103-2

Clark, I. A., Callaghan, M. F., Weiskopf, N., Maguire, E. A., & Mohammadi, S. (2021). Reducing Susceptibility Distortion Related Image Blurring in Diffusion MRI EPI Data. *Front Neurosci*, *15*, 706473. https://doi.org/10.3389/fnins.2021.706473

Clark, I. A., Dalton, M. A., & Maguire, E. A. (2023). Posterior hippocampal CA2/3 volume is associated with autobiographical memory recall ability in lower performing individuals. *Scientific Reports*, *13*(1). https://doi.org/10.1038/s41598-023-35127-2

Clark, I. A., Hotchin, V., Monk, A., Pizzamiglio, G., Liefgreen, A., & Maguire, E. A. (2019). Identifying the cognitive processes underpinning hippocampal-dependent tasks. *J Exp Psychol Gen*, *148*(11), 1861-1881. https://doi.org/10.1037/xge0000582

Clark, I. A., & Maguire, E. A. (2023). Release of cognitive and multimodal MRI data including real-world tasks and hippocampal subfield segmentations. *Sci Data*, *10*(1), 540. https://doi.org/10.1038/s41597-023-02449-9

Clark, I. A., Monk, A. M., & Maguire, E. A. (2020). Characterizing Strategy Use During the Performance of Hippocampal-Dependent Tasks. *Front Psychol*, *11*, 2119. https://doi.org/10.3389/fpsyg.2020.02119

Colgin, L. L. (2011). Oscillations and hippocampal-prefrontal synchrony. *Curr Opin Neurobiol*, *21*(3), 467-474. https://doi.org/10.1016/j.conb.2011.04.006

Coras, R., Pauli, E., Li, J., Schwarz, M., Rossler, K., Buchfelder, M., Hamer, H., Stefan, H., & Blumcke, I. (2014). Differential influence of hippocampal subfields to memory formation: insights from patients with temporal lobe epilepsy. *Brain*, *137*(Pt 7), 1945-1957. https://doi.org/10.1093/brain/awu100

Corkin, S. (2002). What's new with the amnesic patient H.M.? *Nat Rev Neurosci*, *3*(2), 153-160. https://doi.org/10.1038/nrn726

Coutanche, M. N., & Thompson-Schill, S. L. (2013). Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. *Front Hum Neurosci*, *7*, 15. https://doi.org/10.3389/fnhum.2013.00015

Cragg, B. G. (1969). The topography of the afferent projections in the circumstriate visual cortex of the monkey studied by the Nauta method. *Vision Res*, *9*(7), 733-747. https://doi.org/10.1016/0042-6989(69)90011-x

Crivelli-Decker, J., Clarke, A., Park, S. A., Huffman, D. J., Boorman, E. D., & Ranganath, C. (2023). Goal-oriented representations in the human hippocampus during planning and navigation. *Nature communications*, *14*(1), 2946. https://doi.org/10.1038/s41467-023-35967-6

Cukur, T., Huth, A. G., Nishimoto, S., & Gallant, J. L. (2016). Functional Subdomains within Scene-Selective Cortex: Parahippocampal Place Area, Retrosplenial Complex, and Occipital Place Area. *J Neurosci*, *36*(40), 10257-10273. https://doi.org/10.1523/JNEUROSCI.4033-14.2016

Czervionke, L. F., Czervionke, J. M., Daniels, D. L., & Haughton, V. M. (1988). Characteristic features of MR truncation artifacts. *AJR Am J Roentgenol*, *151*(6), 1219-1228. https://doi.org/10.2214/ajr.151.6.1219

D'Argembeau, A. (2013). On the role of the ventromedial prefrontal cortex in self-processing: the valuation hypothesis. *Front Hum Neurosci*, *7*, 372. https://doi.org/10.3389/fnhum.2013.00372

D'Souza, R. D., Wang, Q., Ji, W., Meier, A. M., Kennedy, H., Knoblauch, K., & Burkhalter, A. (2022). Hierarchical and nonhierarchical features of the mouse visual cortical network. *Nat Commun*, *13*(1), 503. https://doi.org/10.1038/s41467-022-28035-y

Dalton, M. A., D'Souza, A., Lv, J., & Calamante, F. (2022). New insights into anatomical connectivity along the anterior-posterior axis of the human hippocampus using in vivo quantitative fibre tracking. *Elife*, *11*. https://doi.org/10.7554/eLife.76143

Dalton, M. A., & Maguire, E. A. (2017). The pre/parasubiculum: a hippocampal hub for scene-based cognition? *Curr Opin Behav Sci*, *17*, 34-40. https://doi.org/10.1016/j.cobeha.2017.06.001

Dalton, M. A., McCormick, C., & Maguire, E. A. (2019). Differences in functional connectivity along the anterior-posterior axis of human hippocampal subfields. *Neuroimage*, *192*, 38-51. https://doi.org/10.1016/j.neuroimage.2019.02.066

Dalton, M. A., Zeidman, P., Barry, D. N., Williams, E., & Maguire, E. A. (2017). Segmenting subregions of the human hippocampus on structural magnetic resonance image scans: An illustrated tutorial. *Brain Neurosci Adv*, *1*, 2398212817701448. https://doi.org/10.1177/2398212817701448

Dalton, M. A., Zeidman, P., McCormick, C., & Maguire, E. A. (2018). Differentiable Processing of Objects, Associations, and Scenes within the Hippocampus. *J Neurosci*, *38*(38), 8146-8159. https://doi.org/10.1523/JNEUROSCI.0263-18.2018

Davachi, L., & DuBrow, S. (2015). How the hippocampus preserves order: The role of prediction and context. *Trends in cognitive sciences*, *19*(2). https://doi.org/10.1016/j.tics.2014.12.004

De Luca, F., McCormick, C., Mullally, S. L., Intraub, H., Maguire, E. A., & Ciaramelli, E. (2018). Boundary extension is attenuated in patients with ventromedial prefrontal cortex damage. *Cortex; a journal devoted to the study of the nervous system and behavior*, *108*. https://doi.org/10.1016/j.cortex.2018.07.002

Dede, A. J., Frascino, J. C., Wixted, J. T., & Squire, L. R. (2016). Learning and remembering real-world events after medial temporal lobe damage. *Proc Natl Acad Sci U S A*, *113*(47), 13480-13485. https://doi.org/10.1073/pnas.1617025113

Degutis, J. K., Chaimow, D., Haenelt, D., Assem, M., Duncan, J., Haynes, J. D., Weiskopf, N., & Lorenz, R. (2024). Dynamic layer-specific processing in the prefrontal cortex during working memory. *Commun Biol*, *7*(1), 1140. https://doi.org/10.1038/s42003-024-06780-8

DeKraker, J., Haast, R. A. M., Yousif, M. D., Karat, B., Lau, J. C., Kohler, S., & Khan, A. R. (2022). Automated hippocampal unfolding for morphometry and subfield segmentation with HippUnfold. *Elife*, *11*. https://doi.org/10.7554/eLife.77945

DeKraker, J., Kohler, S., & Khan, A. R. (2021). Surface-based hippocampal subfield segmentation. *Trends Neurosci*, *44*(11), 856-863. https://doi.org/10.1016/j.tins.2021.06.005

Della Sala, S., Laiacona, M., Spinnler, H., & Trivelli, C. (1993). Autobiographical recollection and frontal damage. *Neuropsychologia*, *31*(8), 823-839. https://doi.org/10.1016/0028-3932(93)90131-i

Denny, C. A., Kheirbek, M. A., Alba, E. L., Tanaka, K. F., Brachman, R. A., Laughman, K. B., Tomm, N. K., Turi, G. F., Losonczy, A., & Hen, R. (2014). Hippocampal memory traces are differentially modulated by experience, time, and adult neurogenesis. *Neuron*, *83*(1), 189-201. https://doi.org/10.1016/j.neuron.2014.05.018

Deshmane, A., Gulani, V., Griswold, M. A., & Seiberlich, N. (2012). Parallel MR imaging. *J Magn Reson Imaging*, *36*(1), 55-72. https://doi.org/10.1002/jmri.23639

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., & Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, *31*(3). https://doi.org/10.1016/j.neuroimage.2006.01.021

Destrieux, C., Fischl, B., Dale, A., & Halgren, E. (2010). Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*, *53*(1), 1-15. https://doi.org/10.1016/j.neuroimage.2010.06.010

Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, *26*, 297-302.

Dickerson, B. C., & Eichenbaum, H. (2010). The episodic memory system: neurocircuitry and disorders. *Neuropsychopharmacology*, *35*(1), 86-104. https://doi.org/10.1038/npp.2009.126

Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, *463*(7281), 657-661. https://doi.org/10.1038/nature08704

Duda, O. R., Hart, P. E., & Stork, D. G. (2001). *Pattern Classification*. Wiley.

Duvernoy, H. M., Delon, S., & Vannson, J. L. (1981). Cortical blood vessels of the human brain. *Brain Res Bull*, *7*(5), 519-579. https://doi.org/10.1016/0361-9230(81)90007-1

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

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. A., & Baker, C. I. (2019). Scene Perception in the Human Brain. *Annu Rev Vis Sci*, *5*, 373-397. https://doi.org/10.1146/annurev-vision-091718-014809

Faust, T. E., Gunner, G., & Schafer, D. P. (2021). Mechanisms governing activity-dependent synaptic pruning in the developing mammalian CNS. *Nat Rev Neurosci*, *22*(11), 657-673. https://doi.org/10.1038/s41583-021-00507-y

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, *1*(1), 1-47. https://doi.org/10.1093/cercor/1.1.1-a

Finn, E. S., Huber, L., & Bandettini, P. A. (2021). Higher and deeper: Bringing layer fMRI to association cortex. *Prog Neurobiol*, *207*, 101930. https://doi.org/10.1016/j.pneurobio.2020.101930

Finn, E. S., Huber, L., Jangraw, D. C., Molfese, P. J., & Bandettini, P. A. (2019). Layer-dependent activity in human prefrontal cortex during working memory. *Nat Neurosci*, *22*(10), 1687-1695. https://doi.org/10.1038/s41593-019-0487-z

Fischl, B., & Dale, A. M. (2000). Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proc Natl Acad Sci U S A*, *97*(20), 11050-11055. https://doi.org/10.1073/pnas.200033797

Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., & Dale, A. M. (2002). Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, *33*(3), 341-355. https://doi.org/10.1016/s0896-6273(02)00569-x

Fischl, B., Stevens, A. A., Rajendran, N., Yeo, B. T., Greve, D. N., Van Leemput, K., Polimeni, J. R., Kakunoori, S., Buckner, R. L., Pacheco, J., Salat, D. H., Melcher, J., Frosch, M. P., Hyman, B. T., Grant, P. E., Rosen, B. R., van der Kouwe, A. J., Wiggins, G. C., Wald, L. L., & Augustinack, J. C. (2009). Predicting the location of entorhinal cortex from MRI. *Neuroimage*, *47*(1), 8-17. https://doi.org/10.1016/j.neuroimage.2009.04.033

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Segonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., & Dale, A. M. (2004). Automatically parcellating the human cerebral cortex. *Cereb Cortex*, *14*(1), 11-22. https://doi.org/10.1093/cercor/bhg087

Frackowiak, R. S. J., Friston, K. J., Frith, C. D., Dolan, R. J., Price, C. J., Zeki, S., Ashburner, J. T., & Penny, W. D. (2004). *Human brain function*. Elsevier Academic Press.

