SUBCORTICAL CIRCUITS FOR PROCEDURAL LEARNING AND MEMORY

E.J. THOMPSON

THESIS SUBMITTED FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY IN NEUROSCIENCE

SUPERVISOR: MARCUS STEPHENSON-JONES

UNIVERSITY COLLEGE LONDON
SAINSBURY WELLCOME CENTRE
SEPTEMBER 2023

DECLARATION

I confirm that the work presented in my own. Where information has been derived from other sources, I confirm that this has been indicated in the text.

Emmett Thompson

ABSTRACT

This thesis explores roles for subcortical circuits in establishing and supporting memory of procedural skills. My investigation considers two distinct but overlapping themes. Firstly, I have explored thalamic and striatal circuits which support skill learning and production online: during wakeful behaviour. **(1)** I outline a novel behavioural paradigm and **(2)** show that both learning, and execution of this 5-step sequence task is contingent on the DLS. **(3)** I then provide evidence that a thalamic region, rILT, supports this function, providing inputs to striatum which can define both the structure and content of motor sequence control. I argue that this circuit may form part of a motor efference feedback loop, chaining action-to-action to support serial order motor skills. In the second part of this thesis, I have examined the processes which support the function of these circuits offline: during rest or sleep. While hippocampal replay is thought to be the substrate of consolidation for episodic memory, little is known about the offline processes that support procedural memory formation. To address this, **(4)** I first identify that an offline mechanism in the DLS is indeed critical for procedural consolidation. **(5)** Then, to determine the neural basis of this mechanism, I describe a novel replay detection method based on an unsupervised point process model. **(6)** I validate this method using ground truth data and argue that extensive methodological testing and unbiased approaches such as the one described, are essential for understanding the true extent of offline dynamics. **(7)** When applying my method, I find fine grained procedural sequences are replayed offline in DLS. This replay shared many features in common with previously identified hippocampal reactivations and persisted throughout all stages of motor skill learning **(8)** Finally, I demonstrate that both procedural memory consolidation and replay in the DLS are independent from hippocampus. Consequently, I find evidence that procedural consolidation is an entirely distinct process from other kinds of neural consolidation.

IMPACT STATEMENT

Understanding how neural circuits support memory formation is key to understanding how the brain orchestrates behaviour. This thesis provides an investigation of how neural circuits in the brain, particularly striatal circuits, support the formation of procedural memories, a typically neglected memory system within the field. Procedural memories are ubiquitous in animal behaviour. This memory system shapes and optimises the way animals move and interact with their environment. Besides supporting daily interactions with the world, In the human brain, our procedural memory system is responsible for our unique ability to produce coherent speech as well as the skilled manipulations needed by craftsmen, musicians, athletes, doctors and even scientists.

This thesis takes a multifaceted approach, exploring both the circuits that support awake learning and execution of motor sequences, as well as the dynamics offline during sleep that support memory consolidation. To study this I have established a novel motor sequence task and training paradigm which allows mice to learn a relatively complicated, but highly controllable movement sequence. The ability to teach a complex movement sequence in mice, and the flexible mutability of the task makes this a valuable contribution to the many excellent paradigms that have been developed previously for the study of procedural behaviours.

Using this task to study awake learning and execution of procedural memory, I find evidence to support a chaining model, in which a subcortical loop circuit involving thalamus and sensory motor striatum constructs procedural movements by linking elemental movement motifs together. These results build on a growing literature suggesting that procedural memories are constructed subcortically from elemental building blocks that can be reused and reordered like syllables in speech. Understanding how these syllable elements are chained together by neural circuitry, is the first step towards understanding how the brain flexibly constructs behaviours from such elements.

In the work in this thesis on role of sleep for memory consolidation, I find evidence of neural replay of procedural activity in the striatum. This work builds on previous knowledge related to other memory systems, providing evidence that procedural memory is supported by similar, but crucially independent, mechanisms. To my knowledge this work represents first direct observations of neural activity related to offline procedural consolidation. These findings may provide a foundation for future research aiming to understand how memory is coordinated within local and global brain systems.

This work can also have an impact outside of fundamental neuroscience research. Many neural pathologies present as failures of procedural coordination including aphasias that cause incomprehensible speech and movement disorders such as Parkinson's disease. Improved knowledge of how these neural systems function normally can help clinicians develop treatments to alleviate symptoms of these pathologies caused by disfunction in procedural circuits. Additionally, the work in this thesis on sleep related memory consolidation mechanisms could be a helpful addition to a growing literature driving froward the development of treatments and technologies combating age and disease related sleep and memory degradation. With an ageing global population, research in this area is increasingly worthwhile.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support of a number of people. I am incredibly grateful to all of them.

Sainsbury Wellcome Centre.

Firstly, I would like to say thank you to all the staff, scientists and students at the SWC who assisted me throughout this process. Completing my PhD at the SWC has been an amazing experience because of incredible people who work there.

Marcus.

You have built an amazing team and it's been real privilege to be part of a young lab with so much ambition. Thank you for your advice and direction throughout; at times when I have been lost in the details your clear thinking has been invaluable. Thank you also for your mentorship and energy; you have backed me 100% during the ups and downs and helped me keep pushing the project all the way to the finish line. We have been a good team and I am grateful for that.

The lab.

I have been very fortunate to have really amazing lab mates through this process. Despite my bizarre lunch habits, and often manic work style you have more than put up with me over the last 5 years. Fred, Hernando, Francesca, Lars, Svenja, Laura, Sthita, Georgie, Matthew, Caterina, Yvonne, Chris, Javsin, Joseph and Benjamin, I've been so lucky to get to know and learn from every one of you. Thank you for all the insight, encouragement, and joy over the last 5 years. Fred & Hernando, thank you especially for all of your time and patience in the early days.

Thank you also to the PPsleep team: Will, Clementine, Tom and Rodrgio for all of your efforts towards the project. I learned a huge amount from you, and it has been great fun working together!

Georgie, Jasvin, Joseph and Benjamin.

Thank you (merci beaucoup) for all for your hard work on this project, and for being foolish enough to trust me as a teacher. Thank you for helping me learn how to be a better mentor; working with you all and helping you improve as scientists has been the most fulfilling part of this process. Also, Georgie and Jasvin, thank you for putting up with me 'hitting the panic button' every 5 minutes throughout most of my third year!

My fellow PhD cohort.

The PhD process can feel very lonely and isolating at times especially during a global pandemic. Sharing the experience with all of you made it far less lonely. Thank you for being there to give perspective and friendship during the highs and lows. Looking forward to future lunch time conversations about dolphins.

To my friends.

Tash, Joe, Sam, Ollie, Marck, Jamie, Dan. Its finally done! I promise to be a better friend and a worse scientist from now on. Thank you all for providing an outlet for me and for being understanding while I did this silly thing even though most of you didn't have the benefit of knowing exactly what I did or why I did it. Thank you for keeping me grounded and making me laugh. Without each of you I would have surely lost my mind over the last 5 years.

Fred, Jean-Christophe, Emilien, Gauthier.

Merci beaucoup pour toutes les pauses en France. C'était précieux pour vraiment échapper à ce projet. Je suis juste triste de ne plus avoir d'excuse pour mes terribles progrès en français !

Alfie.

Thank you for providing the real motivation throughout this PhD…happy Christmas. Also, thanks for being so crap at squash. Knocking you around that court really helped whenever an experiment failed!

Mum and Dad.

Are you finally proud of me now!? I'm kidding of course, but I do hope you are because I want to be able to give something back to the two people who have given me far more support and love than I probably deserve. Mum, thank you for showing me that with relentless energy, motivation and drive you can do almost anything. I have relied on these traits so much throughout this process. Dad, thank you for instilling me with my academic curiosity, helping me to try and think differently and for proving to me that any old fool can write a thesis.

Eugenie.

You have put up with so much so that I could go through this process. I could not have done it without your constant love, support, and companionship. At times when I have been my own worst enemy you were always there for me. I cannot express how grateful I am to you. You have all my love and gratitude.

SUMMARY OF CONTRIBUTIONS

Experimental and project design was devised by Emmett Thompson (ET) and Marcus Stephenson-Jones (MSJ). Experimental work, data analysis and writing were done by the author but often with support from multiple contributors who deserve recognition (see below for specifics).

Chapter 4:

Surgeries and experimental setup were carried out by ET; animal training was done by ET and Georgina Mills (GM). Analysis was completed by ET.

Chapter 5:

Surgeries and animal training were completed by ET and GM. Analysis was done by ET. Histology was carried out by ET and GM.

Chapter 6:

Surgeries were performed by ET. Animal training was done by ET, GM and Jasvin Kaur (JK). Drug cannulation was performed by GM and JK under supervision of ET. Cannulation analysis was done by JK and ET. Neuropixel recording setup was completed by ET and Lars Rollik (LR). Surgical implants and neural recordings were carried out by ET. Adaptation of the PPseq model was by William Dorrell, Clementine Domine, Tom George, Rodrigo Carrasco Davis, LR and ET. PPseq application and analysis was carried out by ET. State space decoding and analysis was completed by Benjamin Waked (BW).

Chapter 7:

Surgeries were carried out by ET and MSJ. Mouse training was done by ET and Joseph Cohen. Surgical implants and neural recordings were carried out by ET. Analysis was completed by ET and BW.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF ABBREVIATIONS

BG Basal ganglia

DLS Dorsolateral Striatum

MSN Medium spiny neuron

VT Ventral thalamus

ILT Intralaminar thalamus

rILT Rostral intralaminar thalamus

UVA Thalamic nucleus Uvaeformis

PF Parafasicular thalamus

LFP Local field potential

REM Rapid eye movement

NREM Non-rapid eye movement

NREM2 Non-rapid eye movement stage 2

SWR Sharp wave ripples

TRN Thalamic reticular nucleus

SO Slow oscillations

δ Delta waves

GRM2 Glutamate metabotropic receptor 2

LMAN Lateral magnocellular nucleus of the anterior nidopallium

CHAPTER 1

1. GENERAL INTRODUCTION

A ubiquitous feature of animal behaviour is the regular and stereotyped expression of motor sequences (Lashley, 1951; Adams 1984). These behaviours are movement patterns, learned for a particular environmental context and performed with high spatial and temporal precision. Whether tying a shoelace or typing a password, motor sequences or skills constitute a large fraction of our daily behavioural repertoire. Learning a skill is marked by a common behavioural transition from novice to expert. For shoelace tying, any child that has graduated past Velcro straps to laces must learn a bafflingly complex hand, wrist, and finger ballet to fasten their shoes. At first tying a shoelace is slow, error prone and unsteady whilst requiring very high cognitive effort. From such humble fumbling eventually evolves a patterned, coordinated hand dance which one can execute with ease. Like any memory, motor skills are primarily learned and strengthened by repetition or practise (Ericsson *et al.* 1993). However, if awake practise can be considered online learning, procedural memory – memory for motor sequences – is also known to be supported by periods of offline consolidation, most significantly during sleep (Diekelmann & Born 2010; Rasch & Born 2013; Schmid *et al.* 2020). Motor sequences are highly prevalent in animal behaviour, and yet the mechanisms that underly procedural memory formation remain elusive. This thesis aims to address this by investigating the circuit mechanisms that support the formation of procedural memory both online and offline.

In this thesis I will address the question of how motor sequence behaviours are learned and controlled by neural circuits. I will focus on two aspects of this question. Firstly, I will explore the circuits that define these behaviours during awake expression with an emphasis on thalamo-striatal contributions to motor sequence expression. Secondly, I will explore the offline mechanisms that support consolidation of motor sequence memory, primarily focusing on striatal neural activity during sleep. With these aims, I will start with an exploration of the current experimental and theoretical background of motor sequence circuits (Chapter 2). In this section I will

provide evidence for the hypotheses that a circuit involving rostral intralaminar thalamus underlies motor skill execution, and that consolidation of motor sequence memory is defined by neural reactivations in striatum. In chapter 3 I will outline the materials and methods used in my own experiments. In chapter 4 I will introduce the novel motor sequence task used throughout my experiments and demonstrate that mouse behaviour on this task has the characteristics of a motor sequence behaviour including being dependent on dorsolateral striatum for both learning and execution. In Chapter 5 I will explore the role of thalamus in this behaviour. In chapter 6 I will show that learning and execution of this task is supported by an offline mechanism and investigate suitable methods for searching for the neural substrate of this mechanism, including describing a novel method for replay detection. In Chapter 7, I will explore offline striatal activity in depth including investigating whether these dynamics are independent of hippocampus. At the conclusion of each chapter, I will discuss specific implications of the experiments I have described and finally, in chapter 8, I conclude with a general discussion. In this final section I will set out my results in broader context with existing literature in the field, comment on future directions for this work and discuss the broad implications of this thesis to our understanding of how procedural memories are formed in the brain.

CHAPTER 2

2. LITERATURE REVIEW

2.1 Motor sequences

The ability of the brain to learn unique complex behavioural sequences arises through the ability to compose discrete elemental motifs of movement into sequential patterns. This structuring exists at multiple timescales. For example, while the fine graine moment-to-moment structuring of movements may be controlled by motor pattern generators, for example in the spinal cord (Tresch et al. 1999), behaviour over minutes can also be described as a higher order sequence of predicable components (Wiltschko *et al.* 2015). Consider the granularity of control in the production of speech. Over short timescales, when forming a simple syllable, the fine-scale kinematics of motor control are a complex sequence of facial, glossal, laryngeal and oesophageal muscle contractions. Over longer timescales, these muscle sequences are combined forming modular representations. Syllables. These modules are then themselves sequenced into words and phrases. How the brain composes motor sequences at multiple scales has been a focus of scientific debate in recent history. In the simplest models, behaviour is learned and controlled by a synfire chain; a neural domino effect whereby the serial order of behaviour is defined by propagation along a single sequential circuit (each element triggering the next), usually with sensory or motor feedback providing the links in this chain. While elegant, the simplicity of these models produces multiple limitations that make them inconsistent with observations about real behaviour. Firstly, domino-like chains should be extremely accurate yet in real behaviour errors are relatively frequent – for example sequence errors such as 'spoonerisms' in speech – suggesting this model is an oversimplification (Lashley, 1951). Also, while a synfire chain mechanisms should always require minimal preparation (the movement sequence is already preconfigured in the chain) there is a great deal of evidence that planning does occur in the brain. For instance, the time to initiate movement sequences increases with the length or complexity of the sequence and both movement properties and neural activity have been shown to have

anticipatory features, which predict future movements (reviewed in Rosenbaum *et al*. 2007). Hence, rather than triggering preconfigured chains, movements are composed in a way that implies the brain prepares them in a more flexible manner. The insufficiency of preconfigured synfire chains is perhaps best illustrated when we consider this inflexibility. While a serial order chain is fixed, in motor sequence behaviours, elemental motifs can in fact be rearranged, reused, and modified to generate novel behavioural sequences without extensive relearning of each element (Tervo *et al*. 2016). For example, different kinds of tennis serve share many common movement elements but are defined by slight adaptations in wrist or arm position during striking of the ball. Hence, for the neural control of motor sequences there must be flexible organisation of these elements – how can a tennis player learn a new serve without overwriting their current one?

A prevalent idea, succeeding more simple serial order models, is that motor sequences are organised via a hierarchical configuration. This is compelling as it captures the notion that multiple scales of compositional planning exist; in these models different scales of behaviour are choreographed independently but in parallel. For example, a higher order queue (first, second, third…) could trigger downstream competing motor elements (Korynsheva *et al*. 2019). At the bottom level of the hierarchy, fine scale kinematic representations dominate. Moving up the hierarchy, these fine scale motor elements are increasingly represented together, combined into modular units often called “motor chunks” (Lashley 1951; Rosenbaum *et al*. 1983; Dezfouli & Balleine 2012). At the top of the hierarchy, actions are represented in broad terms – with little information about fine scale movements remaining. With this kind of organisation, each modular chunk can activate the circuitry needed to generate its composite motor elements. Hence, complex sequences can be formed by adding modules together and novel sequences formed via recombination’s of these modules. Hierarchical control is thought to be a computationally efficient strategy. Recent modelling

work has shown that these systems can reduce action policy complexity and memory load when selecting actions (Lai *et al*. 2022). Further, hierarchical models allow for more flexible learning, independently at different sequential timescales (Maes *et al*. 2021). Experimental evidence that the brain contains a hierarchical controller is also beginning to crystalise. Behavioural work suggests both humans and animals perform motor chunking during sequence learning (Graybiel 1998; Sakai *et al* 2003). Also, multiple rodent studies have illustrated that there are circuits which can control high level behavioural ordering, independent of fine scale motor kinematics (Berridge *et al*. 1987; Jin and Costa, 2015; Geddes *et al*. 2018; Markowitz *et al*. 2018). Further still, multiple studies have shown the brain has abstract representation of action which govern high level behavioural structuring independent from even broad movement control (Wilson *et al*. 2014; Vaidya *et al*. 2021; Samborska *et al*. 2022). Existence of circuits which disambiguate these scales points towards a hierarchy of control mechanisms.

Pragmatically, it is important to recognise that the distinction I have drawn, between serial order and hierarchical models, is to some extent contrived. While it is accepted that simple synfire chains alone cannot account for motor sequence behaviour, serial order circuits do feature in motor sequence control. For example, in the songbird vocal learning circuit, song syllables are thought to be driven by preconfigured synfire chains in singing cortex 'HVC' (Fee & Goldberg 2011). Even in planning models such as competitive queuing a temporal structure (a basic chain of neurons) is required (Korynsheva *et al*. 2019) and the elements of motor activity – which planning circuits order and arrange – are still formed by chain like cascades of neural activity. Hence, modern notions of motor control systems are perhaps more accurately (though overly simplistically) described as a mixture – wherein hierarchical systems can be built out of serial order circuit components. Equally, the degree to which the brain implements control based on hierarchical planning or domino-like serial chaining may depend on the behaviour being controlled

or even the stage of learning. For instance, rodents learning rewarded locations in a cross maze will first direct their behaviour using flexible allocentric (place) coordinates, but later develop an egocentric (response) strategy (Tolman *et al*. 1946). Similarly, it has been shown in both human and non-human primate subjects that sequence learning is initially nonspecific to effector (generalisable) but gradually becomes specific to the overtrained effector (Hikosaka *et al*. 1995, Bapi *et al*. 2000, Rand *et al*. 2000). This transition is skill acquisition and is marked by improved movement speed and accuracy at the cost of flexibility. Control of these phases is thought to occur in distinct parallel circuits (Hikosaka *et al*. 1999) and, since accuracy increases at the cost of flexibility, this may be a shift towards circuits that rely more on serial order control.

In summary, motor sequence behaviours can be described as composed from elemental movement motifs. Control of this composition may be by simple chaining, though evidence from behaviour suggests the brain also has more flexible hierarchical control mechanisms based on planning. These control systems are not entirely separable, and it seems that different behaviours and even learning stages are under different kinds of motor sequence control. In this thesis I will focus on skill learning. Hence, I am concerned with the circuits that control behaviour in a motor coordinate or egocentric framework. In the following sections I aim to explore the circuits that contribute to learning and production of this form of motor control in more detail.

2.2 Neural circuits for motor skill learning and execution

Having now introduced some of the dominant theoretical ideas about how motor sequence behaviours may be organised by the brain. Next, I will introduce the brain regions and neural circuits that are believed to be involved. Learning a motor skill comprises selecting a series of actions and

then evaluating these actions for their efficacy. Hence, acquiring a motor skill is a reinforcement learning process whereby particular movement elements are associated together into a chain (Daw *et al*. 2005; Frank, 2011). In this thesis, I will focus primarily on neural circuits relating to the basal ganglia (BG). The BG are an evolutionarily conserved group of subcortical nuclei which are known to be critical for reinforcement learning and are play an important role in action selection and evaluation (Stephenson-Jones *et al*. 2011). Degeneration here in humans is known to lead to motor deficits related to performing and learning coherent motor sequences (Agostino *et al*. 1992; Bhatia & Marsden 1994; Jackson *et al*. 1995; Freeman et al. 1996; Laforce & Doyon 2001; Vicente et al. 2020).

In particular, I will consider the role of the striatum. The striatum is the major input nucleus, and the largest nucleus of the BG (Lanciego *et al*. 2012). As well as receiving dopaminergic modulation – required for reinforcement learning (Cox & Witten 2019) – the striatum receives extensive sensorimotor inputs from cortex and thalamus (Lanciego *et al*. 2012; Haber, 2016, Hunnicutt *et al*. 2016). These inputs define this brain region as an important integrative node in the distributed mammalian motor control circuit. Perhaps more significant are the downstream targets of the striatum which facilitate dual influence over motor control. Firstly, via BG output nuclei, the striatum has near direct influence over brainstem and midbrain motor nuclei. Secondly, via thalamic feedback the striatum can modulate higher order motor regions such as primary motor cortex. Through this connectivity, the striatum is in an ideal position to form state action associations for learning and executing motor sequences at multiple scales.

2.2.1 Sensorimotor striatum as a locus for motor sequence learning and execution

There is strong evidence that suggests the striatum is an essential component of a mammalian motor sequence circuit. In humans, fMRI imaging has shown

that learned finger tapping sequences engage sensorimotor regions of striatum (Andersen, Madsen and Siebner, 2020). Motor related coding in this region is highly stable across days as would be expected from a region responsible for storing stereotyped movements (Jensen *et al* 2022). Indeed, inactivation of these sensorimotor striatal regions in primates disrupts similar learned finger tapping sequences (Miyachi et al., 1997) and in rodents, lesions to equivalent sensorimotor striatum – the dorsolateral striatum (DLS) – impair both learning and execution of trained motor sequences (Berridge & Whishaw, 1992; Yin 2010; Wolff *et. al.* 2022*)*.

However, while it is broadly agreed that the striatum plays some kind of role in motor sequence learning, a key question is what this exact role is. How does the striatum contribute to motor control? And to what extent is this region in control of sequential motor behaviours? One particularly compelling hypothesis, which captures much of the debate around this question, is that the role the striatum depends on the particular challenges or features of motor learning (Dwhale *et al.* 2021). Viewed within a reinforcement learning paradigm, the striatum is believed to form state action associations (Sutton & Barto 2018). Hence, guided by specific reinforcement during learning, a mapping can be formed between features of state information, and some kind of actionable downstream output. Put more simply, the role of the striatum in motor learning may be flexible, and adaptable to the requirements of a given behaviour. In some circumstances these mappings may be relatively simple. For example, the striatum is known to be essential for contextual action selection – mapping between a learned sensory cue and movement (Lauwereyns et al. 2002). These mappings can be highly specific suggesting the striatum (as a part of the BG) can learn to trigger movements from individual effectors. For example, encoding the direction, timing, and amplitude of saccadic eye movements (Hikosaka & Wurtz 1983). For motor sequence learning, a great deal of evidence suggests that for many learned behaviours, the level of striatal control is in terms of action selection –

ordering the syntax of behavioural chunks. For example, DLS lesions disrupt highly stereotyped grooming patterns, causing out-of-sequence grooming but without altering the specific individual actions involved in grooming (Berridge & Fentress 1987). Similarly, during extended periods of naturalistic behaviour, decoded motor syllables are organised into prolonged sequences of action with a predictable arrangement. Current syllable arrangements are highly predicted by broad fluctuations in DLS dopamine levels during past behaviour suggesting reinforcement guides the DLS to organise the syllable-to-syllable structure of behaviour via a dopamine dependent learning process (Markowitz *et al.* 2023). In line with this, lesion to striatum disrupts this broad structure, altering the order in which syllables are chained to each other, once again without altering the fine scale kinematics of these syllables (Markowitz et al. 2018). Finally, in mice trained on a two-stage lever pressing task, optogenetic manipulations of two genetically distinct populations in DLS were found to have differential control over sequence order. The principal projection neurons in the striatum, medium spiny neurons (MSNs), can be split into two genetically defined populations based on dopamine receptor subtypes: D1 and D2. Activation of D1 MSNs caused ectopic repetition of the current action, while D2 MSN stimulation caused animals to switch to the next movement in the sequence (Geddes et al. 2018). This evidence strongly suggests that the striatal microcircuit is specialised for (or can at least be adapted into) a higher order hierarchical controller, concerned with the broad structure of action.

However, though for many procedural behaviours the DLS is concerned with behavioural ordering, as mentioned before it seems that the motor control mapping the BG learns is dependent on the requirement of behaviour. For example, Studies have also demonstrated that BG circuits are able to modulate the overall vigor of ongoing sequential actions (Desmurget & Turner 2010; Rueda-Orozco & Robbe 2015; Park *et al.* 2020). This implies that besides controlling the timing and movement direction of effectors, BG

outputs can also adaptively modulate gain of that movement in terms of its speed and amplitude. Finally, recent evidence suggests that BG circuits can even learn to produce more complex control signals which specify the fine graine motor kinematics of learned procedural movements (Dhawale *et al*. 2021).

In summary, guided by reinforcement, the striatum can learn to control behaviour across multiple modal axes and at multiple scales. For procedural learning this control mapping is dependent on the challenges of the behaviour being learned, although it appears that often the striatum is required for organising the syntax of behaviour – syllable ordering. Evidence for the DLS as the choreographer of motor sequence behaviours at multiple scales is compelling. However, it is a limited explanation. The DLS is an essential brain region, but only as an integrative nexus in a distributed motor network. Indeed, the striatum is itself not spontaneously active (Lanciego *et al*. 2012) and hence, requires external glutamatergic drive to have any influence on behaviour. Understanding where this drive comes from is key to understanding how the DLS functions as part of wider motor sequence controlling circuitry.

2.2.2 Contribution of cortico-striatal projections to motor sequence learning

One source of dense glutamatergic innervation to DLS is motor cortex (Hunnicutt *et al*. 2016). Like striatum, fMRI imaging has shown this region is highly active during motor sequence tasks in humans (Kami *et al*. 1995; Weistler & Diedrichsen 2013; Yokoi *et al*. 2018). Also like striatum, neurons here are tuned to specific features of movement (Peters *et al*. 2014; Xu et al. 2009), for example, pyramidal cells in primate primary motor cortex encode reaching direction and amplitude during a sequential task (Zimnik & Churchland 2021). In novice rodents, both generalised lesion to motor cortex and pathway specific silencing of motor cortical projections to striatum block

refinement of actions into stereotyped motor skills, leading to the development of highly variable behaviours (Santos *et al*. 2015; Wolff *et*. *al*. 2022). This is in line with the evidence discussed in 2.2.1, that the DLS is the central nexus of motor sequence control. However, an alternate mechanism is that motor cortex alone can control motor sequence behaviours. If this is true, the striatum would only be needed for refining cortical connectivity during learning, or even more simply, the striatum could just be a downstream conduit, only necessary for passing on cortical control signals.

Evidence that the first of these alternate mechanisms could be true comes from Zebra finch song learning. Over the first few months of their life, these birds use a memorised template – their father's song – to refine random vocal babbling into a highly stereotyped melody. Once learned, this song is entirely dependent on equivalent motor cortex, 'HVC', projections to another forebrain nucleus RA. In this circuit, the avian homologue of the striatum, area X, is only needed during learning and functions to bias HVC-RA plasticity towards correct song generation (Fee & Goldberg 2011; Kornfeld *et al*. 2020; Xiao & Roberts 2021). Once the song sequence has been learned and stereotyped, lesions to area X do not impair or alter the song (Scharff *et al*. 2000; Sanchez-Valpuesta *et al*. 2019). Similar mechanisms have been proposed in the mammalian brain. In one such model, the BG entrain cortical plasticity via thalamic feedback such that after learning, cortical dynamics can generate sequences independently from subcortical structures (Logacio *et al*. 2021). In rodents, skill learning has been shown to correlate with intracortical synaptic strengthening within primary motor cortex, giving some credence to this idea (Rioult-Pedotti *et al*. 1998, Xu *et al*. 2009). However, a different interpretation of this data is that early learning of sequence behaviours is primarily facilitated by fast plasticity within motor cortex which is then followed by a slower refinement via the cortico-striatal pathway (Costa *et al*. 2004; Santos *et al*. 2015). If this is true, then the striatum will always been required for motor sequence execution. This an exact role reversal for

the proposed function of these brain regions in songbird vocal learning. Interestingly, although the brain regions do not match, the mechanism seems to be conserved: both systems contain a tutor region needed for learning, and a site where memory is stored. If this true for the mammalian brain, then like area X in the songbird, motor cortex may not be required after learning.

Recent recordings support this notion. Calcium imaging of motor cortical activity across motor skill learning shows that, in early learning, motor cortex is highly correlated with movement. However, as behaviour becomes more refined and less variable, activity in motor cortex decouples from movement suggesting the cortex disengages as an action sequence develops (Kupferschmidt *et al.* 2017; Hwang *et al.* 2019 & 2021). Recent work lesioning motor cortex also provides strong evidence for this. Extensive bilateral lesions prevent learning of a skilled lever press task. However, these same lesions have no effect when executed after learning – when the motor skill has become stereotyped (Kawai et al. 2015). Significantly, lesion to motor cortex does not disrupt task related striatal dynamics (Dwhale *et al.* 2021). Hence, for rodent lever pressing, while DLS is required for both learning and stereotyped execution of motor sequences, motor cortex only seems to be required for learning.

What can explain the opposing circuit layout for songbird singing and rodents lever pressing? A tempting explanation is that this is an evolutionary difference. However, even across different rodent behaviours evidence is not consistent. For example, in some skilled reaching tasks in mice the motor cortex has been shown to always be required for task execution, even after learning (Guo *et al.* 2015). A possible explanation is that the degree to which cortex can disengage is defined by the ability of the basal ganglia to fully control the behaviour – based on the availability of suitable downstream motor controllers. However, in the absence of furthur evidence, for now this is just speculation. In summary, current evidence suggests that for certain behaviours, the motor cortex has the role of tutor, sending inputs to striatum

which are essential during skill learning. In these instances, once a skill has been consolidated, motor cortex is no longer needed and another brain region must be providing glutamatergic drive to the striatum.

2.2.3 The role of thalamo-striatal pathways in motor sequence execution

If motor cortex is not needed for skill execution, then what drives the striatum after learning? Besides motor cortex, the other major input to the sensory motor striatum comes from the thalamus (Hunnicutt *et al.* 2016); though several nuclei in ventral thalamus (VT) send axons to striatum, DLS mainly receives input from the intralaminar thalamus (ILT) (Lacey et al. 2007; Mandelbaum et al. 2019). Single cell tracing in monkeys suggests that more than half of all cells in ILT innervate dorsal striatum exclusively, with a significant proportion of the remaining cortical projecting cells also sending collaterals to striatum (Parent & Parent, 2005). Thalamo-striatal inputs are as numerous as those arising in cortex (Smith *et al.* 2014) and have comparable ability to drive post synaptic spiking – though their synaptic properties do differ (Vandermaelen & Kitai 1980; Smeal *et al.* 2008; Ding *et al.* 2008; Nanda *et al.* 2009; Johansson & Silberberg 2020).

This connectivity makes thalamus an excellent candidate as the main source of glutamatergic drive to striatum during post learning motor sequence execution. However, a role for this thalamo-striatal pathway in controlling action sequences is not well established. Previously, rather than driving activity in striatum, ILT neurons have been suggested as a modulatory input, tuning cortico-striatal connectivity via interactions with cholinergic interneurons (Goldberg & Reynolds 2011). Furthermore, rather than exploring a possible role in motor control, experiments recording from ILT neurons have instead focused on their potential function controlling sensory attention (Matsumoto et al. 2001; Minamimoto *et al.* 2005) and a role linking prefrontal inputs to the striatum for controlling behavioural flexibility (Kato *et al.* 2018).

One system where a role for thalamus in controlling motor sequences is well established is in the songbird vocal learning circuit. Like ILT, an avian thalamic region (UVA) densely innervates the main vocal driver HVC (Ashmore *et al*. 2005). These nuclei show strong syllable modulation, and focal UVA lesions disrupt normal syllable sequencing, and can even abolish syllable singing entirely (Williams & Vicario 1993, Coleman & Vu 2005; Danish *et al*. 2017). While the requirement of UVA for song production is clear, the exact role this nucleus plays is contested. It has been proposed that UVA has a role in organising song composition by modulating syllable order, or even in driving continuous activity in HVC to generate vocal syllables themselves (Alonso et al. 2015). However, thalamic contribution to motor sequence generation in the songbird may also simply be to synchronise activity across hemispheres (Schmitt, 2003). Recent evidence is also conflicting. In one study it was found that for UVA lesioned birds, sleep activity in single HVC hemispheres contains replay of full syllables, and even transitions between syllables (Elmaleh *et al*. 2021). While, these findings suggest UVA does not direct song composition, compelling new evidence has shown the opposite may be true. Inhibiting UVA during song production was shown to truncate singing at syllable boundaries suggesting UVA is important for facilitating syllable-to-syllable transitions. Further, input to HVC from UVA was found to be directed onto specific 'starter cells' which were highly active before syllable onset (Moll *et al*. 2023). Hence, UVA inputs may act to initiate syllables by selecting appropriate HVC starter units. This suggests that the role of UVA is as an organizer, controlling syllable ordering and overall song composition. While compelling, this idea requires further evidence before it can be confidently asserted as true. Nonetheless, though its exact role is yet to be determined, it is clear that the thalamus is a key part of the avian circuit that generates learned vocal sequences.

Could a thalamo-striatal pathway hold a similarly significant position in the equivalent mammalian motor sequence circuit? Recent work in rodents has

begun to explore this possibility. Recordings from striatal projecting neurons in parafasicular nucleus (PF) – a caudal nucleus of ILT – and VT have shown that activity in these nuclei correlate with task initiation and execution in mice completing a sequential lever pressing task (Dias-Hernandez et al. 2018). Further, a causal role for the thalamus in organising the structure of motor behaviours is hinted at from optogenetic manipulations of these pathways. Inhibition to both striatal projecting VT and neurons in the PF delays sequence initiation, while inhibition of the VT pathway mid task extends sequence execution, causing mice to add in extra lever presses during repeated pressing bouts (Dias-Hernandez et al. 2018). This abnormal extension is akin to motif repetition seen during similar striatal manipulations (Geddes et al. 2018), suggesting VT striatal inputs are part of the same functional circuit. Further evidence this is true comes from experiments manipulating a sensory forelimb region of VT. In rats that had previously learned a movement sequence to fit a time interval, muscimol infusions in this region caused animals to overestimate this interval, suggesting a loss of the temporal structure of the previously learned sequence (Hidalgo-Balbuena *et al*. 2019). Even stronger evidence that a thalamo-striatal pathway is responsible for controlling motor sequences comes from presynaptic silencing experiments targeting DLS projecting ILT neurons. Silencing this projection has been shown to lead to rotarod learning deficits in mice (Melief *et al*. 2018), and is catastrophic to a previously learned stereotyped lever pressing sequence in rats (Wolff *et al*. 2022). In the second of these experiments, it was shown that animals were still able to engage with the task but did not improve with continued practice and instead reverted to behaviours seen during early motor exploration. This suggests that the deficit seen is not simply a loss of the ability to initiate the learned sequence, but instead a loss of the circuit that defines both the structure and content of the learned motor sequence.

Though this recent evidence suggests thalamo-striatal projections are responsible for driving established motor sequences, a key question remains

regarding which thalamic regions are responsible. In the silencing experiments of Wolff and colleges (Wolff *et al*. 2022) the broad rostro-caudal extent of ILT was targeted but projections from VT were not silenced. If VT does play a role in controlling motor sequences, it seems likely that it is only a supporting one as VT alone was not sufficient for driving the learned motor sequence. ILT is a diverse set of nuclei which can be anatomically and genetically split into the rostral laminar nuclei (rILT) and the more caudal parafasicular nucleus (PF) (Smith *et al*. 2014; Mandelbaum et al. 2019). This anatomic and genetic delineation also extends to differences in dendritic morphology and firing properties, further suggesting rILT and PF can be considered distinct from each other and hinting that they may have different functional roles (Smith *et al*. 2014). What these roles are can be inferred from further differences between these two regions observed in the synaptic properties of their axonal terminals in the striatum. Both rILT and PF neurons project to MSNs in the striatum but while rILT axons form synapses on dendritic spines, PF boutons predominantly target shafts (Lacey *et al*. 2007). This suggests a modulatory role for PF, while rILT neurons may be better able to drive postsynaptic spiking in striatum. Evidence this is true comes from In vitro optogenetic stimulation in combination with whole cell patch clamp recordings. Ellender *et al*. (2013) showed that rILT activation leads to large amplitude AMPA mediated post synaptic potentials which display short term facilitation. In contrast PF inputs give rise to smaller amplitude NMDA mediated responses with synapses that display short term depression. Since rILT cells tend to spike with high frequency calcium bursts (Lacey *et al*. 2007) these synapses are ideally suited for driving large depolarisations – and therefore spiking – in MSNs. Hence, rILT is an excellent candidate as the main source of glutamatergic drive to DLS needed for execution of learned motor sequences. Since PF inhibition does not impair motor sequence execution (Dias-Hernandez et al. 2018), a strong hypothesis is that silencing rILT projections to DLS alone will impair learned motor sequences.

To summarise, thalamic projections represent a significant source of glutamatergic input to the striatum and increasing evidence suggests thalamic regions could play a vital role in mammalian motor sequence production just as they do in equivalent songbird vocal circuits. Recent work points towards ILT as a key group of nuclei in this regard. These nuclei can be split into rostal and caudal groups with differing neural morphologies, striatal projection patterns, and synaptic properties. I have proposed that the properties of neurons in the rILT make this region particularly interesting as a possible fundamental controller in procedural behaviours. However, it remains unclear whether this is the case as current evidence for this hypothesis is almost entirely circumstantial. As a result, this represents an area that deserves further investigation.

2.2.4 Alternative pathways and extended thalamic circuits for motor sequence learning and execution

Currently the strongest evidence that thalamo-striatal pathways can control learned motor sequence execution comes from the very recent projection silencing experiments discusses in section 2.2.3 (Wolff *et. al.* 2022). However, literature on this subject is sparse and far from conclusive: for example, this population also send collaterals to cortical targets (Parent & Parent, 2005) and recent conflicting evidence suggests there is no contribution of ILT to skilled rotarod running (Kato *et al.* 2021). In the absence of clear direct evidence for the role of ILT, it is therefore helpful to consider two further questions. Firstly, if ILT inputs to striatum are the main drivers of learned motor sequence behaviours, what are the inputs to thalamus that make it suited to this function? Secondly, if not ILT, what other brain regions could be performing this role?

Addressing the second of these questions, besides primary motor cortex and thalamus, DLS receives inputs from other cortical regions (Hunnicutt *et al.* 2016). There is some evidence that somatosensory cortex is involved in motor

learning (Pavildes *et al*. 1993). However, bilateral lesions here have also been shown to have no effect on motor sequence execution (Wolff *et al*. 2022). Another cortical region that has been implicated with a role in controlling learned motor sequences is supplementary or secondary motor cortex (M2). In primates, activity in this region correlates strongly with the movement structure of learned reaching sequences (Shima & Tanji 2000). However, although Inhibition of rodent DLS projecting M2 cells delays sequence onset in a lever press task, this same silencing has been shown to have no effect when applied during sequence execution (Sanchez-Fuentes *et al*. 2021). Furthermore, as discussed previously, extensive motor cortical lesions which include large portions of M2 do not impair motor sequence execution (Kawai et al. 2015; Wolff *et. al*. 2022). Hence, it is doubtful that this region alone makes major contributions to the circuitry that controls learned motor skills.

A final cortical region that should be considered is the prefrontal cortex (PFC). This region sends many inputs to DLS (Hunnicutt *et al*. 2016) and has been implicated in executive functions, controlling the structure and order in which behaviour is composed (Badre & Nee 2018). Recordings from individual cells here support this idea. In a task where primates had to learn multiple sequences of saccades, individual cells showed tuning to the components of the underlying structure of each sequence (Averbeck et al. 2006). Ensembles of these neurons uniquely encoded each learned action sequence, suggesting PFC forms distinct representations for each action sequence. These action state representations even reflect the Bayesian uncertainty of the task; when the rewarded sequence was changed unexpectedly, previous state representations evolved into another as the monkey gathered evidence that a given memorised sequence was required to solve the task. Analysis of population dynamics have shown that overlapping action state representations slowly orthogonalize in both PFC and striatum (Marton et al. 2020), suggesting that PFC - DLS subnetworks may be essential for forming, and separating, hierarchical motor representations akin to those discussed

in chapter 2.2. While the evidence that PFC plays a role in higher order hierarchical structuring of motor behaviours is compelling, there is little to suggest that this region is also responsible for lower-level control of the content of motor skills. This is an important demarcation. While the DLS has been shown to have a role in higher order behavioural ordering (Jin et al. 2014; Geddes *et al*. 2018; Markowitz et al. 2018). Neural activity here also encodes features of movement direction (Hikosaka & Wurtz 1983), amplitude (Rueda-Orozco & Robbe 2015) and even effector specific kinematic representations (Dwhale *et al*. 2021). Inputs from PFC alone therefore cannot account for this.