Frey, U., & Morris, R. G. (1998). Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci*, *21*(5), 181-188. https://doi.org/10.1016/s0166-2236(97)01189-2

Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, *19*(4), 1273-1302. https://doi.org/10.1016/s1053-8119(03)00202-7

Fuentemilla, L., Barnes, G. R., Duzel, E., & Levine, B. (2014). Theta oscillations orchestrate medial temporal lobe and neocortex in remembering autobiographical memories. *Neuroimage*, *85 Pt 2*, 730-737. https://doi.org/10.1016/j.neuroimage.2013.08.029

Gabrieli, J. D., Poldrack, R. A., & Desmond, J. E. (1998). The role of left prefrontal cortex in language and memory. *Proc Natl Acad Sci U S A*, *95*(3), 906-913. https://doi.org/10.1073/pnas.95.3.906

Gaesser, B., Spreng, R. N., McLelland, V. C., Addis, D. R., & Schacter, D. L. (2013). Imagining the future: evidence for a hippocampal contribution to constructive processing. *Hippocampus*, *23*(12), 1150-1161. https://doi.org/10.1002/hipo.22152

Gallichan, D., Scholz, J., Bartsch, A., Behrens, T. E., Robson, M. D., & Miller, K. L. (2010). Addressing a systematic vibration artifact in diffusion-weighted MRI. *Hum Brain Mapp*, *31*(2), 193-202. https://doi.org/10.1002/hbm.20856

Gauthier, J. L., & Tank, D. W. (2018). A Dedicated Population for Reward Coding in the Hippocampus. *Neuron*, *99*(1), 179–193.e7. https://doi.org/10.1016/j.neuron.2018.06.008

Ghosh, V. E., & Gilboa, A. (2014). What is a memory schema? A historical perspective on current neuroscience literature. *Neuropsychologia*, *53*. https://doi.org/10.1016/j.neuropsychologia.2013.11.010

Ghosh, V. E., Moscovitch, M., Melo Colella, B., & Gilboa, A. (2014). Schema representation in patients with ventromedial PFC lesions. *J Neurosci*, *34*(36), 12057-12070. https://doi.org/10.1523/JNEUROSCI.0740-14.2014

Gibbs, J. (1898). Fourier's Series. *Nature*, *59*(200).

Gilboa, A. (2010). Strategic retrieval, confabulations, and delusions: theory and data. *Cogn Neuropsychiatry*, *15*(1), 145-180. https://doi.org/10.1080/13546800903056965

Gilboa, A., Alain, C., Stuss, D. T., Melo, B., Miller, S., & Moscovitch, M. (2006). Mechanisms of spontaneous confabulations: a strategic retrieval account. *Brain*, *129*(Pt 6), 1399-1414. https://doi.org/10.1093/brain/awl093

Gilboa, A., & Marlatte, H. (2017). Neurobiology of Schemas and Schema-Mediated Memory. *Trends Cogn Sci*, *21*(8), 618-631. https://doi.org/10.1016/j.tics.2017.04.013

Gilboa, A., Winocur, G., Grady, C. L., Hevenor, S. J., & Moscovitch, M. (2004). Remembering our past: functional neuroanatomy of recollection of recent and very remote personal events. *Cereb Cortex*, *14*(11), 1214-1225. https://doi.org/10.1093/cercor/bhh082

Gilmore, A. W., Quach, A., Kalinowski, S. E., Gonzalez-Araya, E. I., Gotts, S. J., Schacter, D. L., & Martin, A. (2021). Evidence supporting a time-limited hippocampal role in retrieving autobiographical memories. *Proc Natl Acad Sci U S A*, *118*(12). https://doi.org/10.1073/pnas.2023069118

Godlove, D. C., Maier, A., Woodman, G. F., & Schall, J. D. (2014). Microcircuitry of agranular frontal cortex: testing the generality of the canonical cortical microcircuit. *J Neurosci*, *34*(15), 5355-5369. https://doi.org/10.1523/JNEUROSCI.5127-13.2014

Goyal, A., Miller, J., Watrous, A. J., Lee, S. A., Coffey, T., Sperling, M. R., Sharan, A., Worrell, G., Berry, B., Lega, B., Jobst, B. C., Davis, K. A., Inman, C., Sheth, S. A., Wanda, P. A., Ezzyat, Y., Das, S. R., Stein, J., Gorniak, R., & Jacobs, J. (2018). Electrical Stimulation in Hippocampus and Entorhinal Cortex Impairs Spatial and Temporal Memory. *J Neurosci*, *38*(19), 4471-4481. https://doi.org/10.1523/JNEUROSCI.3049-17.2018

Graham, K. S., Barense, M. D., & Lee, A. C. (2010). Going beyond LTM in the MTL: a synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, *48*(4), 831-853. https://doi.org/10.1016/j.neuropsychologia.2010.01.001

Grande, X., Berron, D., Horner, A. J., Bisby, J. A., Duzel, E., & Burgess, N. (2019). Holistic Recollection via Pattern Completion Involves Hippocampal Subfield CA3. *J Neurosci*, *39*(41), 8100-8111. https://doi.org/10.1523/JNEUROSCI.0722-19.2019

Grande, X., Sauvage, M. M., Becke, A., Duzel, E., & Berron, D. (2022). Transversal functional connectivity and scene-specific processing in the human entorhinal-hippocampal circuitry. *Elife*, *11*. https://doi.org/10.7554/eLife.76479

Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *Neuroimage*, *48*(1), 63-72. https://doi.org/10.1016/j.neuroimage.2009.06.060

Guo, D., & Yang, J. (2020). Interplay of the long axis of the hippocampus and ventromedial prefrontal cortex in schema-related memory retrieval. *Hippocampus*, *30*(3), 263-277. https://doi.org/10.1002/hipo.23154

Gusnard, D. A., Akbudak, E., Shulman, G. L., & Raichle, M. E. (2001). Medial prefrontal cortex and self-referential mental activity: relation to a default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(7), 4259–4264. https://doi.org/10.1073/pnas.071043098

Guzman, S. J., Schlogl, A., Frotscher, M., & Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science*, *353*(6304), 1117-1123. https://doi.org/10.1126/science.aaf1836

Hadar, P. N., Kini, L. G., Coto, C., Piskin, V., Callans, L. E., Chen, S. H., Stein, J. M., Das, S. R., Yushkevich, P. A., & Davis, K. A. (2018). Clinical validation of automated hippocampal segmentation in temporal lobe epilepsy. *Neuroimage Clin*, *20*, 1139-1147. https://doi.org/10.1016/j.nicl.2018.09.032

Hainmueller, T., & Bartos, M. (2020). Dentate gyrus circuits for encoding, retrieval and discrimination of episodic memories. *Nat Rev Neurosci*, *21*(3), 153-168. https://doi.org/10.1038/s41583-019-0260-z

Hardstone, R., Zhu, M., Flinker, A., Melloni, L., Devore, S., Friedman, D., Dugan, P., Doyle, W. K., Devinsky, O., He, B. J., Hardstone, R., Zhu, M., Flinker, A., Melloni, L., Devore, S., Friedman, D., Dugan, P., Doyle, W. K., Devinsky, O., & He, B. J. (2021). Long-term priors influence visual perception through recruitment of long-range feedback. *Nature Communications*, *12*(1). https://doi.org/10.1038/s41467-021-26544-w

Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., Feiner, A., Feng, D., Gaudreault, N., Gerfen, C. R., Graddis, N., Groblewski, P. A., Henry, A. M., Ho, A., Howard, R., . . . Zeng, H. (2019).

Hierarchical organization of cortical and thalamic connectivity. *Nature*, *575*(7781), 195-202. https://doi.org/10.1038/s41586-019-1716-z

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., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proc Natl Acad Sci U S A*, *104*(5), 1726-1731. https://doi.org/10.1073/pnas.0610561104

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

Hasson, U., Furman, O., Clark, D., Dudai, Y., & Davachi, L. (2008). Enhanced intersubject correlations during movie viewing correlate with successful episodic encoding. *Neuron*, *57*(3), 452-462. https://doi.org/10.1016/j.neuron.2007.12.009

Havlicek, M., & Uludag, K. (2020). A dynamical model of the laminar BOLD response. *Neuroimage*, *204*. https://doi.org/ARTN 11620910.1016/j.neuroimage.2019.116209

Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nat Rev Neurosci*, *7*(7), 523-534. https://doi.org/10.1038/nrn1931

Hedouin, R., Commowick, O., Bannier, E., Scherrer, B., Taquet, M., Warfield, S. K., & Barillot, C. (2017). Block-Matching Distortion Correction of Echo-Planar Images With Opposite Phase Encoding Directions. *IEEE Trans Med Imaging*, *36*(5), 1106-1115. https://doi.org/10.1109/TMI.2016.2646920

Heinzle, J., Koopmans, P. J., den Ouden, H. E. M., Raman, S., & Stephan, K. E. (2016). A hemodynamic model for layered BOLD signals. *Neuroimage*, *125*, 556-570. https://doi.org/10.1016/j.neuroimage.2015.10.025

Herweg, N. A., & Kahana, M. J. (2018). Spatial Representations in the Human Brain. *Front Hum Neurosci*, *12*, 297. https://doi.org/10.3389/fnhum.2018.00297

Hickling, A. L., Clark, I. A., Wu, Y. I., & Maguire, E. A. (2024). Automated protocols for delineating human hippocampal subfields from 3 Tesla and 7 Tesla magnetic resonance imaging data. *Hippocampus*, *34*(6), 302-308. https://doi.org/10.1002/hipo.23606

Hodgetts, C. J., Shine, J. P., Lawrence, A. D., Downing, P. E., & Graham, K. S. (2016). Evidencing a place for the hippocampus within the core scene processing network. *Human Brain Mapping*, *37*(11). https://doi.org/10.1002/hbm.23275

Hodgetts, C. J., Voets, N. L., Thomas, A. G., Clare, S., Lawrence, A. D., & Graham, K. S. (2017). Ultra-High-Field fMRI Reveals a Role for the Subiculum in Scene Perceptual Discrimination. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *37*(12). https://doi.org/10.1523/JNEUROSCI.3225-16.2017

Holland, D., Kuperman, J. M., & Dale, A. M. (2010). Efficient correction of inhomogeneous static magnetic field-induced distortion in Echo Planar Imaging. *Neuroimage*, *50*(1), 175-183. https://doi.org/10.1016/j.neuroimage.2009.11.044

Hore, P. J. (1983). A new method for water suppression in the proton NMR spectra of aqueous solutions. *Journal of Magnetic Resonance (1969)*, *54*(3). https://doi.org/10.1016/0022-2364(83)90335-9

Huang, P., Correia, M. M., Rua, C., Rodgers, C. T., Henson, R. N., & Carlin, J. D. (2021). Correcting for Superficial Bias in 7T Gradient Echo fMRI. *Front Neurosci*, *15*, 715549. https://doi.org/10.3389/fnins.2021.715549

Huber, L. (2020). *Removing unwanted venous signal from GE-BOLD maps: Overview of vein removal models and implementations in LAYNII*. Retrieved 04/11/2024 from https://layerfmri.com/2020/04/02/devein/

Huber, L., Handwerker, D. A., Jangraw, D. C., Chen, G., Hall, A., Stuber, C., Gonzalez-Castillo, J., Ivanov, D., Marrett, S., Guidi, M., Goense, J., Poser, B. A., & Bandettini, P. A. (2017). High-Resolution CBV-fMRI Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1. *Neuron*, *96*(6), 1253-1263 e1257. https://doi.org/10.1016/j.neuron.2017.11.005

Hutton, C., Bork, A., Josephs, O., Deichmann, R., Ashburner, J., & Turner, R. (2002). Image distortion correction in fMRI: A quantitative evaluation. *Neuroimage*, *16*(1), 217-240. https://doi.org/10.1006/nimg.2001.1054

Iglesias, J. E., Augustinack, J. C., Nguyen, K., Player, C. M., Player, A., Wright, M., Roy, N., Frosch, M. P., McKee, A. C., Wald, L. L., Fischl, B., Van Leemput, K., & Alzheimer's Disease Neuroimaging, I. (2015). A computational atlas of the hippocampal formation using ex vivo, ultra-high resolution MRI: Application to adaptive segmentation of in vivo MRI. *Neuroimage*, *115*, 117-137. https://doi.org/10.1016/j.neuroimage.2015.04.042

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, *336*(6087), 1454-1458. https://doi.org/10.1126/science.1217230

Jezzard, P., & Balaban, R. S. (1995). Correction for geometric distortion in echo planar images from B0 field variations. *Magn Reson Med*, *34*(1), 65-73. https://doi.org/10.1002/mrm.1910340111

Jezzard, P., & Clare, S. (1999). Sources of distortion in functional MRI data. *Hum Brain Mapp*, *8*(2-3), 80-85. https://doi.org/10.1002/(sici)1097-0193(1999)8:2/3<80::aid-hbm2>3.0.co;2-c

Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci*, *10*(1), 100-107. https://doi.org/10.1038/nn1825

Jia, K., Zamboni, E., Rua, C., Goncalves, N. R., Kemper, V., Ng, A. K. T., Rodgers, C. T., Williams, G., Goebel, R., & Kourtzi, Z. (2021). A protocol for ultra-high field laminar fMRI in the human brain. *STAR Protoc*, *2*(2), 100415. https://doi.org/10.1016/j.xpro.2021.100415

Jones, D. K., & Cercignani, M. (2010). Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR Biomed*, *23*(7), 803-820. https://doi.org/10.1002/nbm.1543

Juchem, C., & de Graaf, R. A. (2017). B(0) magnetic field homogeneity and shimming for in vivo magnetic resonance spectroscopy. *Anal Biochem*, *529*, 17-29. https://doi.org/10.1016/j.ab.2016.06.003

Kaas, J. H., & Lin, C. S. (1977). Cortical projections of area 18 in owl monkeys. *Vision Res*, *17*(6), 739-741. https://doi.org/10.1016/s0042-6989(77)80013-8

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nat Neurosci*, *8*(5), 679-685. https://doi.org/10.1038/nn1444

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci*, *17*(11), 4302-4311. https://doi.org/10.1523/JNEUROSCI.17-11-04302.1997

Kaplan, R., Bush, D., Bisby, J. A., Horner, A. J., Meyer, S. S., & Burgess, N. (2017). Medial Prefrontal-Medial Temporal Theta Phase Coupling in Dynamic Spatial Imagery. *J Cogn Neurosci*, *29*(3), 507-519. https://doi.org/10.1162/jocn_a_01064

Kashyap, S., Ivanov, D., Havlicek, M., Poser, B. A., & Uludag, K. (2018). Impact of acquisition and analysis strategies on cortical depth-dependent fMRI. *Neuroimage*, *168*, 332-344. https://doi.org/10.1016/j.neuroimage.2017.05.022

Kay, K., Jamison, K. W., Vizioli, L., Zhang, R., Margalit, E., & Ugurbil, K. (2019). A critical assessment of data quality and venous effects in sub-millimeter fMRI. *Neuroimage*, *189*, 847-869. https://doi.org/10.1016/j.neuroimage.2019.02.006

Kennedy, H., & Bullier, J. (1985). A double-labeling investigation of the afferent connectivity to cortical areas V1 and V2 of the macaque monkey. *J Neurosci*, *5*(10), 2815-2830. https://doi.org/10.1523/JNEUROSCI.05-10-02815.1985

Kesner, R. P., & Rolls, E. T. (2015). A computational theory of hippocampal function, and tests of the theory: new developments. *Neurosci Biobehav Rev*, *48*, 92-147. https://doi.org/10.1016/j.neubiorev.2014.11.009

Kim, K., & Johnson, M. K. (2015). Activity in ventromedial prefrontal cortex during self-related processing: positive subjective value or personal significance?. *Social cognitive and affective neuroscience*, *10*(4), 494–500. https://doi.org/10.1093/scan/nsu078

Kim, M., & Maguire, E. A. (2018). Hippocampus, Retrosplenial and Parahippocampal Cortices Encode Multicompartment 3D Space in a Hierarchical Manner. *Cereb Cortex*, *28*(5), 1898-1909. https://doi.org/10.1093/cercor/bhy054

Kim, S., Jeneson, A., van der Horst, A. S., Frascino, J. C., Hopkins, R. O., & Squire, L. R. (2011). Memory, visual discrimination performance, and the human hippocampus. *J Neurosci*, *31*(7), 2624-2629. https://doi.org/10.1523/JNEUROSCI.5954-10.2011

Kirwan, C. B., Bayley, P. J., Galvan, V. V., & Squire, L. R. (2008). Detailed recollection of remote autobiographical memory after damage to the medial temporal lobe. *Proc Natl Acad Sci U S A*, *105*(7), 2676-2680. https://doi.org/10.1073/pnas.0712155105

Klein, S. B., Loftus, J., & Kihlstrom, J. F. (2005). Memory and temporal experience: the effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *https://doi.org/10.1521/soco.20.5.353.21125*, *20*(5). https://doi.org/10.1521/soco.20.5.353.21125

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge University Press.

Kobayashi, Y., & Amaral, D. G. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *J Comp Neurol*, *466*(1), 48-79. https://doi.org/10.1002/cne.10883

Kobayashi, Y., & Amaral, D. G. (2007). Macaque monkey retrosplenial cortex: III. Cortical efferents. *J Comp Neurol*, *502*(5), 810-833. https://doi.org/10.1002/cne.21346

Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Curr Biol*, *26*(3), 371-376. https://doi.org/10.1016/j.cub.2015.12.038

Kopelman, M. D., Stanhope, N., & Kingsley, D. (1999). Retrograde amnesia in patients with diencephalic, temporal lobe or frontal lesions. *Neuropsychologia*, *37*(8), 939-958. https://doi.org/10.1016/s0028-3932(98)00143-2

Koster, R., Chadwick, M. J., Chen, Y., Berron, D., Banino, A., Duzel, E., Hassabis, D., & Kumaran, D. (2018). Big-Loop Recurrence within the Hippocampal System Supports Integration of Information across Episodes. *Neuron*, *99*(6), 1342-1354 e1346. https://doi.org/10.1016/j.neuron.2018.08.009

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Front Syst Neurosci*, *2*, 4. https://doi.org/10.3389/neuro.06.004.2008

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., & Bandettini, P. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*(6), 1126-1141. https://doi.org/10.1016/j.neuron.2008.10.043

Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci*, *12*(5), 535-540. https://doi.org/10.1038/nn.2303

Kulaga-Yoskovitz, J., Bernhardt, B. C., Hong, S. J., Mansi, T., Liang, K. E., van der Kouwe, A. J., Smallwood, J., Bernasconi, A., & Bernasconi, N. (2015). Multi-contrast submillimetric 3 Tesla hippocampal subfield segmentation protocol and dataset. *Sci Data*, *2*, 150059. https://doi.org/10.1038/sdata.2015.59

Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends in cognitive sciences*, *20*(7). https://doi.org/10.1016/j.tics.2016.05.004

Kurczek, J., Wechsler, E., Ahuja, S., Jensen, U., Cohen, N. J., Tranel, D., & Duff, M. (2015). Differential contributions of hippocampus and medial prefrontal cortex to self-projection and self-referential processing. *Neuropsychologia*, *73*, 116-126. https://doi.org/10.1016/j.neuropsychologia.2015.05.002

Kuypers, H. G., Szwarcbart, M. K., Mishkin, M., & Rosvold, H. E. (1965). Occipitotemporal Corticocortical Connections in the Rhesus Monkey. *Exp Neurol*, *11*, 245-262. https://doi.org/10.1016/0014-4886(65)90016-6

Kwan, D., Craver, C. F., Green, L., Myerson, J., & Rosenbaum, R. S. (2013). Dissociations in future thinking following hippocampal damage: evidence from discounting and time perspective in episodic amnesia. *J Exp Psychol Gen*, *142*(4), 1355-1369. https://doi.org/10.1037/a0034001

Lamme, V. A., Super, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Curr Opin Neurobiol*, *8*(4), 529-535. https://doi.org/10.1016/s0959-4388(98)80042-1