One way in which these different levels of control could be unified in striatum, however, is through the ILT which also receives dense projections from PFC (Akert & Hartmann 1980; Saalmann, 2014). This bring back the first of the two outstanding questions in this section: if ILT inputs to striatum are the main drivers of learned motor sequence behaviours, what are the inputs to thalamus that make it suited to this function? For the reasons detailed above, receiving inputs from PFC makes ILT extremely suitable as a key node in the motor sequence execution circuit. ILT also receives dense ascending innervation from midbrain and brainstem nuclei. Prominently from the superior colliculus though also from ascending pontine and cerebellar inputs (Yamasaki *et al*. 1986; Cornwall & Philipson 1988; Krout et al. 2002; Bostan & Strick 2018; Mandelbaum *et al* 2019). These motor nuclei are involved in controlling various specific features of motor output (Takakusaki *et al* 2016; Wheatcroft *et al*. 2022). Hence, ILT is a recipient of inputs that carry information about ongoing motor control at a kinematic level. This puts ILT in a unique position; these thalamic nuclei are ideally suited for sending both high level structural information from PFC and low-level kinematic information to striatum for control of both the structure and content of motor sequences.

In fact, that motor output, or efference information, is routed back up to ILT gives rise to a compelling further hypothesis – analogous to similar mechanisms that have been proposed for the function of thalamic feedback in birdsong vocal learning (Alonso *et al*. 2015; Moll *et al*. 2023). Action sequences can be considered as chains of elemental action, such that each element is predicted and potentially even triggered by the proceeding action in the chain. Routing motor information back into the striatum via thalamus could facilitate this by providing the DLS with the information required to form action to action associations. Midbrain and brainstem regions which send inputs to ILT are themselves anatomically downstream of the striatum. ILT is therefore ideal situated to be a key feedback node in this loop architecture and hence, a key node in a circuit which could potentially underly action-to-action chaining for motor sequence generation.

2.2.5 Summary

In this sub-section I have described the broad circuitry that is believed to underpin mammalian motor sequence learning and execution. I have described the DLS as a key nexus in this pathway. Evidence suggests this region is a site where behaviour specific state action mappings are formed for control of motor behaviours at multiple levels. I have discussed that motor sequence learning is a process most likely defined by motor cortical interactions with this region and I have discussed the evidence that thalamo-striatal pathways are a significant feature of procedural circuitry; most likely underlying the generation of already learned motor sequences. I have proposed that the rostral portion of ILT is a key node in this respect. Finally, I have argued that this region may function as an integrative component in a feedback loop; facilitating action-to-action chaining by routing motor efference information from midbrain and brainstem motor controllers to the DLS.

Now that I have introduced the key circuits which are believed to underpin awake – online – motor sequence learning and execution, in the next section I will focus on offline processes (during rest) that are believed to support the function of these circuits.

2.3 Offline consolidation of motor skill memory

Memory for motor sequences is known as procedural memory. Like any memory, procedural skills are primarily learned and strengthened by repetition (Ericsson *et al*. 1993). Repetition of neural activity is thought to engage plasticity mechanisms, engraving activity patterns into neural circuitry, and creating an engram; a representation of that movement sequence in neural form (Josselyn & Tonegawa 2020). While awake practise, physical repetition of a movement sequences, can be considered online learning, procedural memory formation is also thought to be supported by periods of offline consolidation. These are memory improvements that occur in the absence of physical repetition, most significantly thought to occur during sleep (Diekelmann & Born 2010; Rasch & Born 2013; Schmid *et al*. 2020). Sleep is an extremely prominent brain state. On average humans sleep for 7 – 9 hours a night constituting around a third of our entire lifetime (Buysee *et al*. 1989; Krause *et al*. 2017). Still, compared to other mammals, humans are short sleepers. Mice and rats for example have equal ratio of sleep to waking and in some animals, for example certain bat species, sleep periods represent over 80% of their total lifespan (Campbell & Tobler 1984). Even during awake periods, animals spend large proportions of that time in quiescent restful states where neural processes are believed to resemble those observed during sleep (Foster & Wilson 2006; Mendick *et al*. 2011; Fox *et al*. 2013). In this section I will explore how processes during these extensive offline states contribute to the formation and preservation of motor skills.

2.3.1 The role of sleep for procedural memory consolidation

A role for sleep in procedural memory formation is well established (Robertson *et al*. 2004). In humans, performance in finger tapping sequences is improved by periods of post learning sleep (Fischer et al. 2002; Walker et al. 2002; Korman et al. 2007; Nishida & Walker 2007). These improvements are far greater than those seen for the equivalent time spent awake and multiple nights of sleep have been shown to provide larger benefits for skill acquisition than a single night of sleep (Walker et al. 2003-A). A link between sleep and procedural memory has been less extensively studied in animals but evidence is favourable. For example, sleep deprivation has been shown to impair rotarod learning in mice (Yang et al. 2014). Though short periods of post task sleep don't seem to boost performance on this task, a boosting effect has been found for more complex motor tasks in mice (Nagai *et al*. 2017). Likewise, post learning sleep has been shown to improve grasping accuracy in a reaching task in rats (Ramanathan *et al*. 2015). Whether sleep facilitates an actual boost in performance, or simply stabilises motor memory, making it more resistant to interference, is debated (Maltry *et al*. 2021; Walker et al. 2003-A). Most likely the differences here are nuanced and related to the specifics of the motor skill and task being learned. However, what is clear is that offline – practise independent – processes do support motor memory in some form and these processes seem to occur most eminently during sleep.

Sleep is a complex brain state made up of several distinct stages – marked by characteristic oscillatory patterns in local field potential (LFP) (Duran *et al* 2018). For procedural memory it is not entirely clear whether a single sleep state is responsible for consolidation. For instance, there is some evidence for a role for rapid-eye-movement (REM) sleep in supporting motor learning (Smith 2001; Smith *et al*. 2004; Rasch & Born 2013). However, other studies have indicated that, rather than impairing motor learning, pharmacologically

blocking REM sleep improves motor skill acquisition (Rasch *et al*. 2009). If a single sleep state is important for procedural memory consolidation, evidence generally points towards non-rapid-eye-movement (NREM) sleep, specifically stage 2 (NREM2). In humans, overnight improvements in a motor sequence correlate with the amount of NREM2 observed during sleep (Kuriyama *et al*. 2004) and NREM2 deprivation specifically impairs motor learning compared to deprivation of other sleep stages (Smith & MacNeill 1994; Forest & Godbout 2000). Further evidence comes from sensory cuing experiments. Sleep based procedural memory consolidation can be boosted by cuing that memory with presentation of a sound or odour which has previously been associated with it. Cuing memory in this way has been shown to be most effective when targeted to NREM2 sleep stages (Laventure et al. 2016; Cousins et al. 2014). This association between procedural memory and NREM2 seems to be facilitated by a particular oscillatory event, sleep spindles, which occur most prominently during this sleep stage. Sleep spindles are rapid synchronous bursts of activity between 8 and 15Hz. These oscillations are found throughout NREM sleep but are particularly numerous during NREM2 (Laventure et al. 2016). Procedural memory enhancements correlate with the number and density of spindles which are observed during NREM2 (Fogel et al. 2006, Nishida & Walker 2007; Rasch *et al*. 2009; Barakat *et al*. 2012). Further, in sensory cuing experiments, presentation of the cue has been shown to cause a temporally locked increase in observed spindles. Evoking spindles in this way correlates with observed post sleep procedural memory enhancements (Cousins *et al*. 2014; Laventure *et al*. 2016). Hence, the mechanisms which underly procedural memory consolidation most likely occur during NREM2 sleep, and at very least, correlate strongly with spindle events.

A final noteworthy feature of offline memory consolidation is that it is not simply a passive process whereby the brain solidifies all motor experience. Instead, procedural memory consolidation appears to be a highly

choregraphed process, organised around the structure and behavioural significance of memory. In sequential finger tapping tasks, learning a different sequence before sleep blocks sleep induced performance consolidation for the original sequence suggesting overnight memory consolidation mechanisms are active, and orchestrated such that they can be memory selective (Walker et al. 2003-B). Even within memory, there is some evidence that specific parts of a sequence are targeted by consolidation mechanisms. If different tones are paired with different parts of the sequence, memory for specific subsequence movements can be preferentially boosted by playing the associated tones during sleep (Schonauer *et al*. 2014). Furthermore, analysis of post sleep finger tapping performance suggests sleep related gains are specific and targeted to difficult, previously poorly performed, parts of the sequence (Kuriyama *et al*. 2004). One explanation is that poorly performed movements simply have more room for improvement during generalised consolidation. However, this raises the further possibility that offline mechanisms are guided and targeted in a choreographed way by an internal recognition of performance. Similarly, it has been proposed that sleep promotes consolidation for memories that are considered relevant for the future (Diekelmann & Born 2010). Humans show preferential sleep related consolidation of one finger tapping sequences over another if they are told that sequence will be highly rewarded the next day (Fischer & Born 2009). Hence, motor memory consolidation may be targeted to specific relevant memories or even elements of those memories. How this arbitration process works, and how memories are split such that they can be consolidated as single elemental fragments, however, remains unclear.

To summarise, sleep has a crucial role for the consolidation of procedural memories. Sleep dependent processes may even boost memory for motor sequences overnight. Though it may not be the only stage involved, evidence clearly points towards NREM stage 2 sleep as most crucial sleep stage for this process. Consolidation correlates strongly with sleep spindle oscillations

during NREM2 suggesting these LFP events may define memory consolidation epochs. Finally, evidence suggests that sleep dependent memory strengthening is not simply a passive process, but instead seems to be directed and targeted to specific memories or even features of memory. Targeting appears to relate to the behavioural significance of memories suggesting consolidation is high choreographed, though the mechanisms that organise this process are not known.

2.3.2 Neural reactivation as a mechanism for memory consolidation

In the previous section I have described offline procedural memory consolidation as a phenomenon, including specific observed aspects of this process. In the next sections I will explore what is known and what is unknown about the mechanistic neural basis of these phenomena. One area of research that it is useful to draw inspiration from is the study of Episodic memory: memory of places and events. Like procedural memory, episodic memory is also supported by periods of offline consolidation (Inostroza & Born 2013) but has been studied to a far greater extent in recent history. Episodic memory is thought to be dependent on the hippocampus. Bilateral lesions here in humans impair autobiographical memory formation (Scoville & Milner 1957) and in rodents lead to various spatial memory deficits (Morris *et al*. 1982; Olafsdottir *et al*. 2018). Memory consolidation in the hippocampus is believed to be based on reactivations of previously active ensembles. These reactivations, or replay events, were identified by recording place cell assemblies in the hippocampus during sleep (Wilson & McNaughton 1994, Skaggs & McNaughton 1996; Lee & Wilson 2002; O'Neill et al. 2010). In replay, these cell assemblies fire together in a specific temporal order, mimicking activity that previously occurred when an animal moved through a specific spatial trajectory (Lee & Wilson 2002). These reactivation events have been found to occur most frequently during short high frequency (140-250Hz) LFP oscillations termed sharp wave ripples

(SWR) (Wilson & McNaughton 1994; Olafsdottir *et al*. 2018). Replay observed during ripples is often compressed in time, such that activity unfolds up to 20 times faster than the same activity observed during awake behaviour. This compression is thought to be essential for promoting Hebbian plasticity by increasing the degree of temporal correlation between cell firing (Magee & Johnston 1997). That this occurs during sleep is thought to be a way of avoiding interference from external sources of information. This idea is somewhat contradicted by the fact that replay has also been found to occur when animals are awake (Foster & Wilson 2006). However, these awake replay events only occur during periods of quiescence, when there is little movement and such interferences are minimised. More direct evidence for the role of replay in memory consolidation comes from correlation analysis; the number of task specific replay events that occur has been shown to robustly predict how well rats perform on a spatial memory task the next day (Dupret et al. 2010). Causal evidence for this come from closed loop work. Detecting SWR and then electrically interrupting hippocampal activity – and therefore likely interrupting replay – has been shown to impair memory formation (Girardeau et al. 2009; Ego-Stengel & Wilson 2010). Interrupting SWR also disrupts the stability of individual cell spatial coding (Roux et al. 2017), further suggesting activity during these LFP events is essential for consolidating neuronal mappings. Replay also correlates with specific features of memory consolidation. For example, more spatial replay events are observed for rewarded than unrewarded positions (Singer & Frank 2009). However, other features of replay are harder to link directly to features of memory consolidation such as the tendency for replayed neural sequences to appear in reverse order (Foster & Wilson 2006; Diba & Buzsaki 2007). These inversions are not neural patterns seen during awake behaviour so their role in memory consolidation is not clear. Additionally, as well as extended replay events representing continuous spatial trajectories, fragmented, and stationary replay events have been observed (Denovellis *et al*. 2021). Fragmented replay is far less common than spatially coherent reactivations.

However, extended trajectories are also relatively rare such that most replayed activity is stationary – representing a single spatial cluster. While these replay events could easily play a role in memory consolidation, it remains unclear what this role is. Nevertheless, for episodic memory, there is a great deal of evidence that replay is the mechanism by which memories are consolidated. These reactivations may act as repeated simulations of awake activity. Hence, they could allow the brain to form an engram for previous episodic experience via a generative process of engraving activity patterns into hippocampal circuitry.

In summary, for episodic memory, consolidation is attributed to neural replay; reactivation of previously active – behaviourally relevant – ensembles. Episodic consolidation is dependent on the hippocampus, and neural replay tends to occur during SWR oscillations in this region. Replay is often compressed in time which may promote synaptic plasticity, engraving replayed activity into memory engrams. Finally, while replay often propagates in a similar way to awake activity in hippocampus, backwards and fragmented replay have also been identified. How these different kinds of reactivation are organised and how they contribute to memory consolidation is not known.

2.3.3 Neural mechanisms for procedural memory consolidation

Could hippocampal replay also underlie procedural memory consolidation? Some evidence suggests this could be the case. Humans with reduced hippocampal volumes due to brain injury have been found to be worse at learning some motor sequence tasks (Long *et al*. 2018). Further, fMRI imaging studies have shown that hippocampus is active during the early stages of motor skill learning (Schendan *et al*. 2003) and suggest it may be necessary for short term motor skill improvements observed after awake rest (Jacobacci *et al*. 2020). However, if hippocampal replay alone underlies procedural memory consolidation, then this region must have a direct role in motor sequence production. Certainly, the hippocampus is not considered a motor

structure – and was not part of the circuitry discussed in section 2.2. However, in mice, sleep dependent motor skill learning has been shown to lead to increased immediate early expression in hippocampus (Nagai 2016). One suggested mechanism is that hippocampal circuits may define the order of sequential motor output via competitive queuing (Kornysheva *et al*. 2019). However, it seems unlikely that this mechanism can explain sleep dependent procedural consolidation wherein fine motor control becomes more precise. Indeed, hippocampal replay has been more often associated with learning the structure and hierarchy of abstract sequences rather than kinematic features of movement (Lui *et al*. 2019). In fact, it may be more appropriate to consider that hippocampal and procedural circuits function within completely distinct and separate memory systems (White & McDonald, 2002). Strong evidence for this dissociation comes from lesion studies. For example, hippocampal lesion patients with strong declarative memory impairments have been shown to have no procedural memory deficits (Reber & Squire 1998; Hopkins *et al* 2004). Most notably patient HM, who had profound retrograde amnesia, was able to make improvements in a mirror drawing task which tests for procedural learning (Squire, 2009). Moreover, besides being distinct, procedural, and declarative memory may even be somewhat antagonistic. Monkeys with hippocampal lesions have been shown to learn a sequential motor task faster than those with intact hippocampi (Douglas & Pribram 1969). In rodents, animals with bilateral hippocampal ablation are better at motor sequence tasks (Eckart *et al*. 2012; Will *et al*. 2013; Busse & Schwarting 2016[A]; Busse & Schwarting 2016[B]; Schwarting & Busse 2017) and can learn sequential lever pressing tasks faster than intact controls (Jackson & Strong 1969). If lesioned animals can learn to produce stereotyped movements, then hippocampus cannot be a key part of the circuitry that directly forms procedural memories. Since hippocampal lesions do have a slightly detrimental effect on performance when animals are asked to learn very long, highly repetitive, sequences (Christie & Dalrymple-Alford 2004), the hippocampus may only be needed for higher

order, hierarchical, organisation - which is strained during longer repetitive sequences. This dissociation from fine motor control means hippocampal replay cannot directly underly sleep dependent procedural memory improvements.

While strong evidence suggests that replay in the hippocampus is not required for procedural consolidation, could a similar mechanism, replay of experience in motor circuits, underpin this phenomenon? Certainly In multiple animal species, and in humans, replay has been found across cortical areas including primary motor cortex (Ramanathan *et al*. 2015; Rubin *et al*. 2022). Like hippocampal replay these offline reactivations are faster than the same activity observed during awake behaviour and correlate with improvements in motor sequence production (Eichenlaub *et al*. 2020). In line with previously discussed studies showing NREM sleep correlates with motor consolidation (Kuriyama *et al*. 2004; Laventure et al. 2016; Cousins et al. 2014), motor cortical replay occurs most often during NREM sleep and is enhanced during slow wave oscillations – low frequency LFP oscillations defined by synchronous active 'up' and silent 'down' states across neural populations (Gulati *et al*. 2017; Rubin *et al*. 2022). NREM slow wave sleep has been shown to promote the formation of new spines in motor cortex giving credence to the notion that this replay represents a circuit consolidation mechanism (Yang *et al*. 2014). This also provides further evidence that NREM sleep is the key sleep stage for procedural memory consolidation, though it should be noted that REM sleep has been demonstrated to be required for pruning and stabilising recently formed spines (Li *et al*. 2017). Hence, a synergistic mechanism involving both sleep states is likely. Motor cortical reactivations are a compelling mechanism for procedural consolidation, however, as discussed in sections 2.2.1 and 2.2.2 motor sequence learning is believed to be mediated by plasticity in cortico-striatal synapses. If motor cortical reactivations underpin motor sequence consolidation, then this activity must propagate to striatal networks. Evidence this is true comes from

work simultaneously recording EEG and fMRI in humans learning a motor skill. During NREM sleep, increased activity in motor cortex and striatum correlates with improvements in motor memory (Fogel *et al*. 2017). Interestingly offline activity in striatum was associated with spindle events in cortex. As discussed previously, spindle events correlate very strongly with procedural consolidation (Cousins et al. 2014; Laventure et al. 2016) and closed loop transcranial stimulation of motor cortex during sleep spindles has been shown to boost motor memory in humans (Lustenberger *et al*. 2016). Hence, spindles may represent offline coupling between motor cortex and striatum and may therefore be a marker for procedural consolidation. Further, evidence that this is true comes from recordings in rats learning a skilled reaching task (Lemke *et al*. 2021). Blocking offline activity in the striatum impairs learning, while the amount of functional connectivity between cortex and striatum during NREM correlates strongly with motor improvements. Importantly, increased connectivity between cortex and striatum was shown to correlate with cortico-striatal spindle events.

To summarise, unlike other forms of memory, evidence strongly suggests that the hippocampus is not involved in procedural consolidation. Instead, a strong hypothesis is that sequential replay of previous activity in other, motor related circuits, including the striatum, may facilitate this. Procedural reactivations correlate with and may even be triggered by spindle oscillations in a manor analogous to the way in which SWR interact with hippocampal replay. In the next section I will explore the significance of these LFP oscillations for procedural consolidation.

2.3.4 LFP biomarker interactions in procedural memory consolidation

Since sleep spindles correlate with procedural reactivation events, it seems likely that they may contribute in functional way to procedural consolidation. Spindles are thought to be triggered by cortical activity but generated in the thalamus by oscillatory activity with the thalamic reticular nucleus (TRN).