Lawrence, S. J., Norris, D. G., & de Lange, F. P. (2019). Dissociable laminar profiles of concurrent bottom-up and top-down modulation in the human visual cortex. *Elife*, *8*. https://doi.org/10.7554/eLife.44422

Lawrence, S. J. D., van Mourik, T., Kok, P., Koopmans, P. J., Norris, D. G., & de Lange, F. P. (2018). Laminar Organization of Working Memory Signals in Human Visual Cortex. *Curr Biol*, *28*(21), 3435-3440 e3434. https://doi.org/10.1016/j.cub.2018.08.043

Lee, A. C., Brodersen, K. H., & Rudebeck, S. R. (2013). Disentangling spatial perception and spatial memory in the hippocampus: a univariate and multivariate pattern analysis fMRI study. *J Cogn Neurosci*, *25*(4), 534-546. https://doi.org/10.1162/jocn_a_00301

Lee, A. C., Buckley, M. J., Pegman, S. J., Spiers, H., Scahill, V. L., Gaffan, D., Bussey, T. J., Davies, R. R., Kapur, N., Hodges, J. R., & Graham, K. S. (2005). Specialization in the medial temporal lobe for processing of objects and scenes. *Hippocampus*, *15*(6), 782-797. https://doi.org/10.1002/hipo.20101

Lee, A. C., Yeung, L. K., & Barense, M. D. (2012). The hippocampus and visual perception. *Front Hum Neurosci*, *6*, 91. https://doi.org/10.3389/fnhum.2012.00091

Lee, J. K., Ekstrom, A. D., & Ghetti, S. (2014). Volume of hippocampal subfields and episodic memory in childhood and adolescence. *Neuroimage*, *94*, 162-171. https://doi.org/10.1016/j.neuroimage.2014.03.019

Lee, J. K., Fandakova, Y., Johnson, E. G., Cohen, N. J., Bunge, S. A., & Ghetti, S. (2020). Changes in anterior and posterior hippocampus differentially predict item-space, item-time, and item-item memory improvement. *Dev Cogn Neurosci*, *41*, 100741. https://doi.org/10.1016/j.dcn.2019.100741

Leelaarporn, P., Dalton, M. A., Stirnberg, R., Stöcker, T., Spottke, A., Schneider, A., & McCormick, C. (2024). Hippocampal subfields and their neocortical interactions during autobiographical memory. *Imaging Neuroscience*, *2*. https://doi.org/10.1162/imag_a_00105

Lehn, H., Steffenach, H. A., van Strien, N. M., Veltman, D. J., Witter, M. P., & Haberg, A. K. (2009). A specific role of the human hippocampus in recall of temporal sequences. *J Neurosci*, *29*(11), 3475-3484. https://doi.org/10.1523/JNEUROSCI.5370-08.2009

Leopold, A., Krueger, F., dal Monte, O., Pardini, M., Pulaski, S. J., Solomon, J., & Grafman, J. (2012). Damage to the left ventromedial prefrontal cortex impacts affective theory of mind. *Soc Cogn Affect Neurosci*, *7*(8), 871-880. https://doi.org/10.1093/scan/nsr071

Leutgeb, J. K., Leutgeb, S., Moser, M. B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science*, *315*(5814), 961-966. https://doi.org/10.1126/science.1135801

Levine, B., Svoboda, E., Hay, J. F., Winocur, G., & Moscovitch, M. (2002). Aging and autobiographical memory: dissociating episodic from semantic retrieval. *Psychol Aging*, *17*(4), 677-689. https://www.ncbi.nlm.nih.gov/pubmed/12507363

Lieberman, M. D., Straccia, M. A., Meyer, M. L., Du, M., & Tan, K. M. (2019). Social, self, (situational), and affective processes in medial prefrontal cortex (MPFC): Causal, multivariate, and reverse inference evidence. *Neurosci Biobehav Rev*, *99*, 311-328. https://doi.org/10.1016/j.neubiorev.2018.12.021

Lisman, J. E. (1999). Relating Hippocampal Circuitry to Function. *Neuron*, *22*(2). https://doi.org/10.1016/S0896-6273(00)81085-5

Liu, C., Guo, F., Qian, C., Zhang, Z., Sun, K., Wang, D. J., He, S., & Zhang, P. (2021). Layer-dependent multiplicative effects of spatial attention on contrast responses in human early visual cortex. *Prog Neurobiol*, *207*, 101897. https://doi.org/10.1016/j.pneurobio.2020.101897

Lorente de Nò, R. (1934). Studies on the Structure of the Cerebral Cortex II. Continuation of the Study of the Ammonic System. *Journal für Psychologie und Neurologie*, *46*, 113-177.

Lu, H., Hua, J., & van Zijl, P. C. (2013). Noninvasive functional imaging of cerebral blood volume with vascular-space-occupancy (VASO) MRI. *NMR Biomed*, *26*(8), 932-948. https://doi.org/10.1002/nbm.2905

Lund, J. S., Lund, R. D., Hendrickson, A. E., Bunt, A. H., & Fuchs, A. F. (1975). The origin of efferent pathways from the primary visual cortex, area 17, of the macaque monkey as shown by retrograde transport of horseradish peroxidase. *The Journal of comparative neurology*, *164*(3). https://doi.org/10.1002/cne.901640303

Lux, V., Atucha, E., Kitsukawa, T., & Sauvage, M. M. (2016). Imaging a memory trace over half a life-time in the medial temporal lobe reveals a time-limited role of CA3 neurons in retrieval. *Elife*, *5*, e11862. https://doi.org/10.7554/eLife.11862

Maass, A., Schutze, H., Speck, O., Yonelinas, A., Tempelmann, C., Heinze, H. J., Berron, D., Cardenas-Blanco, A., Brodersen, K. H., Stephan, K. E., & Duzel, E. (2014). Laminar activity in the hippocampus and entorhinal cortex related to novelty and episodic encoding. *Nat Commun*, *5*, 5547. https://doi.org/10.1038/ncomms6547

Maguire, E. A. (2001). Neuroimaging studies of autobiographical event memory. *Philos Trans R Soc Lond B Biol Sci*, *356*(1413), 1441-1451. https://doi.org/10.1098/rstb.2001.0944

Maguire, E. A., Frith, C. D., & Morris, R. G. (1999). The functional neuroanatomy of comprehension and memory: the importance of prior knowledge. *Brain*, *122 (Pt 10)*, 1839-1850. https://doi.org/10.1093/brain/122.10.1839

Maguire, E. A., & Mullally, S. L. (2013). The hippocampus: a manifesto for change. *J Exp Psychol Gen*, *142*(4), 1180-1189. https://doi.org/10.1037/a0033650

Maguire, E. A., Mummery, C. J., & Buchel, C. (2000). Patterns of hippocampal-cortical interaction dissociate temporal lobe memory subsystems. *Hippocampus*, *10*(4), 475-482. https://doi.org/10.1002/1098-1063(2000)10:4<475::AID-HIPO14>3.0.CO;2-X

Mandler, J. M. (1984). *Stories, scripts, and scenes: aspects of schema theory* (1st ed.). Psychology Press. https://doi.org/10.4324/9781315802459

Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, *10*(3). https://doi.org/10.1088/0022-3719/10/3/004

Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., & Kennedy, H. (2014).

Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *The Journal of comparative neurology*, *522*(1). https://doi.org/10.1002/cne.23458

Markuerkiaga, I., Barth, M., & Norris, D. G. (2016). A cortical vascular model for examining the specificity of the laminar BOLD signal. *Neuroimage*, *132*, 491-498. https://doi.org/10.1016/j.neuroimage.2016.02.073

Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P. F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *Neuroimage*, *49*(2), 1271-1281. https://doi.org/10.1016/j.neuroimage.2009.10.002

Marr, D. (1971). Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci*, *262*(841), 23-81. https://doi.org/10.1098/rstb.1971.0078

Martinez-Millan, L., & Hollander, H. (1975). Cortico-cortical projections from striate cortex of the squirrel monkey (Saimiri sciureus). A radioautographic study. *Brain Res*, *83*(3), 405-417. https://doi.org/10.1016/0006-8993(75)90833-1

Mayes, A., Montaldi, D., & Migo, E. (2007). Associative memory and the medial temporal lobes. *Trends Cogn Sci*, *11*(3), 126-135. https://doi.org/10.1016/j.tics.2006.12.003

Mayes, A. R., Isaac, C. L., Holdstock, J. S., Hunkin, N. M., Montaldi, D., Downes, J. J., Macdonald, C., Cezayirli, E., & Roberts, J. N. (2001). Memory for single items, word pairs, and temporal order of different kinds in a patient with selective hippocampal lesions. *Cogn Neuropsychol*, *18*(2), 97-123. https://doi.org/10.1080/02643290125897

Mayes, A. R., & Roberts, N. (2001). Theories of episodic memory. *Philos Trans R Soc Lond B Biol Sci*, *356*(1413), 1395-1408. https://doi.org/10.1098/rstb.2001.0941

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, *102*(3). https://doi.org/10.1037/0033-295X.102.3.419

McCormick, C., Barry, D. N., Jafarian, A., Barnes, G. R., & Maguire, E. A. (2020). vmPFC Drives Hippocampal Processing during Autobiographical Memory Recall Regardless of

Remoteness. *Cereb Cortex*, *30*(11), 5972-5987. https://doi.org/10.1093/cercor/bhaa172

McCormick, C., Ciaramelli, E., De Luca, F., & Maguire, E. A. (2018). Comparing and Contrasting the Cognitive Effects of Hippocampal and Ventromedial Prefrontal Cortex Damage: A Review of Human Lesion Studies. *Neuroscience*, *374*, 295-318. https://doi.org/10.1016/j.neuroscience.2017.07.066

McCormick, C., Dalton, M. A., Zeidman, P., & Maguire, E. A. (2021). Characterising the hippocampal response to perception, construction and complexity. *Cortex*, *137*, 1-17. https://doi.org/10.1016/j.cortex.2020.12.018

McCormick, C., Rosenthal, C. R., Miller, T. D., & Maguire, E. A. (2017). Deciding what is possible and impossible following hippocampal damage in humans. *Hippocampus*, *27*(3). https://doi.org/10.1002/hipo.22694

McCormick, C., Rosenthal, C. R., Miller, T. D., & Maguire, E. A. (2018). Mind-Wandering in People with Hippocampal Damage. *J Neurosci*, *38*(11), 2745-2754. https://doi.org/10.1523/JNEUROSCI.1812-17.2018