These oscillations are then spread across the neocortex from thalamus (Klinzing *et al*. 2019). Spindles occur locally in disparate cortical regions and have been theorised as a mechanism for increasing cross connectivity between cortical circuits (Fang *et al*. 2020; Dickey *et al*. 2021). The high frequency nature of these oscillations increases firing rate in pyramidal neurons, temporally locking neural activity together and hence, promoting plasticity (Dickey *et al*. 2021). These features make spindles an ideal substrate for binding disparate cortical representations into unified striatal plasticity. Additionally, it has been shown that specific fractional components of a procedural memory can be independently consolidated (Schonauer *et al*. 2014). Since spindles are local, an intriguing idea is that these events could segment replayed activity in motor cortex to partition procedural memories into constituent fragments for consolidation in cortico-striatal circuits. Whether this is true, however, is not yet known. Spindles may also interact with other LFP events to define consolidation of procedural memories. Reactivation events are highly correlated with coincident nesting of spindles with peaks (synchronously active 'up' states) in slow waves oscillations (Lemke *et al*. 2021; Schreiner *et al*. 2021). Coupling between these two LFP phenomena has been shown to have a role in strengthening memories while similar nesting with slow wave down states - or 'delta oscillations' – leads to memory weakening (Kim *et al*. 2019). How these mechanisms function at a circuit level, however, is not known and open questions remain regarding how the content of neural activity within nested spindle events relates to features of memory consolidation.

If local procedural reactivations do underly procedural memory consolidation and are defined by local LFP phenomena, one final open question relates to how these LFP events are orchestrated. Since, spindles are also believed to underly consolidation of other forms of memory via coupling with cortical slow waves and hippocampal SWR, it is unclear whether procedural memory is mechanistically distinct from other forms of memory, or simply a part of a

global consolidation system. For autobiographical memory, hippocampal damage in humans leads to impaired formation of new memories but spares previous memories (Squire, 2009). Episodic consolidation is therefore thought to be mediated by a transmission of new memories from the hippocampus to the cortex for long term storage (Buzsaki 1989, Siapas & Wilson 1998; Klinzing *et al*. 2019; Spens & Burgess 2023). LFP coupling between these regions is thought to mediate this transfer. Indeed, cortical spindles and hippocampal SWR events are often temporally overlapping, and experiments phase locking these events by optogenetically stimulating TRN to generate spindles that are coincident to SWR has been shown to lead to memory improvements (Latchoumane *et al*. 2017). Further, artificially nesting hippocampal SWR with cortical slow waves has a similar boosting effect on memory for novel objects in rats (Maingret *et al*. 2016). One suggestion is that hippocampal SWR trigger cortical LFP events – and hence replay – across cortex to synchronise and bind disparate cortical regions. If this is true, the hippocampus may also have a role in procedural memory formation, by triggering replay events in motor cortex. Certainly, this role cannot always be needed for procedural memory, given the evidence discussed in the previous section. However, if we entertain that there could be an indirect role for the hippocampus in triggering procedural memory consolidation this may explain why hippocampal replay has been associated with procedural memory (Schendan *et al*. 2003; Schapiro *et al*. 2019 ; Jacobacci *et al*. 2020). Recent correlational evidence for this indirect role for hippocampus has been shown for rats learning a reach to grasp task. In these animals, during early learning sleep, motor cortical slow waves and hippocampal SWR were found to be correlated with each other while slow wave correlations between cortex regions were much weaker. As reaches became stereotyped, these correlations reversed, suggesting initial hippocampal-cortical connectivity leads to greater cortico-cortical connectivity. If learned reaches are then disrupted, meaning animals must adapt their motor plan, these correlations again revert to that seen in early learning suggesting the hippocampus could,

in some circumstances have a role in consolidating early motor exploration into stereotyped procedural memories (Kim *et al*. 2022). However, since these results do not provide any causal evidence it is very hard to reconcile this hypothesis considering the lesion work mentioned previously (Jackson & Strong 1969; Douglas & Pribram 1969; Reber & Squire 1998; Hopkins *et al* 2004; Squire, 2009; Eckart *et al*. 2011; Will *et al*. 2013; Busse & Schwarting 2016[A]; Busse & Schwarting 2016[B]; Schwarting & Busse 2017). If the hippocampus is important for triggering procedural reactivations, then why does hippocampal damage not impair procedural memory formation? Furthermore, the notion that hippocampal SWR trigger LFP events in cortex is itself disputed and the opposite may even be true. While there are clear bidirectional influences, recent work suggests that cortical slow waves and spindles are the best predictors of nested SWR in hippocampus (Dickey *et al*. 2021; Staresina et al 2023). These conflicting lines of evidence underpin the fact that little is understood about how offline consolidation is organised by memory systems. To what extent episodic and procedural memory systems interact and whether they are truly independent remains entirely unclear. Furthermore, whatever role (or lack of role) the hippocampus has in procedural memory formation, open question remains over the mechanisms by which spindles and slow wave events are organised in relation to procedural consolidation.

In this section I have described a role for offline processing in procedural memory consolidation. The mechanism that underlies this phenomenon is likely to be akin to hippocampal replay of neural activity, though most likely occurs in cortico-striatal circuits and may be entirely distinct from hippocampal consolidation. Procedural reactivations are probably triggered by spindle oscillations though the structure and content of these events at the circuit level in striatum is not known. Further, sleep spindles may interact with other oscillatory events including slow wave oscillations. How these interactions define procedural consolidation and alter the content of

procedural reactivations is also not known. Finally, to what extent these LFP phenomena define procedural consolidation, and how they are themselves orchestrated in relation to other memory consolidation systems remains an open question.

2.4 Summary and research aims

In this chapter I have first outlined the broad circuitry that is believed to underpin mammalian motor sequence learning and execution with an emphasis on the role of circuits that interact with dorsolateral striatum. Though the exact role this region has in controlling motor behaviours in debated, evidence is clear that the DLS is a key node in the mammalian motor sequence circuit. The motor cortex appears to be an essential partner to this region, proving glutamatergic drive to striatum which leads to the formation of state action associations during learning. After learning, however, a role for the cortex is less clear. Mounting evidence suggests that in certain behavioural contexts, thalamic regions may take over as the main glutamatergic drivers of striatal activity. However, multiple questions remain. Firstly, it is not clear which specific thalamic regions are involved in this process. Secondly, if thalamo-strital projections are essential for motor sequence production, it is not known what features of motor control this pathway contributes to.

While it is not known which specific thalamic regions essential, recent compelling evidence suggests that the ILT could be a key group of nuclei in this regard. I have proposed that the projection patterns and synaptic features of neurons in the rostral part of ILT makes these cells well suited for this role. These neurons have access to both higher order structural information from PFC and kinematic motor efference information from midbrain and brainstem motor controllers. Hence, I have proposed this thalamic region may function as an integrative component in feedback loop; facilitating action-to-action chaining by routing this information to the DLS.

With these proposals as guiding hypotheses, this thesis has the following aims related to these ideas.

- Firstly, I aim to find evidence showing whether projections to the striatum from rILT are required for normal motor sequence learning and execution.
- Secondly, I aim to understand how thalamo-striatal projections contribute to motor control during both learning and execution of motor sequence behaviours.

I also explored the role of offline consolidation in motor sequence learning and execution. Like other forms of memory, sleep dependent processes are essential for normal procedural memory formation. The mechanisms that underlie this phenomenon may occur in the DLS and are likely to be similar to those found for episodic memory consolidation – hippocampal replay of episodic neural activity. However, whether sequential replay of procedural activity occurs is not known. Evidence also suggests that procedural consolidation may be entirely independent of the hippocampus but the extent to which this is true remains an unresolved matter.
In relation to these unknowns, this thesis has the following aims:

- Firstly, I aim to determine whether neural replay of procedural activity exists. I will investigate the content of consolidation related activity in striatum and determine to what extent procedural consolidation mechanisms reflect those observed for other memory systems.
- Secondly, I aim to explore whether there is a role for the hippocampus in procedural memory consolidation or whether it is consolidated by a process which is entirely distinct from other forms of memory.

CHAPTER 3

3. MATERIALS AND METHODS

3.1 Animals

Adult male and female Mice (8-50 weeks) from the following mouse lines were used: C57BL/6J wild-type, GRM2-Cre (MGI ID: 5311758) and Ai75D (tdTomato, JAX Stock No: 025106). Mice were housed in HVC Cages with free access to chow and water on a 12:12 h inverted light:dark cycle and tested during the dark phase. Mice used in sleep recording experiments were housed on a non-inverted light cycle and tested during the light phase (normal daylight hours). For behavioural experiments, mice were water deprived. Animals had access to water during each training session, and otherwise 1mL of water per mouse was administered. Water was supplemented as needed if the weight of the mouse was below 85%. All experiments were performed in accordance with the UK Home Office regulations Animal (Scientific Procedures) Act 1986 and the Animal Welfare and Ethical Review Body (AWERB). Animals in test and control groups were randomly selected.

3.2 Behavioural procedures

3.2.1 Training arenas

Water restricted mice were trained in custom built behavioural arenas measuring approximately 16cm x 19cm x 24cm (width, length, height). Box walls were made of 0.5cm thick translucent white or transparent red acrylic and had 8 poke ports mounted on the front wall. Mouse ports (sanworks, ID 1010) protruded 2cm from the wall into the area and were arranged in a 4 x 2 grid such that neighbouring ports (vertical and horizontal) were 3cm apart (centroid to centroid). Ports contained side mounted infrared LED and sensor to detect poke events and a back mounted visible light LED to illuminate the port. Each port also contained a waterspout for reward delivery. Poke events were registered by Sanworks port PCBs (ID: 1004) connected to a Bpod (ID: 1027) programmed with a custom behavioural protocol (MATLAB). Water delivery was triggered by Bpod via a connected Miniature Solenoid (Lee

Company). Sounds were played at port entry via a speaker (DigiKey part number: HPD-40N16PET00-32-ND) and amplifiers (DigiKey part number: 668-1621-ND).

3.2.2 Sequence task

Mice were rewarded for completing the full 5 step poke sequence. No reward was given if animals missed a step in the sequence, but animals were not punished for adding extra steps into the sequence. Single trial events were defined as all poking activity that led to reward delivery at the final port, hence, within a single trial animals could make multiple attempts at completing the sequence or add additional elements to the sequence and still eventually receive reward. The task was self-paced though if an animal initiated a trial but did not register a poke into any port for 30s this trial timed out, was left unrewarded, and a new trial was cued. To reduce the number of false pokes – where animals attempted to but didn't break a port IR beam – it was necessary to signal to the animals if a given poke had been registered: across all levels, when animals entered a port (breaking the IR beam) a short duration buzzing sound was played.

3.2.3 Automated training

Poke sequences were shaped by an automatic protocol with 50 training levels of set difficulty. Mice started from the lowest level (level 1) and progressed up to the final task (level 50). During training, performance was assessed every 10 trials (see "training performance") and this metric determined whether mice progressed up a level (performance > 90%), regressed down a level (performance < 20%), or remained at the same training level (performance > 20% & < 90%). In Early task levels (1 – 12) mice were rewarded for performing each step of the 5-step sequence. With progression to higher levels, reward steadily decreased and then switched off port by port until the mice were only rewarded for reaching the final port (levels 12 - 50). In early task levels (1- 12) each step in the sequence was also visually guided by successively

illuminating port LEDs which were switched off port by port once the animal had successfully poked. After this (levels 12 – 49), across successive levels LEDs brightness was steadily dimmed port by port and eventually turned off permanently, except for the initial port in the sequence which remained illuminated at the start of each trial to signal a new trial was available. At the final stage of the task (level 50) only the first port in the sequence was illuminated. As with previous levels once the mouse poked and began the trial this LED was switched off and for the duration of these trials no other port LEDs were switched on.

3.2.4 Behavioural testing

Even after reaching the final task (level 50), during training animals were able to drop down to lower levels if they performed badly. However, in circumstances where it was necessary to test animal performance at the full task, mice were held at level 50 for the duration of the session. For AP5 infusion experiments, to increase sensitivity to performance changes during test sessions, training performance was assessed (and training level updated) every 4 trials.

3.3 Surgical procedures

Mice were anesthetized with Isoflurane (0.5–2.5% in oxygen, 1 l/min) - also used to maintain anaesthesia. Carpofen (5 mg/kg) was administered subcutaneously before the procedure. Craniotomies were made using a 1-mm dental drill (Meisinger, HP 310 104 001 001 004). Injections were delivered using pulled glass pipettes (Drummond, 3.5”) on a stereotaxic frame (Leica, Angle TwoTM).

3.3.1 Viral injections

For striatal lesions, initial lesions were excitotoxic via injection of NMDA (2mg/100mL), though for most animals shown lesion was achieved by

injecting a 1:1 mix of AAV2/1-hSyn-Cre (1014 vg/ml) and AAV2/5-EF1a-DIO-taCasp3-T2A-TEVp (1014 vg/ml) as this proved more successful. The mix was diluted 5 times in saline buffer prior to injection. For control animals, saline or GFP virus AAV2/5-CAG-EGFP (1043 vg/ml) was injected instead. In each hemisphere 4 injections (~80nl each) at 3 different depths were made to distribute the viruses as evenly as possible and to provide enough coverage, injections were targeted to the DLS or DMS dependent on experimental group. For retrograde JAWS mediated optogenetic manipulations, 4 injections (~80nl each) of rgAAV8-EF1a-loxFAS-JAWS-KCG-EGFP-ER2 were injected into the DLS at 3 different depths to distribute the virus evenly and provide coverage. For injections into the DLS, insertions were made at coordinates AP: 0.2 to 0.8mm ML: 2.5 to 2.7mm DV: -3.0mm to -3.7mm (where a range is given, injections were given at regular spacing between these values). For injections into the DMS, insertions were made at coordinates AP: 0.2 to 0.8mm ML: 1.8mm DV: -3.0mm to -3.7mm (where a range is given, injections were given at regular spacing between these values).

For hippocampus lesion experiments this same cre-caspase mixture as used in the striatal lesion experiments was injected. Control animals underwent the same surgical procedures except saline was injected instead. Injections were made bilaterally at 13 locations per hemisphere (see table 3.3.1a). After surgery, animals were given at least 3 weeks of recovery before training was started.

For GRM2 lesions, Cre-dependent caspase virus AAV2/5-EF1a-DIO-taCasp3-T2A-TEVp (1014 vg/ml) was injected to achieve lesion. Control animals underwent the same surgical procedures but a GFP virus AAV2/5-CAG-DIO-EGFP (3x1012 vg/ml), was injected instead. Injections were made bilaterally at 8 sites per hemisphere (see table 3.3.1b).

AP (mm)	ML (mm)	DV (mm)
-1.22	+/- 0.75	- 1.8
-1.58	+/- 1.54	- 1.87
	+/- 0.81	- 2.06
	+/- 0.81	- 1.66
- 2.18	+/- 0.9	- 1.83
	+/- 1.8	- 2.08
		- 1.58
	+/- 2.54	- 2.59
		- 2.09
- 2.8	+/- 1.57	- 2.03
		- 1.53
	+/- 2.43	- 2.4
		- 1.9
	+/- 3.17	- 4.04
		- 3.54
		- 3.04
		- 2.54
	+/- 2.31	- 4.66
- 3.64	3.02	- 4.33
		- 3.83
		- 3.33
		- 2.83
	2.53	- 3.90
		- 3.40
		- 2.90
		- 2.40

Table 3.3.1a: Positions of injections into hippocampus

AP (mm)	ML (mm)	DV (mm)
- 1.4	+/- 0.3	- 3.7
	+/- 0.7	- 3.4
1.6	+/- 0.2	- 3.73
	+/- 0.4	- 3.63
	+/- 0.7	- 3.55
1.85	+/- 0.2	- 3.88
	+/- 0.4	- 3.8
	+/- 0.7	- 3.35
		- 3.55

Table 3.3.1b: Positions of injections into intralaminar thalamus

3.3.2 Cannula and fibreoptic implantation

For Cannulation experiments, 5mm 26-Gauge cannulas (P1 technologies, cat number: C315GS-5/SPC) were implanted at coordinates anterior posterior (AP) 0.5mm, medial lateral (ML) 2.9mm and dorsal ventral (DV) 2.0mm from pial surface, with a 10 degrees tilt. For optogenetic experiments, we implanted flat optical fibers of 200µm diameter (Newdoon: FOC-C-200-1.25-0.37-7) at AP -1.6mm, ML -1.9mm, DV 2.8mm (20 degrees tilt). Implants were affixed using light-cured dental cement (3m Espe Relyx U200) and dental cement (Super-Bond C&B Bulk-mix, Sun Medical) and the wound sutured (6-0, Vicryl Rapide).

3.3.3 Neuropixel probe implantation

Prior to training animals underwent an initial surgery where the skull was exposed, coated with a thin layer of dental cement, and marked for later skull levelling. A craniotomy was drilled over an arbitrarily chosen region of posterior cortex and a ground pin was implanted. To replicate the weight and size of the eventual implant, mice were trained on the task with a size and

weight matched dummy implant, this was fixed during initial surgery to the cement layer using silicon (Kwik-sil: World Precision Instruments). For probe implantation, the dummy implant was removed, and the skull was levelled using previously noted skull markings. A craniotomy and durotomy was made at coordinates AP 0.8mm, ML = 2.1mm and the probe was implanted to a depth of 4.0mm at a 10-degree angle. The external grounding wire was fixed to the previously implanted skull pin. The craniotomy was then sealed using Duragel (Cambridge Neurotech) and a 3D printed casing was fitted around the probe for protection. Implants were performed using a retrievable system and, in some cases, animals were reimplanted with a second probe. At the end of the experiment, probes were recovered for future reuse.

3.4 Electrophysiological recordings

Animals were habituated to the size and weight of the implant by first training on the task with a size and weight matched dummy implant fixed to the skull. Dummy implants were constructed from a 3D printed plastic casing, aluminium implant cassette and surgical tape. During training, to simulate eventual recording conditions animals were tethered to an overhead cable connected to a motorised rotary joint (Doric, B330-1027-001). Mice were also habituated to sleeping in their home cage while connected to this tether. To increase sleep incidence during recordings, all mice (except for one) were housed with a normal light dark cycle. Neuropixel 1.0 (phase3B) probes were implanted through motor cortex and striatum and after implantation daily recording sessions were conducted as continuous (2-6 hours) recordings across sleep and task epochs. Prior to implantation probes were coated in dye (DiI) for later visualisation. Brains were imaged using a serial section 2-photon microscope, we then registered these volumes to the allen atlas (brain-reg) and the dye tracks were located and traced in common atlas coordinates (brain-reg segment). Recordings were acquired using neuropixels acquisition hardware (imec neuropixels 1.0 headstage, interface

cable and PXIe Acquisition Module) with open-Ephys software. Post-acquisition spike sorting was done using Kilosort3 (Pachitariu et al. 2023). Spike sorting was confirmed using Phy2 (https://github.com/cortex-lab/phy/) but data was not manually curated.

3.5 Pharmacological manipulations

3.5.1 Muscimol

For thalamic muscimol injections, bilateral cranial openings were performed over midline thalamus and a headbar was positioned stereotactically. A landmark that aligned to bregma allowed for future injections in stereotactic positions. Closures of cranial openings were prevented by covering the exposed brain with Duragel (Cambridge Neurotech). Skull was covered with kwiksil (World Precision Instruments), which was removed before every injection. Before each training session, animals were headfixed while awake and ~30nL of either muscimol (Sigma-Aldrich) at 0.2mg/ml or saline buffer were injected for experimental and control sessions. To trace injection sites during histology, injection pipettes were coated in dye (DiI - Sigma). After a 15-minute home cage rest period, animals were tested on the task.

For striatum Muscimol infusions, ~30nL of either muscimol (Sigma-Aldrich) at 0.2mg/ml or saline buffer were infused via implanted cannulas. The infusion system consisted of a 1µl Hamilton syringe (Merck), plastic tubing (P1 technologies cat no. 8F023X050P01) and injection cannulas (P1 technologies cat no. C200IS-5/SPC). Tubing was filled with mineral oil to ensure an air-tight setup for accurate volume administration. Animals were briefly headfixed and infused at a rate of 10nl/min for 5 minutes per cannula. Animals were then allowed to rest in their home cage for 10-15 minutes and tested on the task. Between muscimol infusion experiments animals were given recovery break of at least a day and task behaviour was assessed on this break day ensure performance returned to that pre-infusion. All animals were habituated to headfixing prior to experiment onset.

3.5.2 AP5

We adapted methods described by Lemke et al. (2021). We bilaterally infused 450nl saline or 450nl of 5µg/µl D-AP5 (Bio-Techne, diluted in saline), via cannulae implanted into the dorsolateral striatum. Immediately after training, animals were headfixed and infusions were carried out at a rate of 65-90 nl/min for 5-7 minutes per cannula. After infusion animals were returned to their home cage. In the test session the next day (approximately 24 hours later), animals were trained on a performance sensitive version of the behavioural task (see Behavioural training). This allowed for higher sensitivity in detecting changes to task performance.

Infusions during learning were done from levels 12 to 49 (after the reward guided habituation phase: levels 1 to 11). Infusions were done if animals climbed at least 3 levels that session but were not done on consecutive days. Before the experiment, animals were habituated to head-fixing. Infusions of saline and AP5 were alternated throughout the learning curve of each animal. After learning, once animals reached stable expert performance (level 50) for at least 4 days, infusions of either saline or AP5 were given for 4 consecutive days. All mice were used for both experimental groups and so before switching the type of infusion given animals were trained until at least 4 days of stable expert performance was seen. One animal was unable to reach level 50 with stable performance so was excluded from this experiment.

3.6 Optogenetic manipulations

For opto-inhibition of striatum projecting rILT neurons, in 12% of trials, randomly selected, a sustained 1s pulse of red (632nm) light was delivered after mice initiated the first poke into port 1 of a trial. Light intensity was calibrated to 10mW at the fibre tip.

3.7 Immunohistochemistry

Brain slices were all stained following the same procedure: Blocking in staining solution (PBS + 1%BSA + 0.5%Triton-X) for 1 hour. Primary antibody(s) (1:1000 in staining solution) for 2-4 hours at room temperature or overnight at 4 degrees celsius while rocking. 15 minutes wash with staining solution. Second antibody(s) (1:1000 in staining solution) and DAPI for 2 hours at room temperature while rocking. Slices were then washed in PBS and mounted using Prolong or ProGold mounting medium (Thermo-Fischer). Primary antibodies used were NeuN (abcam, ab177487), GFAP (abcam, ab13970), PH3 (Meck, 07424) and DARP-32 (Biotechne MAB4230). Secondary antibodies used were Alexa-488 anti-mouse (Thermo Fischer, A-11001), Alexa-488 anti-chicken (Thermo Fischer, A-11039), Alexa-488 anti-mouse (Thermo Fischer, A28175) and Alexa-647 anti-rabbit (Thermo Fischer, A-21245).

3.8 Tissue processing and image analysis

At the end of experiments, animals were euthanized via intraperitoneal (IP) injection (10 ml/kg pentobarbital) and brain tissue fixed via vascular perfusion (4% paraformaldehyde) and collected for histology.

3.8.1 Injection and implantation placements

Brains were imaged using a serial section (Mayerich et al., 2008) two-photon (Ragan et al., 2012). Our microscope was controlled by ScanImage Basic (Vidrio Technologies, USA) using BakingTray, a custom software wrapper for setting up the imaging parameters:

- https://github.com/SainsburyWellcomeCentre/BakingTray,https://doi.org/10.5281/zenodo.363160 9

Images were assembled using StitchIt:

- https://github.com/SainsburyWellcomeCentre/StitchIt,https://zenodo.org/badge/latestdoi/57851444

The 3D coordinates of the injections, fiber, cannula and neuropixel probe placements were determined by aligning brains to the Allen Reference Atlas (Allen Reference Atlas – Mouse Brain. Available from atlas.brain-map.org.) using brainreg (Tyson et al., 2021,2022), and visualized using custom python code and brainrender (Claudi et al., 2021).

3.8.2 Quantification of chronic lesions

For striatal lesions brains were sliced using a cryotome at a thickness of 40um and with a vibratome at 100um for hippocampal lesions. 15 to 20 slices covering the entire region at regular intervals were selected for NeuN and GFAP staining. Slices were mounted in standard glass slides using standard mounting medium, and subsequently imaged in the Slide Scanner (Zeiss) using a 20x objective. Lateral and medial striatum were defined by the extent of axons from prelimbic/cingulate projections and motor cortical projections respectively (Allen projection experiment numbers: 157711748, 112514202, 180720175 & 180709942). The lesioned areas for each of these regions was determined manually for each slice using Brainreg segment. Mice with lesions that had more than 20% of volume overlapping with cortex were excluded from analysis.

3.9 Models, metrics & analysis

3.9.1 performance measure

A trial was defined as all the poke events that occurred proceeding reward or trial time out (no pokes for 30s) This meant a single trial could contain multiple attempts at the sequence. For all post hoc analysis, performance was calculated per trial and involved segmenting sequence pokes into ‘attempts’ which were temporally relevant: a sequence of pokes which were within 2s of each other. In other words, if a nose poke occurred more than 2s

after the previous poke, it defined the start of a new attempt. Attempts were considered as starting only from the initiation port (port 1). Any attempts that didn't contain this port at all were ignored and for attempts that did, any pokes that occurred before the first poke into port 1 were excluded. An attempt was assigned a value of 1 if it contained the perfect correct poke sequence (repeat pokes were ignored) into the correct poke sequence. If the attempt contained errors, it was assigned a value 0. The mean score across these attempts was then calculated for each trial giving a score for each trial. The average of these scores is reported.

3.9.2 training performance

During training a simplified measure was used to calculate ongoing trial by trial task performance used for updating training level. For a given trials, a score of 1 was given if (excluding repeat pokes) animals completed the full sequence without adding in additional ports. For anything else, a trial was marked with a zero. Performance was scored as a mean over a window of 10 trials except for AP5 test sessions where this window was reduced to 4 trials.

3.9.3 video analysis

Videos were captured at 60fps and mouse movements were tracked using DeepLabCut (Mathis *et al*. 2018). Tracking points below 98% confidence interval were excluded and replaced by interpolating between accepted points. In task movement variability was first calculated individually for different task subsequence (movement vectors). To achieve this, only trajectories that passed close to each port in the subsequence (within a 1cm radius) in order and with appropriate timing (within a 2s port-to-port time window) were considered. Trajectories were averaged to find the mean trajectory curve. This curve which was then segmented into 1000s of spatial bins to be used as a reference and the distances between each data point in each trajectory and their closest spatial bin were noted. These distances were then used to calculate the standard deviation (movement variability) of each

tracked trajectory from the mean curve. To create the standard space motif occurrence plots during analysis of PPseq output data similar analysis was done. However, average trajectory curves were generated for the entire task sequence rather than individual subsequence chunks.

3.9.4 Data analysis

Unless stated otherwise, all data analysis was done using custom python scripts.

3.9.5 Statistical analysis

A paired t-test (2 groups) or one-sample ANOVA (3 or more groups) were carried out when the assumptions for a normal distribution of observation within groups (Shapiro-Wilks test) were satisfied. Otherwise, the non-parametric Mann-Whitney U test or Kruskal-Wallace test were used. When there were unequal observations between groups, an independent t-test was used, given the assumptions were satisfied.

For the statistical analysis of training learning curves, the data was down sampled into bins of 100 trials. Each animal's learning curve was randomly reassigned to the lesion or control group to look at the mean difference between controls and lesions. This shuffling of learning curves was repeated 10000 times and all 10000 means were used to find the 95% confidence interval for the difference in the data due to chance. A difference between control and lesion groups outside this interval suggests it is statistically significant.

A linear regression model was used to analyse the effect of infusion, saline or AP5, and training level, on the change in levels.

3.9.6 PPseq model

Running PPSeq (Williams *et al.* 2021) required setting 12 hyperparameters. As described in the chapter 6.4, 7 of these parameters were fixed and the remaining 5 were chosen in by grid search. We used cross-validation: a

subset of spikes was held-out from the data, the rest used to train the model, then the log-likelihood of the held out spikes was measured under the trained model. The hyperparameters that lead to the highest held out log-likelihood were taken as those which best capture the structure in the data and could predict heldout spikes. We chose our selected model from the top 20 models (all within error of each other) by visual inspection of the output labelling. Together this specified the 12 values that we used for our subsequent analysis. We performed this hyperparameter fitting on the data from one animal, then used the same values, occasionally scaled for the number of neurons and average firing rate in the data, for all other mice.

The chosen PPseq model was applied to each recording session individually. Striatal neurons were first filtered to remove high and low firing rate units (Fano factor 0.5 -12) and a 600s period of high task engagement awake activity was chosen to fit the model on. Limited by compute power, it was not possible to run PPseq on the entire offline (post task) period. Instead, multiple candidate sleep periods (At least 500s and up to 1500s in length) were selected and PPseq was run on concatenated data from these periods. These periods were manually chosen for times when average firing rate and animal velocity (based on video tracking) were low: indicative of sleep. When applying PPseq to sleep, all the free parameters of the model were fixed per their fit from the awake data except the model was also permitted to fit an additional 2 motifs (latent sequences). Hence, after applying PPseq to sleep, up to latent 8 motifs could be reported by the model. This step was included to allow the PPseq model better flexibility and higher fidelity when fitting to the sleep data. By allowing for new motifs that could be fit to sleep data we aimed to prevent forced fitting of the awake motifs to noisy background spikes.

After running PPseq, the awake and sleep motifs (candidate replay events) were pre-processed before further analysis was performed. Firstly, individual

spike labelling by the model was scored for confidence by comparing across the last 100 iterations of the model (after the log likelihood model fit had plateaued). If a spike was labelled as contributing to a motif type in less than 75% of iterations, then it was classified as background. Single replay events were defined by binning PPseq labelled spikes into 20ms time bins and grouping bins together if they were adjacent and contained spikes for a given motif type. Replay events were then excluded from analysis if they did not contain at least 5 spikes from at least 3 different neurons. Coactive replay events were defined as any sequence of events that occurred within 500ms of each other.

3.9.7 Bayesian state space decoder

The decoder used was as described in Denovelis *et al.* (2021). Models were trained for each recoded session to predict a two-dimensional prior (video tracking position) from spiking. Training data for each model was filtered to only include successful movement trajectories. Filtering was done spatially: only trajectories that passed close to each port (within a 1cm radius) in order and with appropriate timing (within a 2s port-to-port time window) were considered. Spatial bins were 3 times the average distance travelled in one time bin (20ms), on average this corresponded to approximately 400 bins.

Replay detection with the decoder was done by applying the trained model to short data segments of interest; usually 1-5s of data. Detected events were defined as true replay or noise by quantifying the spatial coherence of the decoded position. This was defined by whether the number of spatial bins necessary to explain the prior position up to a 95% confidence interval was within a threshold value. Thresholds were calculated prior to applying the decoder to sleep data. For each event type in each recorded session the threshold number of bins was calculated from the distribution of 95% posterior density for more than 200 hundred awake events and periods of random noise activity. Based on these distributions, the threshold value was

set to maximize the number of true positive while minimizing the number of false positive events.

3.9.8 Synthetic data testing

For synthetic data tests synthetic replay data was generated by implanting PPseq detected motifs into background noise. For each test, motifs were extracted from the corresponding PPseq labelled awake dataset by manually setting an inclusion zone generated by a time window centred on the middle of the detected motifs. For decoder tests all neuron spikes in each inclusion zone were extracted. For PPseq tests values were chosen for top and bottom neuron IDs excluding neurons outside of these values. Motifs were then filtered for representative, non-contaminated motifs. This was done by excluding the top and bottom 25% of motifs based on total motif spikes and the number of contaminant (other motif) spikes in these windows. Motifs that did not occur regularly in the labelled data were excluded from this analysis. From all extracted motifs, 200 motifs were then chosen (selected to maximise equal numbers from all motif types extracted) for implantation into 600s of noise. When required motifs were then manipulated and altered. For warp values which stretched the motifs, fewer than 200 were implanted to avoid overlap between motifs. In rare cases where less than 200 motifs were originally extracted in total then random motifs from underrepresented groups were duplicated to make up this number. The chosen motifs were then randomly ordered and implanted into noise equally spaced apart. Noise data was deleted where motifs were implanted such that the implanted motifs replaced spikes in these regions. Background noise was generated by randomly permuting neuron IDs from awake activity. Hence, spiking content in background noise and the original spike data was identical, but the neuron Ids given to PPseq were shuffled.

CHAPTER 4

---

4. LEARNING AND EXECUTION OF A NOVEL MOTOR SEQUENCE TASK IS DEPENDENT ON DORSOLATRAL STRIATUM

4.1 introduction

This thesis sets out to determine the circuit mechanisms that underpin the formation of procedural memories. To investigate this in mice required a task in which animals learned a procedural motor skill. As described in chapter 2.1, motor sequence behaviours are best described as stereotyped movement chains, composed from a concatenation of elemental motifs (Lashley, 1951; Berridge *et al*. 1987; Graybiel, 1998; Markowitz *et al*. 2018). Hence, to understand how these behavioural chains are constructed by procedural circuits, it is helpful to study behaviour that is composed of multiple easily discretised steps such as typing a pin code or typing a password. To this end I aimed to develop a task in which mice learned a highly stereotyped motor sequence, composed from elemental movements with a clear structure.

4.2 A novel 5-step motor sequence task for mice

In my task, mice were presented with an array of 8 (2x4) wall mounted ports and had to complete a sequence of 5 nose pokes for reward (**figure 4.2a**). Water reward was given on completion of the full sequence and pokes into erroneous ports were not punished. Hence, animals could add additional port pokes into the sequence and still receive reward. Mice were instructed by an automated training protocol which guided behaviour with light and reward (see methods) (**figure 4.2b**). The automated protocol consisted of 50 set training levels of gradually decreasing light and reward guidance up to the final task (level 50) when mice had to complete the full sequence from memory only. As animals performed the task they were automatically promoted to higher levels or allowed to fall down levels based on training performance (see methods). Training level progression across trials formed a standardised learning progression curve for each mouse (**figure 4.2c**).

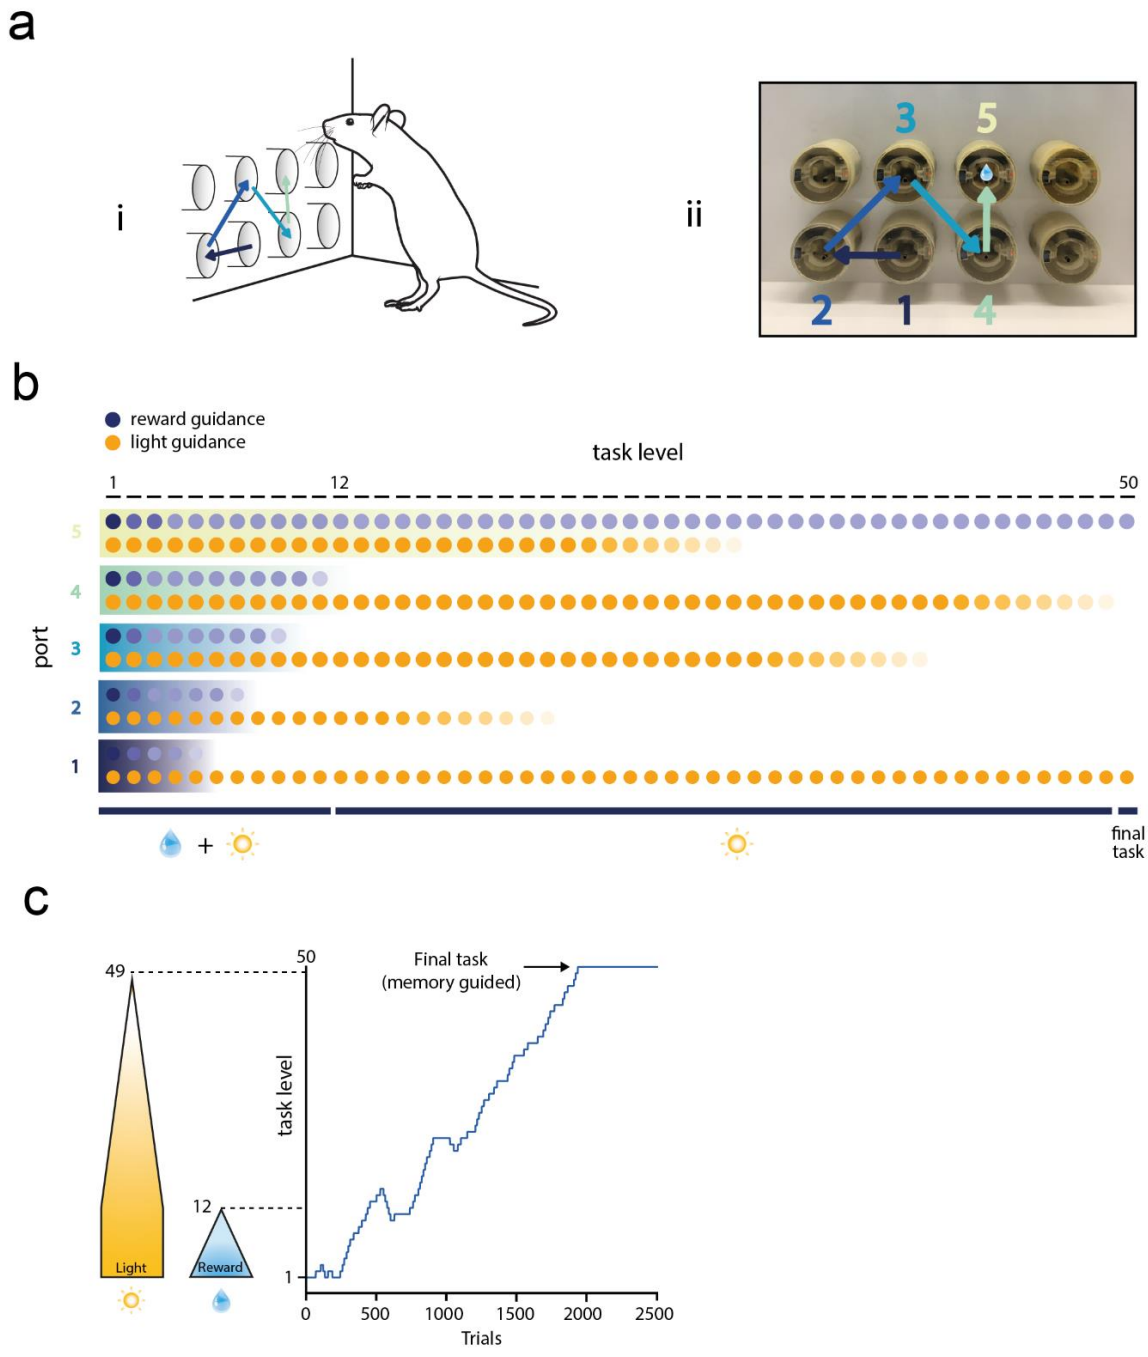

Figure 4.2: A novel 5 step motor sequence task for mice

Legend on next page

Figure 4.2: A novel 5 step motor sequence task for mice

(a) i: Task schematic. Mice are presented with an array of 8 wall mounted ports within a closed arena. ii: Photograph of the 8-port poke wall with poke sequence overlayed. Mice initiated trials by poking into port 1 and the had to complete a further 4 sequential pokes in the order illustrated. **(b)** schematic showing light guidance and reward delivery at each port in the sequence for each training level. In early levels port order is guided by both light and reward, by level 50 only port 1 is illuminated prior to trial initiation and only port 5 is rewarded after memory guided completion of the full sequence. **(c)** Example animal learning curve showing training level progression and regression across trials.