McDermott, K. B., Szpunar, K. K., & Christ, S. E. (2009). Laboratory-based and autobiographical retrieval tasks differ substantially in their neural substrates. *Neuropsychologia*, *47*(11), 2290-2298. https://doi.org/10.1016/j.neuropsychologia.2008.12.025

McGibney, G., Smith, M. R., Nichols, S. T., & Crawley, A. (1993). Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. *Magnetic resonance in medicine*, *30*(1). https://doi.org/10.1002/mrm.1910300109

McKenzie, S., & Eichenbaum, H. (2011). Consolidation and reconsolidation: two lives of memories? *Neuron*, *71*(2), 224-233. https://doi.org/10.1016/j.neuron.2011.06.037

McKinnon, G. C. (1993). Ultrafast interleaved gradient-echo-planar imaging on a standard scanner. *Magn Reson Med*, *30*(5), 609-616. https://doi.org/10.1002/mrm.1910300512

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in neurosciences*, *10*(10). https://doi.org/10.1016/0166-2236(87)90011-7

Medicine, A. E. C. o. (2014a). *The Gradient Echo Pulse Sequence and Modified Flip Angle (34 of 56)* YouTube. https://www.youtube.com/watch?v=yuj85NdU85c&list=PLPcImQzEnTpz-5TzxyyoYSbiAa9xdd89l&index=34&t=291s

Medicine, A. E. C. o. (2014b). *Introducing MRI* YouTube Playlist, YouTube. https://www.youtube.com/playlist?app=desktop&list=PLPcImQzEnTpz-5TzxyyoYSbiAa9xdd89l&cbrd=1

Mendoza-Halliday, D., Major, A. J., Lee, N., Lichtenfeld, M. J., Carlson, B., Mitchell, B., Meng, P. D., Xiong, Y. S., Westerberg, J. A., Jia, X., Johnston, K. D., Selvanayagam, J., Everling, S., Maier, A., Desimone, R., Miller, E. K., & Bastos, A. M. (2024). A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. *Nat Neurosci*, *27*(3), 547-560. https://doi.org/10.1038/s41593-023-01554-7

Miller, T. D., Chong, T. T., Aimola Davies, A. M., Johnson, M. R., Irani, S. R., Husain, M., Ng, T. W., Jacob, S., Maddison, P., Kennard, C., Gowland, P. A., & Rosenthal, C. R. (2020). Human hippocampal CA3 damage disrupts both recent and remote episodic memories. *Elife*, *9*. https://doi.org/10.7554/eLife.41836

Monk, A. M., Barry, D. N., Litvak, V., Barnes, G. R., & Maguire, E. A. (2021). Watching Movies Unfold, a Frame-by-Frame Analysis of the Associated Neural Dynamics. *eNeuro*, *8*(4). https://doi.org/10.1523/ENEURO.0099-21.2021

Monk, A. M., Dalton, M. A., Barnes, G. R., & Maguire, E. A. (2021). The Role of Hippocampal-Ventromedial Prefrontal Cortex Neural Dynamics in Building Mental Representations. *J Cogn Neurosci*, *33*(1), 89-103. https://doi.org/10.1162/jocn_a_01634

Morgan, P. S., Bowtell, R. W., McIntyre, D. J., & Worthington, B. S. (2004). Correction of spatial distortion in EPI due to inhomogeneous static magnetic fields using the reversed gradient method. *Journal of magnetic resonance imaging : JMRI*, *19*(4). https://doi.org/10.1002/jmri.20032

Mormann, F., Kornblith, S., Cerf, M., Ison, M. J., Kraskov, A., Tran, M., Knieling, S., Quian Quiroga, R., Koch, C., & Fried, I. (2017). Scene-selective coding by single neurons in the

human parahippocampal cortex. *Proc Natl Acad Sci U S A*, *114*(5), 1153-1158. https://doi.org/10.1073/pnas.1608159113

Morra, J. H., Tu, Z., Apostolova, L. G., Green, A. E., Toga, A. W., & Thompson, P. M. (2010). Comparison of AdaBoost and Support Vector Machines for Detecting Alzheimer's Disease Through Automated Hippocampal Segmentation. *IEEE Transactions on Medical Imaging*, *29*(1). https://doi.org/10.1109/TMI.2009.2021941

Moscovitch, M. (1989). Confabulation and the frontal systems: Strategic versus associative retrieval in neuropsychological theories of memory. In H. L. R. I. F. I. M. Craik (Ed.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 133–160). Lawrence Erlbaum Associates, Inc.

Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic Memory and Beyond: The Hippocampus and Neocortex in Transformation. *Annu Rev Psychol*, *67*, 105-134. https://doi.org/10.1146/annurev-psych-113011-143733

Moscovitch, M., & Melo, B. (1997). Strategic retrieval and the frontal lobes: evidence from confabulation and amnesia. *Neuropsychologia*, *35*(7), 1017-1034. https://doi.org/10.1016/s0028-3932(97)00028-6

mriquestions.com. *BOLD and Brain Activity: Does the BOLD response result from the firing of nerve cells?* . mriquestions.com. https://mriquestions.com/does-boldbrain-activity.html

Mugler, J. P., 3rd, Bao, S., Mulkern, R. V., Guttmann, C. R., Robertson, R. L., Jolesz, F. A., & Brookeman, J. R. (2000). Optimized single-slab three-dimensional spin-echo MR imaging of the brain. *Radiology*, *216*(3), 891-899. https://doi.org/10.1148/radiology.216.3.r00au46891

Mullally, S. L., Intraub, H., & Maguire, E. A. (2012). Attenuated boundary extension produces a paradoxical memory advantage in amnesic patients. *Curr Biol*, *22*(4), 261-268. https://doi.org/10.1016/j.cub.2012.01.001

Mullally, S. L., & Maguire, E. A. (2013). Exploring the role of space-defining objects in constructing and maintaining imagined scenes. *Brain Cogn*, *82*(1), 100-107. https://doi.org/10.1016/j.bandc.2013.02.013

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

Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr Opin Neurobiol*, *7*(2), 217-227. https://doi.org/10.1016/s0959-4388(97)80010-4

Nadel, L., Winocur, G., Ryan, L., & Moscovitch, M. (2007). Systems consolidation and hippocampus: two views. *Debates in Neuroscience*, *1*(2). https://doi.org/10.1007/s11559-007-9003-9

Nawa, N. E., & Ando, H. (2019). Effective connectivity within the ventromedial prefrontal cortex-hippocampus-amygdala network during the elaboration of emotional autobiographical memories. *Neuroimage*, *189*. https://doi.org/10.1016/j.neuroimage.2019.01.042

Nawa, N. E., & Ando, H. (2020). Effective connectivity during autobiographical memory search. *Brain Behav*, *10*(8), e01719. https://doi.org/10.1002/brb3.1719

Neunuebel, J. P., & Knierim, J. J. (2014). CA3 retrieves coherent representations from degraded input: direct evidence for CA3 pattern completion and dentate gyrus pattern separation. *Neuron*, *81*(2), 416-427. https://doi.org/10.1016/j.neuron.2013.11.017

Noppeney, U., Patterson, K., Tyler, L. K., Moss, H., Stamatakis, E. A., Bright, P., Mummery, C., & Price, C. J. (2007). Temporal lobe lesions and semantic impairment: a comparison of herpes simplex virus encephalitis and semantic dementia. *Brain*, *130*(Pt 4), 1138-1147. https://doi.org/10.1093/brain/awl344

Norris, D. G., & Polimeni, J. R. (2019). Laminar (f)MRI: A short history and future prospects. *Neuroimage*, *197*, 643-649. https://doi.org/10.1016/j.neuroimage.2019.04.082

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain research*, *34*(1). https://doi.org/10.1016/0006-8993(71)90358-1

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press.

O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Velasco-Hernández, G., Krpalkova, L., Riordan, D., & Walsh, J. (2019). Deep Learning vs. Traditional Computer Vision. *Computer Vision Conference*. https://doi.org/10.1007/978-3-030-17795-9

Oliva, A., Fernandez-Ruiz, A., Buzsaki, G., & Berenyi, A. (2016). Spatial coding and physiological properties of hippocampal neurons in the Cornu Ammonis subregions. *Hippocampus*, *26*(12), 1593-1607. https://doi.org/10.1002/hipo.22659

Olman, C. A., Inati, S., & Heeger, D. J. (2007). The effect of large veins on spatial localization with GE BOLD at 3 T: Displacement, not blurring. *Neuroimage*, *34*(3). https://doi.org/10.1016/j.neuroimage.2006.08.045

Olman, C. A., & Yacoub, E. (2011). High-field FMRI for human applications: an overview of spatial resolution and signal specificity. *Open Neuroimag J*, *5*, 74-89. https://doi.org/10.2174/1874440001105010074

Olsen, R. K., Carr, V. A., Daugherty, A. M., La Joie, R., Amaral, R. S. C., Amunts, K., Augustinack, J. C., Bakker, A., Bender, A. R., Berron, D., Boccardi, M., Bocchetta, M., Burggren, A. C., Chakravarty, M. M., Chetelat, G., de Flores, R., DeKraker, J., Ding, S. L., Geerlings, M. I., . . . Hippocampal Subfields, G. (2019). Progress update from the hippocampal subfields group. *Alzheimers Dement (Amst)*, *11*, 439-449. https://doi.org/10.1016/j.dadm.2019.04.001

Olsen, R. K., Moses, S. N., Riggs, L., & Ryan, J. D. (2012). The hippocampus supports multiple cognitive processes through relational binding and comparison. *Frontiers in human neuroscience*, *6*. https://doi.org/10.3389/fnhum.2012.00146

Ortiz-Tudela, J., Turan, G., Vilas, M., Melloni, L., & Shing, Y. L. (2024). Schema-driven prediction effects on episodic memory across the lifespan. *Philos Trans R Soc Lond B Biol Sci*, *379*(1913), 20230401. https://doi.org/10.1098/rstb.2023.0401

Palombo, D. J., Amaral, R. S., Olsen, R. K., Müller, D. J., Todd, R. M., Anderson, A. K., & Levine, B. (2013). KIBRA polymorphism is associated with individual differences in hippocampal subregions: evidence from anatomical segmentation using high-resolution MRI. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *33*(32). https://doi.org/10.1523/JNEUROSCI.1406-13.2013

Palombo, D. J., Bacopulos, A., Amaral, R. S. C., Olsen, R. K., Todd, R. M., Anderson, A. K., & Levine, B. (2018). Episodic autobiographical memory is associated with variation in the size of hippocampal subregions. *Hippocampus*, *28*(2). https://doi.org/10.1002/hipo.22818