4.3 Mice learn to complete the 5-step sequence task by producing stereotyped movements

Under the automated training protocol animals were able to learn to complete the full task from memory. On average, mice reached the final memory guided version of the task (level 50) after roughly 2000 (mean = 1942 +/- 111, SEM) trials with all animals able to reach the final task in under 4000 trials (**figure 4.3a, b**). After reaching the full task, mice completed the sequence with highly stereotyped timing as is demonstrated by the regularity of nose pokes into each port relative to sequence onset (**figure 4.3c**). Port-to-port transition times were very fast and similar across trials the full extent of the task, though the final transition (transition 4), which was most proximal to reward, was completed significantly faster than any other (**figure 4.3d**). This speed difference cannot be explained by port-to-port distance since transition 4 was equal in length to transition 1. As well as completing the task

rapidly, sequence completion across mice was very accurate such that trained animals made few port-to-port transition errors (**figure 4.3e**) and were also highly task focused, rarely poking into task irrelevant ports (**figure 4.3g**). In some training sessions video data was also captured. For these sessions, analysis was extended to examine movements via tracking of the head position during the task. While novice animals with less task experience produced relatively variable movements from trial-to-trial (**figure 4.3f[i]**), expert mice produced far more stereotyped reproducible movements (**figure 4.3f[ii]**). Overall movement variability for a given session was calculated across 4 subsequence movement vectors (**see figure4.3f [i & ii]**) and reported as the average of the standard deviations of all tracking trajectories from the overall mean trajectory for each subsequence. Comparing this to the number of trials the mouse had completed previously revealed that as animals gained more task experience (completed more trials), trial-to-trial movement variability decreased (**figure 4.3f**) suggesting movements became more stereotyped with practice. Variability plateaued after around 4000 to 8000 trials. Since all mice reached the final task level before completing 4000 trials, this suggests animals continued to improve for a short time after reaching the full task. Rapid, highly stereotyped and accurate movements are hallmarks of learned procedural skills. Since animals perform this task with highly stereotyped movements, both in terms of timing and spatial accuracy, suggests that when completing this task mice learned a procedural skill. This is a crucial outcome which makes this task highly appropriate for studying the mechanisms that underpin the formation of procedural memories. In my experiments therefore, the procedural memory I am studying is that of a motor skill which encompasses full body movements via a series of stereotyped postural adjustments.

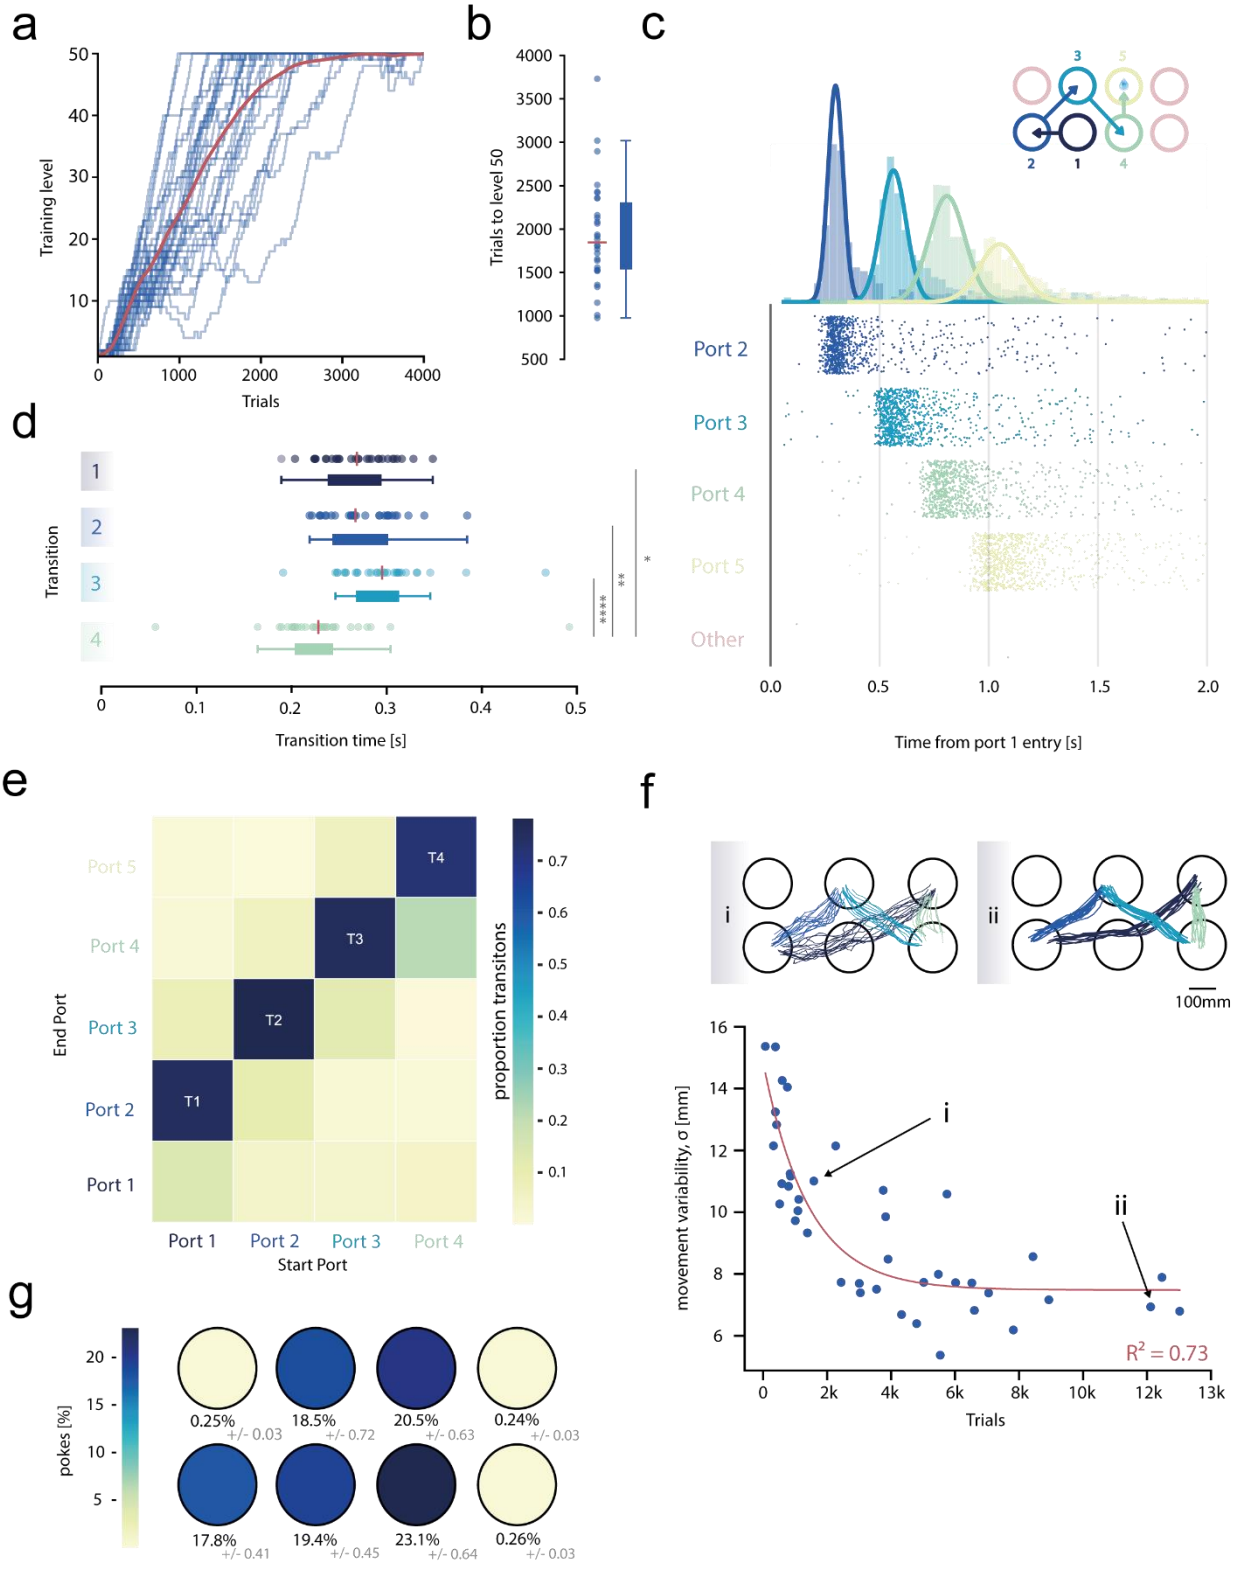

Figure 4.3: Mice learn to complete the 5-step sequence task by producing stereotyped movements

Legend on next page

Figure 4.3: Mice learn to complete the 5-step sequence task by producing stereotyped movements

(a) Training level progression curves (blue) for multiple animals with mean learning curve overlayed (red) (n = 33 animals). **(b)** Number of trials before the final task (level 50) was reached for the animals show in (a) (mean = 1942 +/- 111, SEM) **(c)** Example poke times from a single session of a trained animal that had reached the final task. Port-poke-in times are shown (points) for each port relative to trial initiation for all sequence related ports, poke data is combined for the 3 remaining sequence irrelevant (other) ports. Data is summarised above by a density histogram (binned at 25ms intervals). **(d)** Transitions intervals between ports (poke out to poke in) for trained expert animals completing the 4 task relevant transitions (Kruskal-Wallis, $p = 4.4e^{-7}$, post-hoc Dunn comparisons test, displayed stars: $p = 1.8e^{-7}$, 0.0007 & 0.006). **(e)** Transition heatmap showing mean port to port transition proportions for multiple trained animals. Each transition in the task sequence is represented by its constituent start (x-axis) and end (y-axis) ports and the desired correct transitions are marked (T1 - 4). **(f)** Bottom: average movement variability (standard deviation from average trajectory) across all subsequence task movements for mice at different levels of task experience (n = 8 mice). Above: example tracking data showing subsequence trajectories from a relatively naïve mouse (i) and an expert animal (ii) (15 trajectories per subsequence shown were chosen randomly from all trajectories, excluding those with standard deviations outside of the interquartile range of standard deviations) (Number of trials is total completed prior the tracking session, the tracking point was the centre of the head). **(g)** Mean port poke occurrences for trained animals (SEM for each port shown in grey). Data (d - f) is summary data taken from trials 3000 to 3500 for the animals shown in (a).

4.4 Lesion to dorsolateral striatum impairs formation of procedural memory

As discussed in chapter 2.2, the DLS is thought to be a key region, necessary for both learning and execution of motor skills (Miyachi et al., 1997; Yin and Knowlton 2006; Yin 2010; Wolff *et. al.* 2022). I next aimed to test if this is true for my task. To do this I performed lesions to the striatum by injecting either viral caspase or NMDA. Mice were injected bilaterally into the DLS or DMS and then trained on the task a few weeks later (**figure 4.4a**) (for lesion extent see **appendix figure 1**). Compared to saline injected control mice, DLS lesioned animals showed impaired task learning (**figure 4.4b**). All control mice reached the final level while only one lesioned animal was as able to reach this criterion within 4000 trials. Shuffle analysis reveals that the lesioned cohort learned the first 20 levels at a comparable rate to controls (roughly 1500 trials) and learning only diverged after this point. This divergence corresponds with training levels where port lights were switched off entirely (**figure 4.2b**). After these light guided levels animals must increasingly rely on memory guided strategies and so the deficit observed in lesioned animals suggests they had issues in forming memory of the motor sequence. To test if these observed motor learning deficits are region specific within striatum, I also ablated the medial portion of the dorsal striatum (DMS) (**figure 4.4c**) (for lesion extent see appendix fig1). Lesions to DMS did not impair task learning compared to saline injected control animals (**figure 4.4d**). All but one DMS lesioned mouse reached the final task within 4000 trials and these mice learned at a comparable rate to control animals. Overall, I find that lesions to the dorsolateral part of striatum uniquely impair motor sequence learning in my task and this deficit specifically corresponds to problems when task execution is dependent on a memory guided strategy.

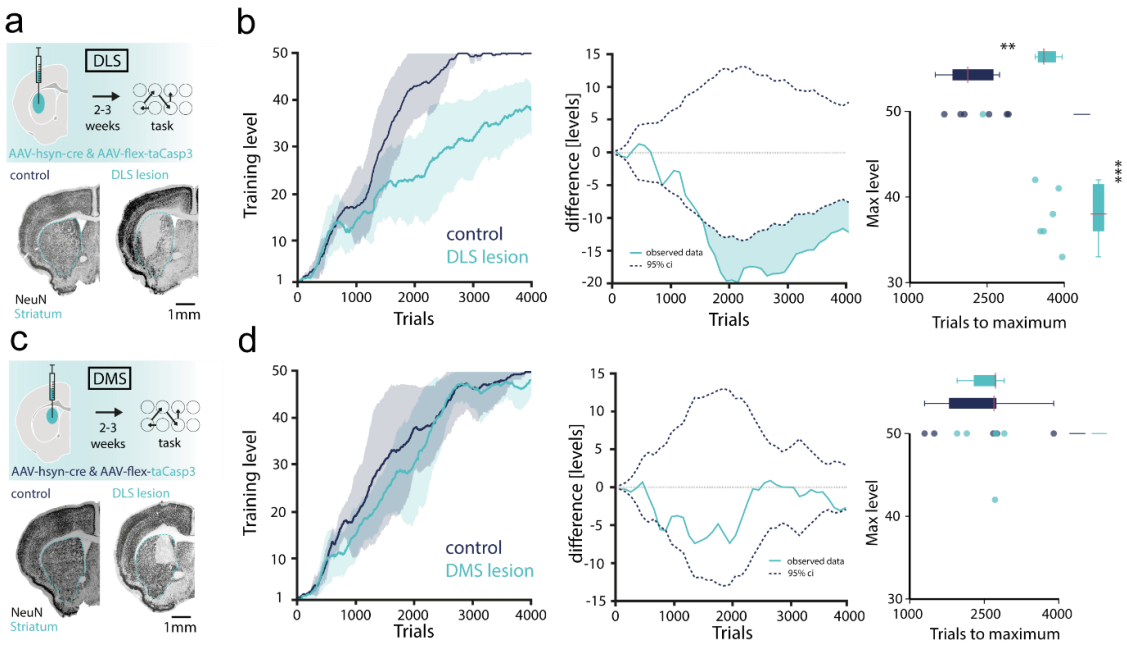

Figure 4.4: Lesion to dorsolateral striatum impairs formation of procedural memory

(a) Top: schematic diagram showing experimental design for DLS lesions. Injection surgery was performed in naïve mice, 2-3 weeks prior to training onset. Bottom: example slice histology showing lesion extent. **(b)** Left: average training level progression curves for control and lesion animal groups (shaded area denotes standard deviation). Middle: differences in performance between the groups. Dotted lines indicate the 95% confidence interval for shuffled data (see methods). Right: maximum training level obtained within 4000 trials against trials taken to reach maximum (max level, p = 0.001, trials to max, p = 0.0008, independent t-test) (lesion; n = 7 mice, control; n = 7 mice). **(c-d)** Same as panels a-b but for DMS lesion (lesion; n= 6 mice, control; n = 6 mice) .

4.5 Lesion to dorsolateral striatum causes previously acquired procedural memory to be lost

If the DLS is required for procedural memory formation, this region may also be responsible for long term memory storage, after initial learning has taken place. To test if the DLS was also required for post learning task execution, I ablated this region in highly trained expert animals using a viral-mediated Caspase-based strategy (**figure 4.5a**). Lesions to the DLS were catastrophic to the previously learned behaviour. Animals still attempted the task by poking into multiple ports, however, average task performance - based on a measure of proportion perfectly accurate sequence completion (see methods) – fell from 0.8 (approximately 4 out of 5 sequence attempts correct) to just 0.2 (1 out of 5 sequence attempts correct) (**figure 4.5b**). Poor task performance post lesion was explained by a dramatic increase in the proportion of port-to-port transition errors animals made, suggesting mice lost memory of the accurate transition structure of the task (**figure 4.5c,d**). Classifying error types into repeats (consecutive pokes into the same port), mis-sequences (incorrect, but task relevant, port transitions) and other errors (poke transitions to task irrelevant ports), I found the increase in errors post lesion was generalised; not marked by increased occurrence of a specific type of error (**figure 4.5e**). As well as losing memory of task structure, lesioned mice also completed correct port-to-port transitions more slowly and with greater temporal variability indicating a loss of motor stereotypy (**figure 4.6f**). In line with this, for a subset of animals that were filmed before and after lesion, tracking analysis on four correct subsequence movements showed that after DLS lesion, trial-to-trial movement variability increased (**figure 4.5h**). Despite loss of procedural aspects of task memory, mice were still able to move to all ports and remembered which ports were involved in the task: rarely poking into task irrelevant ports and favouring the final (reward) port (**figure 4.5g**). Overall, these experiments provide strong

evidence that DLS is a key site for both procedural memory formation and long-term storage. Ablation of this region caused loss of previously learned procedural memories while leaving other aspects of task memory intact. These findings are supported by a complimentary acute DLS perturbation experiment that we performed (**Appendix figure2**). Infusions of muscimol (GABA-A receptor agonist) into the DLS also caused expert mice to temporarily forget the previously accurate transition structure of the task. Muscimol temporarily silences neural activity and hence, in line with my lesion findings, temporary loss of DLS function caused acute loss of procedural memory of the task.

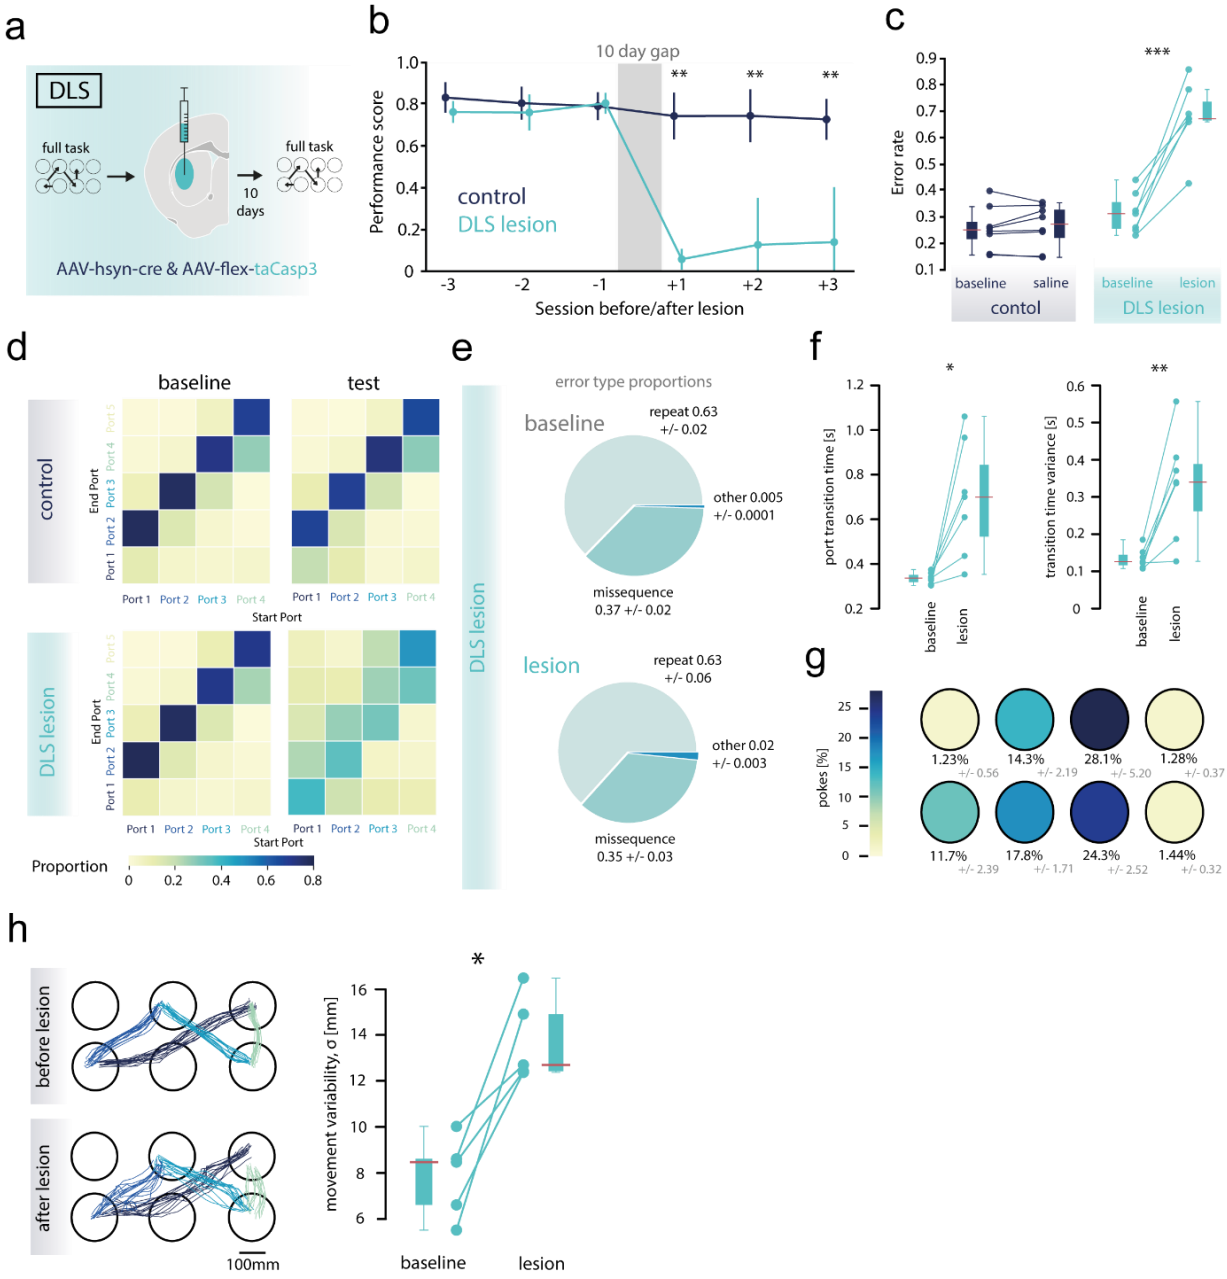

Figure 4.5: Lesion to dorsolateral striatum causes previously acquired procedural memory to be lost.