Park, S., & Chun, M. M. (2009). Different roles of the parahippocampal place area (PPA) and retrosplenial cortex (RSC) in panoramic scene perception. *Neuroimage*, *47*(4), 1747-1756. https://doi.org/10.1016/j.neuroimage.2009.04.058

Petcharunpaisan, S., Ramalho, J., & Castillo, M. (2010). Arterial spin labeling in neuroimaging. *World J Radiol*, *2*(10), 384-398. https://doi.org/10.4329/wjr.v2.i10.384

Peters, J., & D'Esposito, M. (2016). Effects of Medial Orbitofrontal Cortex Lesions on Self-Control in Intertemporal Choice. *Curr Biol*, *26*(19), 2625-2628. https://doi.org/10.1016/j.cub.2016.07.035

Pfaffenrot, V., Bouyeure, A., Gomes, C. A., Kashyap, S., Axmacher, N., & Norris, D. (2024). Characterizing BOLD activation patterns in the human hippocampus with laminar fMRI. *bioRxiv*. https://doi.org/10.1101/2024.07.04.602065

Piefke, M., Weiss, P. H., Zilles, K., Markowitsch, H. J., & Fink, G. R. (2003). Differential remoteness and emotional tone modulate the neural correlates of autobiographical memory. *Brain*, *126*(Pt 3), 650-668. https://doi.org/10.1093/brain/awg064

Pipitone, J., Park, M. T., Winterburn, J., Lett, T. A., Lerch, J. P., Pruessner, J. C., Lepage, M., Voineskos, A. N., Chakravarty, M. M., & Alzheimer's Disease Neuroimaging, I. (2014). Multi-atlas segmentation of the whole hippocampus and subfields using multiple automatically generated templates. *Neuroimage*, *101*, 494-512. https://doi.org/10.1016/j.neuroimage.2014.04.054

Poiret, C., Bouyeure, A., Patil, S., Grigis, A., Duchesnay, E., Faillot, M., Bottlaender, M., Lemaitre, F., & Noulhiane, M. (2023). A fast and robust hippocampal subfields segmentation: HSF revealing lifespan volumetric dynamics. *Front Neuroinform*, *17*, 1130845. https://doi.org/10.3389/fninf.2023.1130845

Polimeni, J. R., Fischl, B., Greve, D. N., & Wald, L. L. (2010). Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. *Neuroimage*, *52*(4), 1334-1346. https://doi.org/10.1016/j.neuroimage.2010.05.005

Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Curr Biol*, *23*(17), R764-773. https://doi.org/10.1016/j.cub.2013.05.041

Price, C. J. (2000). The anatomy of language: contributions from functional neuroimaging. *J Anat*, *197 Pt 3*(Pt 3), 335-359. https://doi.org/10.1046/j.1469-7580.2000.19730335.x

Pykett, I. L., & Rzedzian, R. R. (1987). Instant images of the body by magnetic resonance. *Magn Reson Med*, *5*(6), 563-571. https://doi.org/10.1002/mrm.1910050607

Ranganath, C., & Ritchey, M. (2012). Two cortical systems for memory-guided behaviour. *Nature reviews. Neuroscience*, *13*(10), 713–726. https://doi.org/10.1038/nrn3338

Ranjeva, J. P., Franconi, J. M., Manelfe, C., & Berry, I. (1997). Magnetization transfer with echo planar imaging. *MAGMA*, *5*(4), 259-265. https://doi.org/10.1007/BF02595043

Read, M. L., Berry, S. C., Graham, K. S., Voets, N. L., Zhang, J., Aggleton, J. P., Lawrence, A. D., & Hodgetts, C. J. (2024). Scene-selectivity in CA1/subicular complex: Multivoxel pattern analysis at 7T. *Neuropsychologia*, *194*, 108783. https://doi.org/10.1016/j.neuropsychologia.2023.108783

Reagh, Z. M., & Ranganath, C. (2023). Flexible reuse of cortico-hippocampal representations during encoding and recall of naturalistic events. *Nat Commun*, *14*(1), 1279. https://doi.org/10.1038/s41467-023-36805-5

Rebola, N., Carta, M., & Mulle, C. (2017). Operation and plasticity of hippocampal CA3 circuits: implications for memory encoding. *Nat Rev Neurosci*, *18*(4), 208-220. https://doi.org/10.1038/nrn.2017.10

Redondo, R. L., & Morris, R. G. (2011). Making memories last: the synaptic tagging and capture hypothesis. *Nat Rev Neurosci*, *12*(1), 17-30. https://doi.org/10.1038/nrn2963

Rekkas, P. V., & Constable, R. T. (2005). Evidence that autobiographic memory retrieval does not become independent of the hippocampus: an fMRI study contrasting very recent

with remote events. *J Cogn Neurosci*, *17*(12), 1950-1961. https://doi.org/10.1162/089892905775008652

Richmond, L. L., Gold, D. A., & Zacks, J. M. (2017). Event perception: Translations and applications. *J Appl Res Mem Cogn*, *6*(2), 111-120. https://doi.org/10.1016/j.jarmac.2016.11.002

Rissman, J., Greely, H. T., & Wagner, A. D. (2010). Detecting individual memories through the neural decoding of memory states and past experience. *Proc Natl Acad Sci U S A*, *107*(21), 9849-9854. https://doi.org/10.1073/pnas.1001028107

Ritchey, M., Libby, L. A., & Ranganath, C. (2015). Cortico-hippocampal systems involved in memory and cognition: the PMAT framework. *Progress in brain research*, *219*, 45–64. https://doi.org/10.1016/bs.pbr.2015.04.001

Robertson, R. G., Rolls, E. T., & Georges-Fran ois, P. (1998). Spatial view cells in the primate hippocampus: effects of removal of view details. *J Neurophysiol*, *79*(3), 1145-1156. https://doi.org/10.1152/jn.1998.79.3.1145

Robin, J., Buchsbaum, B. R., & Moscovitch, M. (2018). The Primacy of Spatial Context in the Neural Representation of Events. *J Neurosci*, *38*(11), 2755-2765. https://doi.org/10.1523/JNEUROSCI.1638-17.2018

Robin, J., & Moscovitch, M. (2017). Details, gist and schema: hippocampal–neocortical interactions underlying recent and remote episodic and spatial memory. *Current Opinion in Behavioral Sciences*, *17*. https://doi.org/10.1016/j.cobeha.2017.07.016

Robin, J., & Olsen, R. K. (2019). Scenes facilitate associative memory and integration. *Learn Mem*, *26*(7), 252-261. https://doi.org/10.1101/lm.049486.119

Rockland, K. S. (2019). What do we know about laminar connectivity? *Neuroimage*, *197*, 772-784. https://doi.org/10.1016/j.neuroimage.2017.07.032

Rockland, K. S., & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Res*, *179*(1), 3-20. https://doi.org/10.1016/0006-8993(79)90485-2

Rodriguez, P. F. (2010). Neural decoding of goal locations in spatial navigation in humans with fMRI. *Hum Brain Mapp*, *31*(3), 391-397. https://doi.org/10.1002/hbm.20873

Roediger, H. L., & Mcdermott, K. B. (1995). Creating False Memories - Remembering Words Not Presented in Lists. *Journal of Experimental Psychology-Learning Memory and Cognition*, *21*(4), 803-814. https://doi.org/Doi 10.1037/0278-7393.21.4.803

Rolls, E. T. (2023). Hippocampal spatial view cells for memory and navigation, and their underlying connectivity in humans. *Hippocampus*, *33*(5). https://doi.org/10.1002/hipo.23467

Rolls, E. T., & Treves, A. (1994). Neural networks in the brain involved in memory and recall. *Prog Brain Res*, *102*, 335-341. https://doi.org/10.1016/S0079-6123(08)60550-6

Romero, J. E., Coupe, P., & Manjon, J. V. (2017). HIPS: A new hippocampus subfield segmentation method. *Neuroimage*, *163*, 286-295. https://doi.org/10.1016/j.neuroimage.2017.09.049

Rose, D., & Lowe, I. (1982). Dynamics of adaptation to contrast. *Perception*, *11*(5), 505–528. https://doi.org/10.1068/p110505

Rosenbaum, R. S., Moscovitch, M., Foster, J. K., Schnyer, D. M., Gao, F., Kovacevic, N., Verfaellie, M., Black, S. E., & Levine, B. (2008). Patterns of autobiographical memory loss in medial-temporal lobe amnesic patients. *J Cogn Neurosci*, *20*(8), 1490-1506. https://doi.org/10.1162/jocn.2008.20105

Rowland, C. A. (2014). The effect of testing versus restudy on retention: a meta-analytic review of the testing effect. *Psychol Bull*, *140*(6), 1432-1463. https://doi.org/10.1037/a0037559

Roy, M., Shohamy, D., & Wager, T. D. (2012). Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends Cogn Sci*, *16*(3), 147-156. https://doi.org/10.1016/j.tics.2012.01.005

Ruthotto, L., Kugel, H., Olesch, J., Fischer, B., Modersitzki, J., Burger, M., & Wolters, C. H. (2012). Diffeomorphic susceptibility artifact correction of diffusion-weighted magnetic

resonance images. *Phys Med Biol*, *57*(18), 5715-5731. https://doi.org/10.1088/0031-9155/57/18/5715

Schacter, D. L. (2012). Adaptive constructive processes and the future of memory. *Am Psychol*, *67*(8), 603-613. https://doi.org/10.1037/a0029869

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nat Rev Neurosci*, *8*(9), 657-661. https://doi.org/10.1038/nrn2213

Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., & Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. *J Neurosci*, *35*(41), 13904-13911. https://doi.org/10.1523/JNEUROSCI.2618-15.2015

Schlichting, M. L., Mumford, J. A., & Preston, A. R. (2015). Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nat Commun*, *6*, 8151. https://doi.org/10.1038/ncomms9151

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry*, *20*(1), 11-21. https://doi.org/10.1136/jnnp.20.1.11

Sekeres, M. J., Winocur, G., & Moscovitch, M. (2018). The hippocampus and related neocortical structures in memory transformation. *Neurosci Lett*, *680*, 39-53. https://doi.org/10.1016/j.neulet.2018.05.006

Sellitto, M., Ciaramelli, E., & di Pellegrino, G. (2010). Myopic discounting of future rewards after medial orbitofrontal damage in humans. *J Neurosci*, *30*(49), 16429-16436. https://doi.org/10.1523/JNEUROSCI.2516-10.2010

Seltzer, B., & Pandya, D. N. (1991). Post-rolandic cortical projections of the superior temporal sulcus in the rhesus monkey. *J Comp Neurol*, *312*(4), 625-640. https://doi.org/10.1002/cne.903120412

Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., Yu, B. M., Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., & Yu, B. M. (2022). Feedforward and feedback interactions between visual

cortical areas use different population activity patterns. *Nature Communications*, *13*(1). https://doi.org/10.1038/s41467-022-28552-w

Shao, X., Li, A., Chen, C., Loftus, E. F., & Zhu, B. (2023). Cross-stage neural pattern similarity in the hippocampus predicts false memory derived from post-event inaccurate information. *Nat Commun*, *14*(1), 2299. https://doi.org/10.1038/s41467-023-38046-y

Shipp, S. (2023). Computational components of visual predictive coding circuitry. *Front Neural Circuits*, *17*, 1254009. https://doi.org/10.3389/fncir.2023.1254009

Siapas, A. G., Lubenov, E. V., & Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron*, *46*(1), 141-151. https://doi.org/10.1016/j.neuron.2005.02.028

Simons, J. S., Verfaellie, M., Hodges, J. R., Lee, A. C., Graham, K. S., Koutstaal, W., Schacter, D. L., & Budson, A. E. (2005). Failing to get the gist: reduced false recognition of semantic associates in semantic dementia. *Neuropsychology*, *19*(3), 353-361. https://doi.org/10.1037/0894-4105.19.3.353

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J. M., & Matthews, P. M. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, *23 Suppl 1*, S208-219. https://doi.org/10.1016/j.neuroimage.2004.07.051

Soderlund, H., Moscovitch, M., Kumar, N., Mandic, M., & Levine, B. (2012). As time goes by: hippocampal connectivity changes with remoteness of autobiographical memory retrieval. *Hippocampus*, *22*(4), 670-679. https://doi.org/10.1002/hipo.20927

Spalding, K. L., Bergmann, O., Alkass, K., Bernard, S., Salehpour, M., Huttner, H. B., Bostrom, E., Westerlund, I., Vial, C., Buchholz, B. A., Possnert, G., Mash, D. C., Druid, H., & Frisen, J. (2013). Dynamics of hippocampal neurogenesis in adult humans. *Cell*, *153*(6), 1219-1227. https://doi.org/10.1016/j.cell.2013.05.002

Spalding, K. N., Jones, S. H., Duff, M. C., Tranel, D., & Warren, D. E. (2015). Investigating the Neural Correlates of Schemas: Ventromedial Prefrontal Cortex Is Necessary for Normal

Schematic Influence on Memory. *J Neurosci*, *35*(47), 15746-15751. https://doi.org/10.1523/JNEUROSCI.2767-15.2015

Spano, G., Pizzamiglio, G., McCormick, C., Clark, I. A., De Felice, S., Miller, T. D., Edgin, J. O., Rosenthal, C. R., & Maguire, E. A. (2020). Dreaming with hippocampal damage. *Elife*, *9*. https://doi.org/10.7554/eLife.56211

Spatz, W. B., Tigges, J., & Tigges, M. (1970). Subcortical projections, cortical associations, and some intrinsic interlaminar connections of the striate cortex in the squirrel monkey (Saimiri). *The Journal of comparative neurology*, *140*(2). https://doi.org/10.1002/cne.901400203

Spiers, H. J., & Maguire, E. A. (2006). Thoughts, behaviour, and brain dynamics during navigation in the real world. *Neuroimage*, *31*(4), 1826-1840. https://doi.org/10.1016/j.neuroimage.2006.01.037

Spiers, H. J., Maguire, E. A., & Burgess, N. (2001). Hippocampal amnesia. *Neurocase*, *7*(5), 357-382. https://doi.org/10.1076/neur.7.5.357.16245

Spreng, R. N., Mar, R. A., & Kim, A. S. (2009). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. *J Cogn Neurosci*, *21*(3), 489-510. https://doi.org/10.1162/jocn.2008.21029

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol Rev*, *99*(2), 195-231. https://doi.org/10.1037/0033-295x.99.2.195

Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Curr Opin Neurobiol*, *5*(2), 169-177. https://doi.org/10.1016/0959-4388(95)80023-9

Squire, L. R., Genzel, L., Wixted, J. T., & Morris, R. G. (2015). Memory consolidation. *Cold Spring Harb Perspect Biol*, *7*(8), a021766. https://doi.org/10.1101/cshperspect.a021766

Squire, L. R., van der Horst, A. S., McDuff, S. G., Frascino, J. C., Hopkins, R. O., & Mauldin, K. N. (2010). Role of the hippocampus in remembering the past and imagining the future.

Proc Natl Acad Sci U S A, *107*(44), 19044-19048. https://doi.org/10.1073/pnas.1014391107

Steinvorth, S., Corkin, S., & Halgren, E. (2006). Ecphory of autobiographical memories: an fMRI study of recent and remote memory retrieval. *Neuroimage*, *30*(1), 285-298. https://doi.org/10.1016/j.neuroimage.2005.09.025

Steinvorth, S., Levine, B., & Corkin, S. (2005). Medial temporal lobe structures are needed to re-experience remote autobiographical memories: evidence from H.M. and W.R. *Neuropsychologia*, *43*(4), 479-496. https://doi.org/10.1016/j.neuropsychologia.2005.01.001

Stirnberg, R., & Stocker, T. (2021). Segmented K-space blipped-controlled aliasing in parallel imaging for high spatiotemporal resolution EPI. *Magn Reson Med*, *85*(3), 1540-1551. https://doi.org/10.1002/mrm.28486

Summerfield, J. J., Hassabis, D., & Maguire, E. A. (2010). Differential engagement of brain regions within a 'core' network during scene construction. *Neuropsychologia*, *48*(5). https://doi.org/10.1016/j.neuropsychologia.2010.01.022

Svoboda, E., McKinnon, M. C., & Levine, B. (2006). The functional neuroanatomy of autobiographical memory: a meta-analysis. *Neuropsychologia*, *44*(12), 2189-2208. https://doi.org/10.1016/j.neuropsychologia.2006.05.023

Taylor, K. J., Henson, R. N., & Graham, K. S. (2007). Recognition memory for faces and scenes in amnesia: dissociable roles of medial temporal lobe structures. *Neuropsychologia*, *45*(11), 2428-2438. https://doi.org/10.1016/j.neuropsychologia.2007.04.004

Tigges, J., Spatz, W. B., & Tigges, M. (1973). Reciprocal point-to-point connections between parastriate and striate cortex in the squirrel monkey (Saimiri). *Journal of Comparative Neurology*, *148*(4). https://doi.org/10.1002/cne.901480406

Tompary, A., & Davachi, L. (2017). Consolidation Promotes the Emergence of Representational Overlap in the Hippocampus and Medial Prefrontal Cortex. *Neuron*, *96*(1), 228-241 e225. https://doi.org/10.1016/j.neuron.2017.09.005

Tranel, D., & Jones, R. D. (2006). Knowing “What” and Knowing “When”. *Journal of Clinical and Experimental Neuropsychology*, *28*(1), 43-66. https://doi.org/10.1080/13803390490919344

Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374-391. https://doi.org/10.1002/hipo.450040319

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., Witter, M. P., & Morris, R. G. (2007). Schemas and memory consolidation. *Science*, *316*(5821), 76-82. https://doi.org/10.1126/science.1135935

Tu, Z., & Bai, X. (2010). Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *IEEE Trans Pattern Anal Mach Intell*, *32*(10), 1744-1757. https://doi.org/10.1109/TPAMI.2009.186

Tu, Z., Zheng, S., Yuille, A. L., Reiss, A. L., Dutton, R. A., Lee, A. D., Galaburda, A. M., Dinov, I., Thompson, P. M., & Toga, A. W. (2007). Automated extraction of the cortical sulci based on a supervised learning approach. *IEEE Trans Med Imaging*, *26*(4), 541-552. https://doi.org/10.1109/TMI.2007.892506

Turk-Browne, N. B. (2019). The hippocampus as a visual area organized by space and time: a spatiotemporal similarity hypothesis. *Vision research*, *165*. https://doi.org/10.1016/j.visres.2019.10.007

Ugurbil, K. (2016). What is feasible with imaging human brain function and connectivity using functional magnetic resonance imaging. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *371*(1705). https://doi.org/10.1098/rstb.2015.0361

Ungerleider, L. G., & Haxby, J. V. (1994). 'What' and 'where' in the human brain. *Current opinion in neurobiology*, *4*(2). https://doi.org/10.1016/0959-4388(94)90066-3

Van Essen, D. C., & Gallant, J. L. (1994). Neural mechanisms of form and motion processing in the primate visual system. *Neuron*, *13*(1), 1-10. https://doi.org/10.1016/0896-6273(94)90455-3

van Kesteren, M. T., Fernandez, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proc Natl Acad Sci U S A*, *107*(16), 7550-7555. https://doi.org/10.1073/pnas.0914892107

van Kesteren, M. T., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in neurosciences*, *35*(4). https://doi.org/10.1016/j.tins.2012.02.001

van Kesteren, M. T. R., Rignanese, P., Gianferrara, P. G., Krabbendam, L., & Meeter, M. (2020). Congruency and reactivation aid memory integration through reinstatement of prior knowledge. *Sci Rep*, *10*(1), 4776. https://doi.org/10.1038/s41598-020-61737-1

Van Leemput, K., Bakkour, A., Benner, T., Wiggins, G., Wald, L. L., Augustinack, J., Dickerson, B. C., Golland, P., & Fischl, B. (2009). Automated segmentation of hippocampal subfields from ultra-high resolution in vivo MRI. *Hippocampus*, *19*(6), 549-557. https://doi.org/10.1002/hipo.20615

van Mourik, T., Koopmans, P. J., & Norris, D. G. (2019). Improved cortical boundary registration for locally distorted fMRI scans. *PLoS One*, *14*(11), e0223440. https://doi.org/10.1371/journal.pone.0223440

van Mourik, T., van der Eerden, J., Bazin, P. L., & Norris, D. G. (2019). Laminar signal extraction over extended cortical areas by means of a spatial GLM. *PLoS One*, *14*(3), e0212493. https://doi.org/10.1371/journal.pone.0212493

Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature reviews. Neuroscience*, *10*(11). https://doi.org/10.1038/nrn2733

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W., & Mishkin, M. (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

Vass, L. K., & Epstein, R. A. (2013). Abstract representations of location and facing direction in the human brain. *J Neurosci*, *33*(14), 6133-6142. https://doi.org/10.1523/JNEUROSCI.3873-12.2013

Viard, A., Piolino, P., Desgranges, B., Chetelat, G., Lebreton, K., Landeau, B., Young, A., De La Sayette, V., & Eustache, F. (2007). Hippocampal activation for autobiographical memories over the entire lifetime in healthy aged subjects: an fMRI study. *Cereb Cortex*, *17*(10), 2453-2467. https://doi.org/10.1093/cercor/bhl153

Viskontas, I. V., McAndrews, M. P., & Moscovitch, M. (2000). Remote episodic memory deficits in patients with unilateral temporal lobe epilepsy and excisions. *J Neurosci*, *20*(15), 5853-5857. https://doi.org/10.1523/JNEUROSCI.20-15-05853.2000

Vizioli, L., De Martino, F., Petro, L. S., Kersten, D., Ugurbil, K., Yacoub, E., Muckli, L., Vizioli, L., De Martino, F., Petro, L. S., Kersten, D., Ugurbil, K., Yacoub, E., & Muckli, L. (2020). Multivoxel Pattern of Blood Oxygen Level Dependent Activity can be sensitive to stimulus specific fine scale responses. *Scientific Reports*, *10*(1). https://doi.org/10.1038/s41598-020-64044-x

von Economo, C., & Koskinas, G. N. (1925). *Die Cytoarchitektonik der Hirnrinde des erwachsenen Menschen*. Springer.