Legend on next page

Figure 4.5: Lesion to dorsolateral striatum causes previously acquired procedural memory to be lost.

(a) Schematic diagram showing experimental design for DLS lesions performed in trained expert mice. **(b)** Task performance scores (see methods) for the 3 sessions prior to and following injection surgery (ANOVA $p = 1.29e^{-26}$, Tukey HSD: p = 0.001, 0.001 & 0.001. lesion, n = 7 mice; control, n = 7 mice). **(c)** Proportion of port-to-port transition errors in the 3 sessions before and 3 sessions after injection surgery (p = 0.004, paired t-test). **(d)** Transition histograms before and after injection surgery. **(e)** Error type as a proportion of all errors made for the lesion group before and after injection surgery (average is mean +/- SEM). **(f)** left: port to port transitions times (port out to port in) for correct transitions before and after injection surgery for the lesion group (p = 0.01, paired t-test). Right: variance (standard deviation) of transitions times before and after injection surgery for the lesion group (p = 0.006, paired t-test). **(g)** Average percentage port poke occurrences across all lesioned animals for the 3 sessions after injection surgery (grey numbers are SEM). **(h)** Left: example animal movement tracking for 4 subsequence movement vectors before and after lesion (15 trajectories per subsequence shown were chosen randomly from all trajectories, excluding those with standard deviations outside of the interquartile range of standard deviations. Tracking point was the centre of the head). Right: Average movement variability (standard deviation from mean trajectory) across the four subsequence trajectories before and after lesion to DLS (P = 0.0085, paired t-test, n = 5)

4.6 Discussion

4.6.1 A 5 step sequence task for mice

This thesis sets out to determine the circuit mechanisms that underpin the formation of procedural memories. Investigating this first required a standardised procedural behaviour in mice. I chose to develop a novel 5 step sequence task with a clear compositional structure akin to typing a pin code. This compositionality offered several benefits. Firstly, since, procedural motor skills are stereotyped movement chains, composed from a concatenation of elemental motifs (Lashley, 1951; Berridge *et al*. 1987; Graybiel, 1998; Markowitz *et al*. 2018), the clear piece-wise structure of the task has use in understanding how this chaining process occurs. Additionally, the task setup is highly flexible and so allows for future work mutating the sequence. For example, the port order can be altered very simply to add repeating elements, or even to allow training of multiple sequences with overlapping or shared motor elements. Though outside the scope of this thesis, these manipulations are likely to be very important for future work understanding how the brain hierarchically uses and reuses motor elements across and within sequence chains. Finally, I hoped that the piece-wise structure of the task, and the ability to guide animals with both light and reward throughout the task based on an automated training protocol, would make a relatively complex behaviour for mice, simple and quick to train. Indeed, I found mice were able to learn the sequence task quickly and, once expert, these animals produced am accurate series of postural movements. As well as being accurate, these movements were very fast and highly stereotyped, as would be expected from a procedural skill. Interestingly, despite being the same distance as other movements in the sequence, I found the final movement in the sequence was completed significantly faster across animals. It is unclear why this was the case. An intriguing possibility is that animal's movement vigour is modulated by expected reward proximity,

as had been suggested previously (Takikawa *et al*. 2002; Rigoux & Guigon 2012; Shadmehr *et al*. 2016; Manohar *et al*. 2017; Summerside *et al*. 2018) However, with the data I have shown I cannot rule out that the upwards movement required by this final part of the sequence simply has a biomechanical advantage. While the answer to this question lies outside the scope of this thesis, if motor sequence elements are invigorated differentially, this task could be useful for future work determining how this modulation is achieved within procedural circuits.

4.6.2 A role for the DLS in task learning

Having established a task in which mice learn a procedural skill, the next aim was to determine in which region of the brain memory for this skill is learned and stored. There is strong evidence that the striatum is key nexus, involved in both motor sequence learning and post learning execution (Miyachi et al., 1997, Berridge & Whishaw, 1992; Yin and Knowlton 2006; Yin 2010; Wolff *et. al*. 2022*)* I find evidence that this is also true for mice learning the sequence task since I found that specific lesion of this region resulted in impaired task learning and execution. The learning deficits observed only started on average at around training level 20. This may be explained by the training guidance provided to the animals during the initial stages of the task. While guided by lights, the circuits which support procedural memory may not be as crucial. However, as this guidance begins to disappear past level 20, the need for procedural memory becomes more prominent, leading to the observed deficit.

Though animals with DLS lesions learned the task more slowly, they were not completely impaired. Unfortunately, the experiment was curtailed before performance completely plateaued in the lesion group so it is unclear whether these animals could have reached the full task. However, at the very least these animals were able to reach high training levels, indicating that they could perform the task using a partially memory guided strategy. If this region

is needed for learning the task, why could these animals learn the task at all? Several explanations which are equally viable (and not mutually exclusive) can rationalize this observation. Firstly, the DLS is a large area and my lesions were not complete. It is possible that the remaining cells in this region were sufficient to support learning – albeit at a slower rate due to the reduced population size. Secondly, after chronic lesion, it is possible that compensatory mechanisms in other brain regions are able take over and support task learning in the absence of the DLS. For example, the motor cortex is known to support motor skill learning. It is thought that in a healthy brain this region disengages after early learning, and subcortical circuits take over (Kawai et al. 2015; Kupferschmidt *et al.* 2017; Hwang *et al.* 2019 & 2021; Wolff et al., 2022). In the absence of a healthy DLS, motor cortex may continue to support learning throughout training. If mice with lesion to DLS are learning the task via a compensatory mechanism, it is unclear to what extent - if any – these compensatory circuits support normal learning in a healthy brain. Hence, from these data is it only appropriate to conclude that the DLS plays a role in procedural learning, and not that it is uniquely responsible for this. Nonetheless, these results are in line with previous research and do support the notion that the DLS is a key brain region for procedural memory formation.

4.6.3 A role for the DLS in task execution

Besides procedural memory formation, I also found that this region was crucial for long term storage of procedural memory - after learning has taken place. Two complimentary experiments (lesions and muscimol infusions) corroborated this finding. Since only procedural aspects of task memory were dependent on DLS – for example, mice still poked into task related ports – it suggests that the DLS is specifically involved in procedural memory storage. This result is in line with the notion different memory types are independently controlled by distinct circuits (White & McDonald 2002).

One limitation is that the animals used in these experiments were not directly tested for their ability to properly control bodily movements. Hence, whether animals lost procedural memory for the task, or merely lost the ability to properly express procedural memory cannot be stated with complete certainty. However, besides during task execution, no obvious motor control deficits were observed. Furthermore, even during task attempts animals were still able to poke into all ports and able to complete subsequence movements alone, albeit in a less stereotyped manor. Hence, it seems highly unlikely that a generalised motor control issue was responsible for the deficits observed in these animals. The observation that only task related movements were impaired by DLS lesions is in line with previous work (Mizes *et al*. 2022). Further, this notion also fits within the current conceptions of how the DLS is thought to contribute to motor control. As described in chapter 2, the striatum is thought to learn state-action associations and learn how to contribute to behaviour in a manner that is dependent on the demands of the behaviour in question. In my task, DLS lesions appeared to cause three distinct deficits that have all previously been observed in other task contexts. Firstly, in other task settings the striatum has been suggested to control action selection (Park *et al*. 2020). This was also true in my task. Animals lost serial ordering of the sequence, implying the DLS had been contributing to the control of high level behavioural ordering in my task. Secondly, in previous work the DLS has been shown to contribute to the low level – kinematic – control of movements (Dhawale *et al*. 2021). In my experiments the stereotyped port-to-port transition movements were lost and animals reverted to more variable movements suggesting the DLS had also been contributing to kinematic motor control. Finally, the striatum has been suggested to be a controller of motor vigour (Jurado-Parras *et al*. 2020). This also seems to be true in my task where animals performed port-to-port movements more slowly after DLS lesions.

4.6.4 Summary

In summary, in this chapter I described a novel motor skill learning task for mice. I have shown that animals learn this task well and produce stereotyped procedural movements from memory. I find that a region of motor cortex recipient dorsolateral striatum is transcriptionally active after task practice and, in line with previous work, I find lesions to this region impaired both task learning execution. Altogether, the work in this chapter demonstrates that the dorsolateral striatum is a key region, supporting both procedural memory formation and execution. However, the striatum is not spontaneously active – requiring a source of glutamatergic input. Therefore, open questions remain as to what drives striatal activity during procedural behaviours. As discussed in chapter 2.2, while motor cortex is thought to support early learning, evidence suggests that a thalmo-striatal pathway is required for both forming and performing procedural memories. I have proposed that thalamic inputs may be necessary for chaining motor elements as part of an efference loop architecture. In the next chapter I will build on the work in this chapter, extending the procedural circuit by exploring a possible role for the thalamus in supporting procedural memory for my task.

CHAPTER 5

5. THALAMIC CONTRIBUTIONS TO MOTOR SEQUENCE LEARNING AND EXECUTION

5.1 introduction

So far, I have established a novel procedural sequence task and demonstrated that mice perform this task well, learning to produce a stereotyped series of postural adjustments from memory. I have provided strong evidence that both formation, and long term storage of procedural memory for this task involve the DLS. However, as outlined in chapter 2.2, the striatum is not spontaneously active and hence, requires glutamatergic drive from an external source to be active. If the striatum is the site where procedural memories are formed, it therefore cannot act alone. While motor cortical inputs to striatum have been shown to be necessary for procedural memory formation, evidence suggests that these inputs are no longer needed for executing an already learned, perfected, motor skill (Kawai *et al*. 2015, Dwhale *et al*. 2021). The other major source of glutamatergic drive to the striatum comes from thalamus. Several thalamic nuclei could provide this input for motor sequence production, though multiple lines of evidence point favourably towards the rlLT (see chapter 2.3). I have proposed this region forms part of an efference feedback loop, routing motor information from midbrain and brainstem nuclei to the striatum. Such feedback would enable the formation of action-to-action associations and hence, facilitate chaining of elemental actions into stereotyped procedural motor sequences.

5.2 Midline thalamus including rlLT is required for motor sequence chaining

If thalamic input to the striatum facilitates action-to-action chaining, then activity in the thalamus will be needed in order to generate the full continuous extent of the learned motor sequence. To broadly investigate whether this is true I injected muscimol (a GABA-A agonist) into the thalamus of trained animals. I specifically targeted midline thalamus in order to ensure the rlLT was silenced in these experiments (**figure 5.2a**). Microinjections were

performed via predrilled drilled sealable craniotomies (see methods) and animals were placed into the behavioural area 5 -15 minutes later and tested on the task. After saline injections mice performed comparably to non-injected baseline sessions, making few port-to-port transition errors (**figure 5.2b,c**). Injection of muscimol, however, caused acute task deficits characterised by reduced ability to produce the previously known correct port-to-port transitions (**figure 5.2b,c**). These effects wore off and at a timescale expected for the breakdown of muscimol (Baraldi *et al*. 1979) after which animal performance returned to normal. Analysing port-to-port transitions during the first hour of task performance, it is clear that silencing midline thalamus with muscimol did not cause a generalised increase in transition errors, but instead had a distinctive effect on port-to-port transition structure (**figure 5.2d**). Muscimol specifically increased one type of error: repeat errors; repeated pokes into the same port within 2 seconds of initial port entry (**figure 5.2e**). This increase in port repeats was an acute effect of muscimol (**figure 5.2f**) and was far more prevalent than any other type of transition error mice made (**figure 5.2g**).

Overall, I find broadly silencing midline thalamus caused mice to lose the transition structure of the task, such that animals could no longer make correct port to port transitions and instead repeatedly poked into task relevant ports.

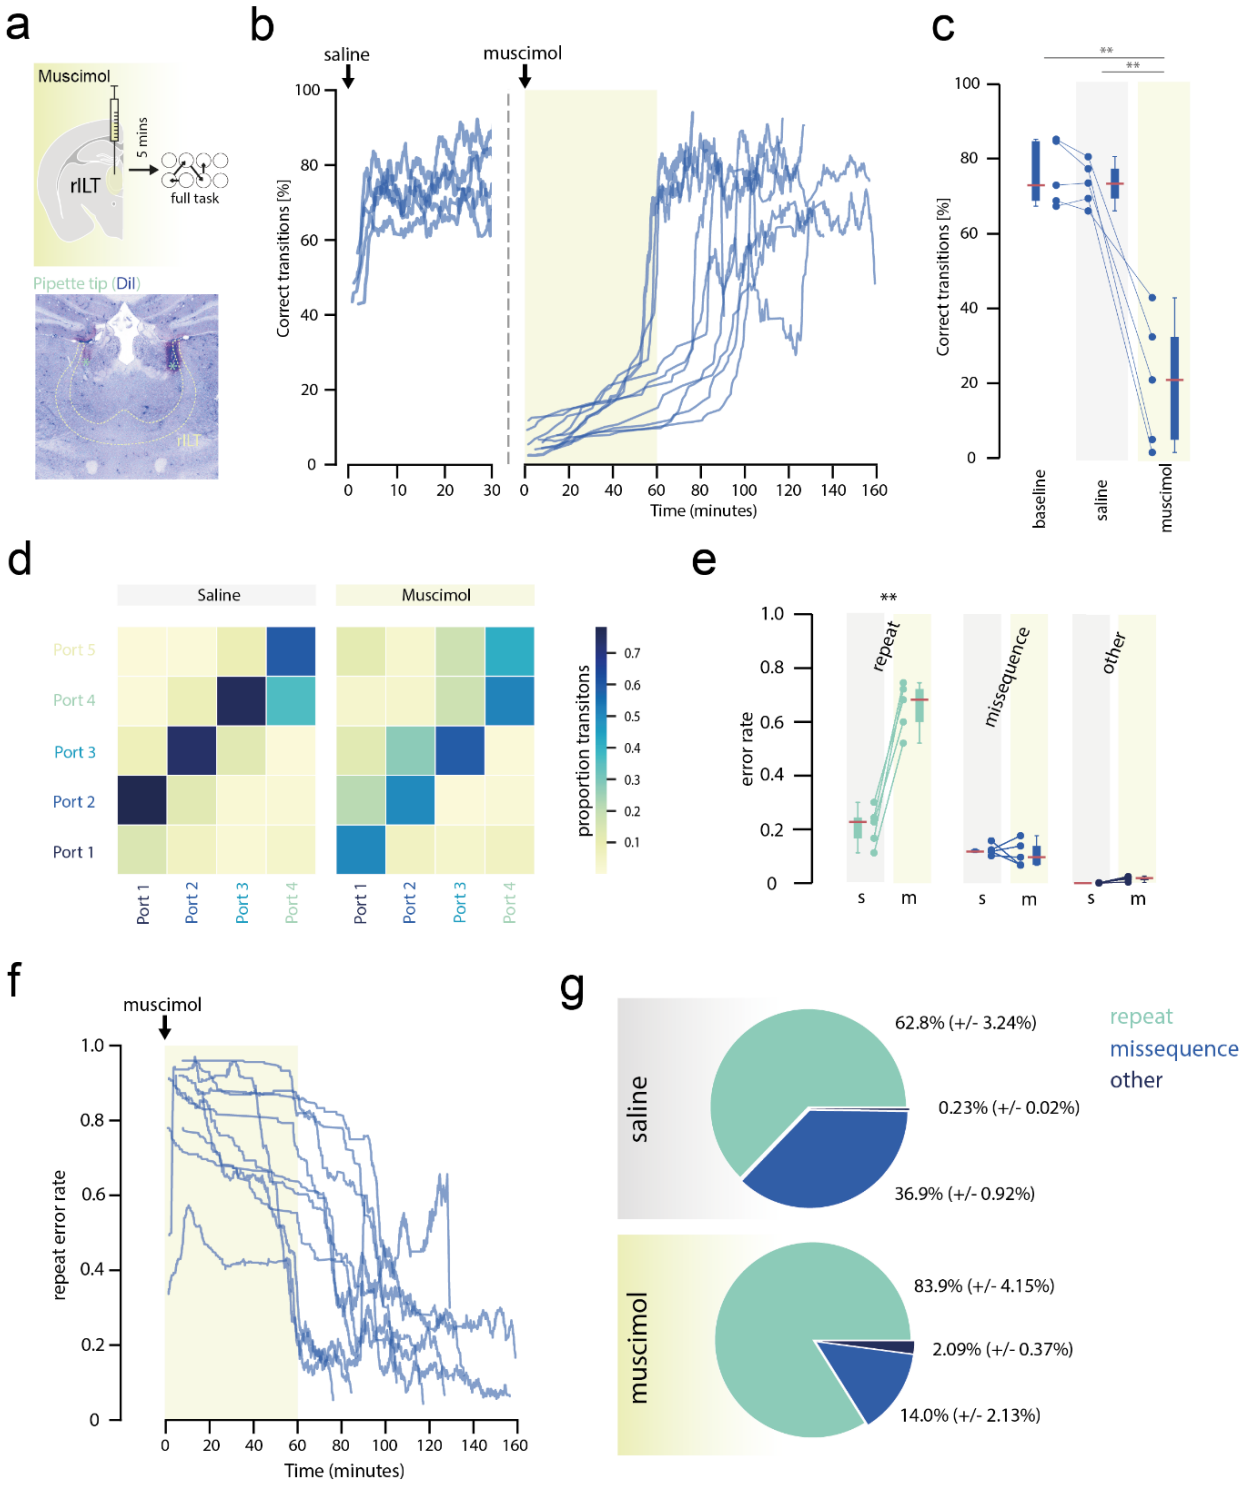

Figure 5.2: Midline thalamus including rILT is required for motor sequence chaining

Legend on next page

Figure 5.2: Midline thalamus including rILT is required for motor sequence chaining

(a) Top: schematic showing experimental design. Bottom: example histology showing DiI residue from injection pipette tips (purple) and rILT outline (yellow). **(b)** Correct port-to-port transitions against time after saline injection (left) and muscimol (right) (2 muscimol sessions performed per animal, n = 5 mice). **(c)** Average correct port-port transitions during the first 60 minutes of behaviour (red lines = mean, $p = 1.0e^{-5}$,ANOVA. Tukey HSD: p = 0.001, p = 0.001). **(d)** Transition heatmaps in 60-minute window after injection. **(e)** Error type occurrences as a proportion of all port-to-port transitions ($p = 1.6e^{-15}$,ANOVA. Tukey HSD: p = 0.001). **(f)** Repeat error occurrences per transition plotted against time for muscimol sessions. **(g)** Average transition error types as a percentage of all errors (+/- SEM).