Waehnert, M. D., Dinse, J., Weiss, M., Streicher, M. N., Waehnert, P., Geyer, S., Turner, R., & Bazin, P. L. (2014). Anatomically motivated modeling of cortical laminae. *Neuroimage*, *93 Pt 2*, 210-220. https://doi.org/10.1016/j.neuroimage.2013.03.078

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *Neuroimage*, *137*, 188-200. https://doi.org/10.1016/j.neuroimage.2015.12.012

Wang, H., Suh, J. W., Das, S. R., Pluta, J. B., Craige, C., & Yushkevich, P. A. (2013). Multi-Atlas Segmentation with Joint Label Fusion. *IEEE transactions on pattern analysis and machine intelligence*, *35*(3), 611–623. https://doi.org/10.1109/TPAMI.2012.143

Wang, H., & Yushkevich, P. A. (2013). Multi-atlas segmentation with joint label fusion and corrective learning-an open source implementation. *Frontiers in neuroinformatics*, *7*, 27. https://doi.org/10.3389/fninf.2013.00027

Wang, J., Nasr, S., Roe, A. W., & Polimeni, J. R. (2022). Critical factors in achieving fine-scale functional MRI: Removing sources of inadvertent spatial smoothing. *Hum Brain Mapp*, *43*(11), 3311-3331. https://doi.org/10.1002/hbm.25867

Wang, S. H., & Morris, R. G. (2010). Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. *Annu Rev Psychol*, *61*, 49-79, C41-44. https://doi.org/10.1146/annurev.psych.093008.100523

Warren, D. E., Jones, S. H., Duff, M. C., & Tranel, D. (2014). False recall is reduced by damage to the ventromedial prefrontal cortex: implications for understanding the neural correlates of schematic memory. *J Neurosci*, *34*(22), 7677-7682. https://doi.org/10.1523/JNEUROSCI.0119-14.2014

Warrington, O., Graedel, N. N., Callaghan, M. F., & Kok, P. (2024). Communication of perceptual predictions from the hippocampus to the deep layers of the parahippocampal cortex. *bioRxiv*. https://doi.org/10.1101/2024.03.28.587186

Webb, C. E., Turney, I. C., & Dennis, N. A. (2016). What's the gist? The influence of schemas on the neural correlates underlying true and false memories. *Neuropsychologia*, *93*(Pt A), 61-75. https://doi.org/10.1016/j.neuropsychologia.2016.09.023

Wisse, L. E., Kuijf, H. J., Honingh, A. M., Wang, H., Pluta, J. B., Das, S. R., Wolk, D. A., Zwanenburg, J. J., Yushkevich, P. A., & Geerlings, M. I. (2016). Automated Hippocampal Subfield Segmentation at 7T MRI. *AJNR. American journal of neuroradiology*, *37*(6). https://doi.org/10.3174/ajnr.A4659

Wisse, L. E. M., Daugherty, A. M., Olsen, R. K., Berron, D., Carr, V. A., Stark, C. E. L., Amaral, R. S. C., Amunts, K., Augustinack, J. C., Bender, A. R., Bernstein, J. D., Boccardi, M., Bocchetta, M., Burggren, A., Chakravarty, M. M., Chupin, M., Ekstrom, A., de Flores, R., Insausti, R., . . . Hippocampal Subfields, G. (2017). A harmonized segmentation protocol for hippocampal and parahippocampal subregions: Why do we need one and what are the key goals? *Hippocampus*, *27*(1), 3-11. https://doi.org/10.1002/hipo.22671

Witter, M. P., & Amaral, D. G. (1991). Entorhinal cortex of the monkey: V. Projections to the dentate gyrus, hippocampus, and subicular complex. *J Comp Neurol*, *307*(3), 437-459. https://doi.org/10.1002/cne.903070308

Witter, M. P., & Amaral, D. G. (2020). The entorhinal cortex of the monkey: VI. Organization of projections from the hippocampus, subiculum, presubiculum, and parasubiculum. *J Comp Neurol*, *529*(4), 828-852. https://doi.org/10.1002/cne.24983

Wolbers, T., & Buchel, 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., Wiener, J. M., Mallot, H. A., & Buchel, C. (2007). Differential recruitment of the hippocampus, medial prefrontal cortex, and the human motion complex during path integration in humans. *J Neurosci*, *27*(35), 9408-9416. https://doi.org/10.1523/JNEUROSCI.2146-07.2007

Wong-Riley, M. (1978). Reciprocal connections between striate and prestriate cortex in squirrel monkey as demonstrated by combined peroxidase histochemistry and autoradiography. *Brain research*, *147*(1). https://doi.org/10.1016/0006-8993(78)90781-3

Wood, M. L., & Henkelman, R. M. (1985). Truncation artifacts in magnetic resonance imaging. *Magn Reson Med*, *2*(6), 517-526. https://doi.org/10.1002/mrm.1910020602

Xie, L., Wisse, L. E. M., Wang, J., Ravikumar, S., Khandelwal, P., Glenn, T., Luther, A., Lim, S., Wolk, D. A., & Yushkevich, P. A. (2023). Deep label fusion: A generalizable hybrid multi-atlas and deep convolutional neural network for medical image segmentation. *Med Image Anal*, *83*, 102683. https://doi.org/10.1016/j.media.2022.102683

Yang, G., Pan, F., & Gan, W. B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature*, *462*(7275), 920-924. https://doi.org/10.1038/nature08577

Yeung, L. K., Olsen, R. K., Hong, B., Mihajlovic, V., D'Angelo, M. C., Kacollja, A., Ryan, J. D., & Barense, M. D. (2019). Object-in-place Memory Predicted by Anterolateral Entorhinal Cortex and Parahippocampal Cortex Volume in Older Adults. *J Cogn Neurosci*, *31*(5), 711-729. https://doi.org/10.1162/jocn_a_01385

Yu, Y., Huber, L., Yang, J., Jangraw, D. C., Handwerker, D. A., Molfese, P. J., Chen, G., Ejima, Y., Wu, J., & Bandettini, P. A. (2019). Layer-specific activation of sensory input and

predictive feedback in the human primary somatosensory cortex. *Sci Adv*, *5*(5), eaav9053. https://doi.org/10.1126/sciadv.aav9053

Yushkevich, P. A., Amaral, R. S., Augustinack, J. C., Bender, A. R., Bernstein, J. D., Boccardi, M., Bocchetta, M., Burggren, A. C., Carr, V. A., Chakravarty, M. M., Chetelat, G., Daugherty, A. M., Davachi, L., Ding, S. L., Ekstrom, A., Geerlings, M. I., Hassan, A., Huang, Y., Iglesias, J. E., . . . Hippocampal Subfields, G. (2015). Quantitative comparison of 21 protocols for labeling hippocampal subfields and parahippocampal subregions in in vivo MRI: towards a harmonized segmentation protocol. *Neuroimage*, *111*, 526-541. https://doi.org/10.1016/j.neuroimage.2015.01.004

Yushkevich, P. A., Pluta, J. B., Wang, H., Xie, L., Ding, S. L., Gertje, E. C., Mancuso, L., Kliot, D., Das, S. R., & Wolk, D. A. (2015). Automated volumetry and regional thickness analysis of hippocampal subfields and medial temporal cortical structures in mild cognitive impairment. *Hum Brain Mapp*, *36*(1), 258-287. https://doi.org/10.1002/hbm.22627

Yushkevich, P. A., Wang, H., Pluta, J., Das, S. R., Craige, C., Avants, B. B., Weiner, M. W., & Mueller, S. (2010). Nearly automatic segmentation of hippocampal subfields in in vivo focal T2-weighted MRI. *Neuroimage*, *53*(4). https://doi.org/10.1016/j.neuroimage.2010.06.040

Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychol Bull*, *127*(1), 3-21. https://doi.org/10.1037/0033-2909.127.1.3

Zeidman, P., Lutti, A., & Maguire, E. A. (2015). Investigating the functions of subregions within anterior hippocampus. *Cortex*, *73*, 240-256. https://doi.org/10.1016/j.cortex.2015.09.002

Zeidman, P., & Maguire, E. A. (2016). Anterior hippocampus: the anatomy of perception, imagination and episodic memory. *Nat Rev Neurosci*, *17*(3), 173-182. https://doi.org/10.1038/nrn.2015.24

Zeidman, P., Mullally, S. L., & Maguire, E. A. (2015). Constructing, Perceiving, and Maintaining Scenes: Hippocampal Activity and Connectivity. *Cereb Cortex*, *25*(10), 3836-3855. https://doi.org/10.1093/cercor/bhu266

Zeng, H., & Constable, R. T. (2002). Image distortion correction in EPI: comparison of field mapping with point spread function mapping. *Magn Reson Med*, *48*(1), 137-146. https://doi.org/10.1002/mrm.10200

Zwanenburg, J. J., Versluis, M. J., Luijten, P. R., & Petridou, N. (2011). Fast high resolution whole brain T2* weighted imaging using echo planar imaging at 7T. *Neuroimage*, *56*(4). https://doi.org/10.1016/j.neuroimage.2011.03.046

Appendix

I acknowledge the use of ChatGPT-3.5 and ChatGPT-4 (OpenAI, https://openai.com/) to:

- draft text in Sections 2.8.2.1. and 4.3.6., which I subsequently fact checked and manually edited;
- edit and proofread the final draft.