5.3 Optogenetic silencing of DLS projecting rILT leads to sequence chaining deficits

The results from my muscimol experiments provide intriguing evidence that the thalamus has a role in chaining actions together. Though promising, if a region of midline thalamus does indeed facilitate action-to-action chaining for learned procedural memories, several problems remain. Firstly, the muscmiol experiment was not area specific: the spread of muscmiol is highly general and so, despite the rILT being targeted, the particular thalamic region or subregions which provide input to striatum cannot be determined from this experiment. Secondly, and importantly, silencing thalamus in this way is not pathway specific: besides silencing inputs to the striatum, injected muscimol will have had other off target downstream effects. Most notably intralaminar thalamic nuclei have dense bilateral connectivity with cortex (Catsman-Berrevoets & Kuypers, 1978; Royce, 1983; Velayos *et al*. 1989; Sadikot et al 1992; Parent & Parent, 2006). Hence, silencing the thalamus in this way may have also altered excitability across extended regions of cortex.

To overcome some of these experimental limitations, I aimed to disrupt only striatal projecting cells in rILT during task execution. To do this, I injected an AAV-retrograde-JAWS into the striatum to trigger expression of inhibitory opsin in upstream, striatum projecting cells. For control mice this virus was replaced with an AAV-retrograde-GFP. Fibres were implanted bilaterally over the rILT (**figure 5.3a**) and mice were trained to expert level and habituated to the fibre optic tethers. During stimulation sessions, on 12% of trials, nose pokes into the first port in the sequence triggered a continuous 10mW red laser for 1s. Inhibition of striatum projecting rILT neurons via laser stimulation in mice expressing JAWS caused animals to make significantly more errors on the first transition in the sequence (**figure 5.3b**). To better understand what this effect on behaviour was, analysis was extended to the types of errors made during laser stimulation trials. If rILT inputs to striatum provide motor efference information – such that current action is chained to (and triggers) the next action in the sequence, then inhibiting these inputs should break this chain. Indeed, in line with the effects of the muscimol experiment, I find optogenetic perturbation of striatum projecting rILT neurons led to a specific increase in port repeat errors (**figure 5.3c**). As before, mice not only more often failed to produce the next sequential action after poking into the first port, but also more often repeated the action that was previously taken. Inhibition also had a notable and consistent effect on transition speed, causing mice to complete the first port-to-port movement more slowly (**figure 5.3d**). After the laser was turned off no deficits were observed, and mice were able to complete the sequence normally.

To summarise, I find that optogenetic inhibition of striatum projecting neurons in rILT led to reduced movement speed and increased repeat poke errors in trained mice. Notably, this result is in line with findings from the previously discussed muscimol experiments. Together these results are in line with the hypothesised role of rILT as a key nucleus involved in action-to-action chaining.

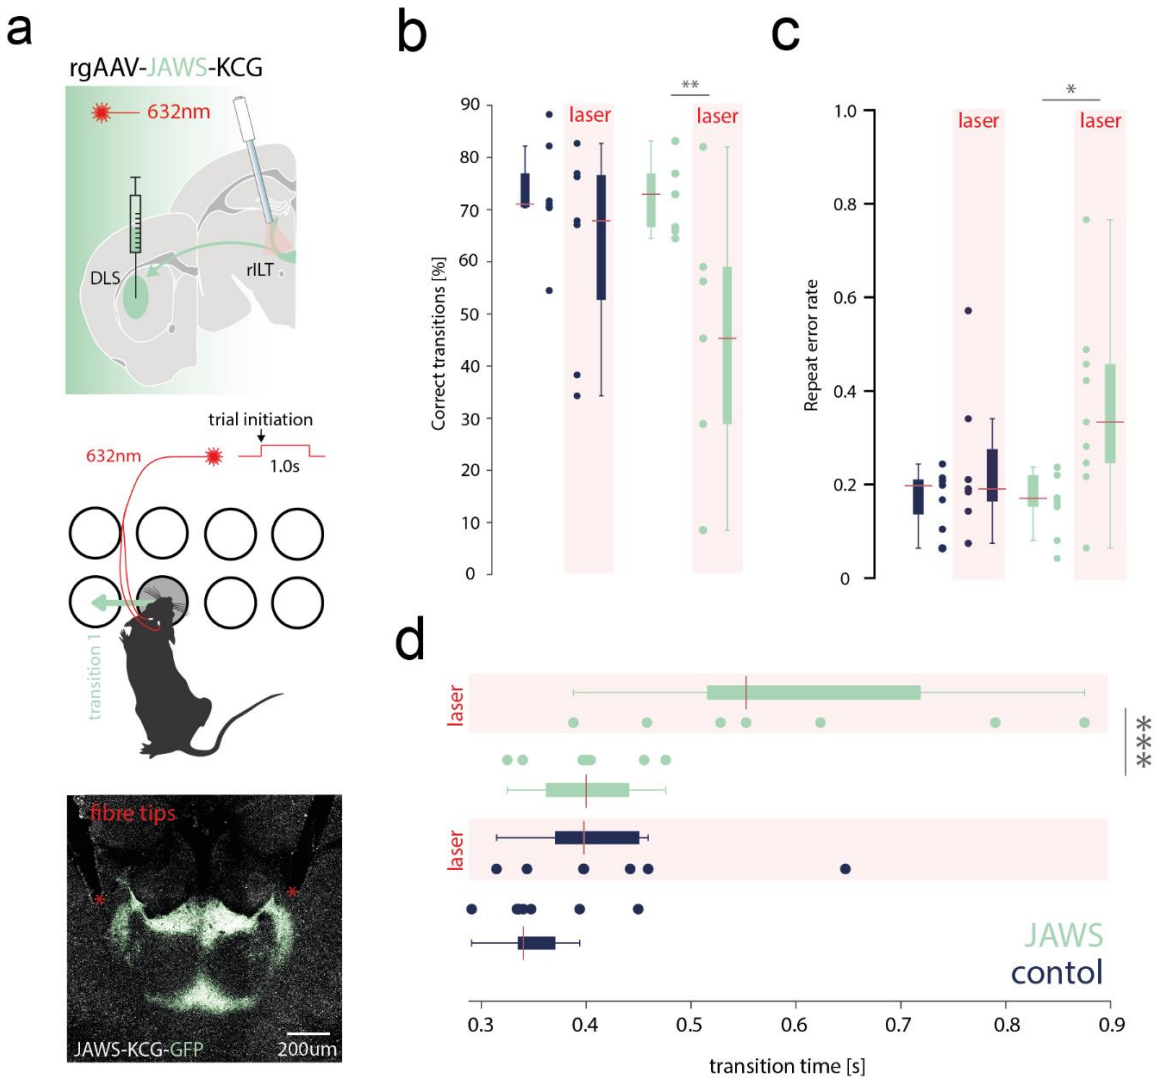

Figure 5.3: Optogenetic silencing of DLS projecting rILT leads to sequence chaining deficits

(a) Top: schematic showing virus strategy and fibre placement. Middle: schematic showing experimental design. Laser was triggered on 12% of trials by port 1 poke in events. Bottom: florescence image showing virus expression in rILT and bilateral fibre locations (fibre tips are marked by red stars) **(b)** Percentage correct port-to-port transitions during laser and baseline sessions for control and opsin mice. (p = 0.005, paired t-test) **(c)** Repeat error type occurrences as a proportion of all transitions. (p = 0.008, paired t-test) **(d)** Port to port transition times (p = 0.0098, paired t-test, n = 7 mice per group)

5.4 Ablation of GRM2, a genetically identified population in rILT, impairs motor sequence learning

As outlined in chapter 2.3, I hypothesise that projections to striatum from rostral rather than caudal ILT (PF) are essential for action-to-action chaining. Thus far I have shown some evidence that long term storage of learned motor skills is contingent upon a thalamic region and, using an optogenetic approach, I have provided evidence that striatal projecting rILT cells are required for motor sequence chaining. However, several issues provoke further experimentation. Firstly, though striatum projecting cells were specifically targeted, it's unclear whether these neurons also send collaterals to other regions; for example to cortex. It cannot therefore be ruled out as a possibility that the behavioural effects of inhibition were mediated by suppression of a non-striatal pathway. Secondly, JAWS was also expressed by cells in the caudal nucleus of ILT, the PF. Though fibreoptic ferules were implanted over rILT, it cannot be ruled out that inhibition of the PF (due to light spreading caudally) could have been responsible for the effect seen. Finally, the optogenetic effects observed were not particularly strong or consistent. This suggests that the approach used was possibly not strong enough, or was not able to influence a large enough proportion of the functional population to consistently disrupt activity in this circuit.

Stronger evidence to determine what the role rILT has in controlling motor sequence behaviours requires a better approach. Ideally this approach should aim to be pathway specific but also have improved capacities to segment rostral and caudal ILT nuclei, and greater capability to perturb the circuit during behaviour. These objectives are difficult to realise experimentally. Firstly, separating rostal and caudal ILT is challenging. These regions are adjacent, share many common inputs and any functional separation is most likely highly overlapping between cells in the two subregions. Further, while other optogenetic or chemo-genetic tools could offer greater control over thalamic activity, it seems likely that the main

limiting factor lies in the ability to influence a large enough proportion of the relevant functional population in rILT. This nucleus is a long thin sheet of cells and so light from an implanted ferule can only be targeted onto a portion of it and this shape also makes expressing injected opsin in this region very difficult. While retrograde viral injections in the DLS overcome this issue to some extent, the efficacy of retrograde transduction is unlikely to be complete (Haenraets *et al*. 2017). Additionally, although rILT axons are high ramified in the striatum (Parent & Parent, 2006), the large volume of this region makes it difficult to target thalamic projections in great quantities.

Driven by these constraints, I was motivated by the possibility that a genetically defined rILT subpopulation exists which is essential for motor sequence generation. Finding a genetic marker could allow for a better ability to target cells in the rILT for inhibition and hence, help solve the experimental limitations I have outlined. A candidate subpopulation are glutamate metabotropic receptor 2 expressing cells (GRM2). GRM2 has dense expression in rostral IL thalamus (centromedial (CM) and centrolateral (CL) nuclei, although does have some sparse expression in caudal (PF) nuclei and surrounding thalamic regions. (**See appendix figure 3a**). The anatomical distribution of this population is promising, however, for this approach to be useful, I first had to determine whether these neurons are involved in motor sequence execution. If GRM2 cells are required for driving motor sequence execution, then mice should not be able to learn the task when this population is ablated. To test this, I used a GRM2-tdtomato-Cre mouse line and aimed to ablate the GRM2-expressing cells in the intralaminar thalamus using a viral-mediated caspase strategy (**figure 5.4a**). Slice histology revealed many of the GRM2 positive cells in CM, CL and PF were lesioned (**figure 5.4b**). Compared to control animals of the same mouse line (injected with AAV-flexed-GFP), we found lesions to GRM2 prevented animals from reaching the full task (**figure 5.4c**). All control mice reached the final level while only one lesioned animal was as able to reach this criterion within 4000 trials. Shuffle

analysis reveals that the lesioned cohort learned the early stages of the task slowly but mostly within normal range and learning only diverged after around 2500 trials at which point lesioned mice plateau (**figure 5.4d**). This divergence corresponds with training levels where port lights are switched off entirely (**figure 4.2b**). After these light guided levels, animals must increasingly rely on memory guided strategies and so the deficit observed in lesioned animals suggests they had issues in forming a memory guided strategy to solve the motor sequence. This is consistent with the hypothesised role for this population in driving learned motor sequences and suggests that procedural memory formation is reliant on this population of cells. To better understand how ablating GRM2 cells in ILT impairs the ability of animals to learn and execute the task, performance for lesioned and control animals were compared for trials after learning (4000 to 5000 trials when group average performance had plateaued). Lesioned animals still learned which ports were task relevant (**figure 5.4e**) but made more port-to-port transitions errors than controls suggesting they had deficits in remembering the transition structure of the task (**figure 5.4f**). Interestingly, almost all this difference is explained by errors which occurred towards the end of the sequence (transitions between ports 3, 4 and 5) (**figure 5.4g**). Unlike the previous experiments discussed in this chapter, analysis of the different kinds of errors showed that lesioned mice did not have a specific increase in repeat errors, and hence, the deficit observed was generalised (**figure 5.4j**). Besides deficits in the transition structure of the task, poor task performance could also be explained by motor deficits or reduced ability to learn to stereotype movements. Lesioned animals completed port-to-port transitions in comparable intervals to control mice suggesting there were no gross motor deficits (**figure 5.4h**) However, trial-to-trial variability for these time intervals was higher for the lesioned cohort (**figure 5.4i**) indicating lesioned mice may be less stereotyped in their movements. To confirm if this was true, tracking analysis was completed on a subset of control and lesion animals that were filmed during their final behavioural session (at the end of all training). While

control animals had highly stereotyped overlapping transition trajectories, mice with lesions to the GRM2 population in ILT failed to produce stereotyped movements and showed far more trail-to-trial movement variability (**figure 5.4k**).

To summarise, motivated by experimental constraints I have identified a genetic marker, GRM2, with potential utility as a tool for manipulating functionally relevant neurons. I have tested whether this marker is suitable and found that lesions to this population impaired the ability of mice to learn to generate stereotyped reproducible movements, a key hallmark of procedural skills.

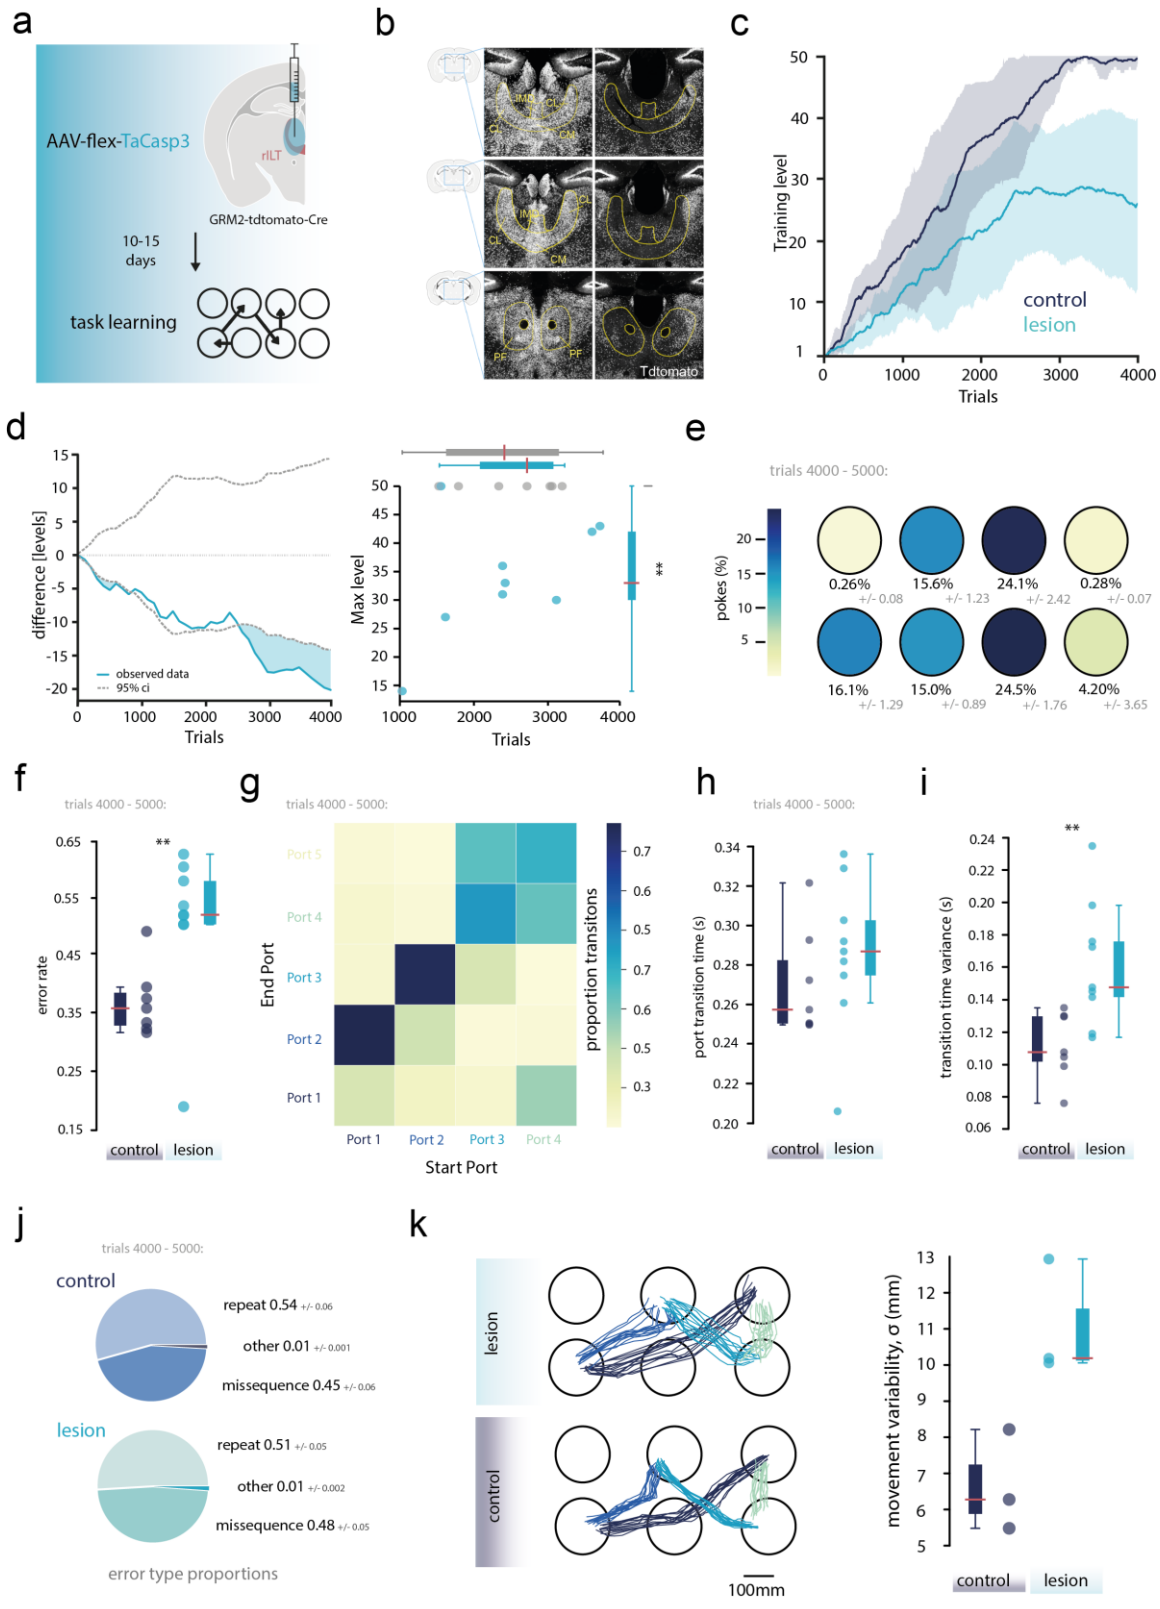

Figure 5.4: Ablation of GRM2, a genetically identified population in rILT, impairs motor sequence learning

Legend on next page

Figure 5.4: Ablation of GRM2, a genetically identified population in rILT, impairs motor sequence learning

(a) Schematic showing experimental design, including virus strategy and injection location**. (b)** Example histology showing GRM2 positive cells (tdtomato) in a non-injected baseline animal (left) and a lesioned mouse (right) **(c)** Learning curves for control and lesioned animals (shaded area denotes standard deviation, lesion group, n = 9, control group n = 7) **(d)** Left: differences in performance between the groups. Dotted lines indicate the 95% confidence interval for the shuffled data (see methods). Right: maximum training level obtained within 4000 trials against trials taken to reach maximum (p = 0.001, independent t-test). **(e)** Mean port poke proportions for all lesioned animals (trials 4000-5000). **(f)** Port transition heatmap, averaged for all lesion mice (trials 4000-5000). **(g)** Error rate as proportion of port-to-port transitions (trials 4000-5000). **(h)** mean port-to-port transition time (p = 0.4 independent t-test) trials 4000-5000). **(i)** Mean transition time variance (standard deviation, p = 0.008 independent t-test, trials 4000-5000). **(j)** Error type breakdown as a proportion of total errors made (reported error is SEM, trials 4000-5000). **(k)** Left: example tracking traces for a control and lesion animal (15 traces shown per subsequence transition were selected randomly from traces within the interquartile range of standard deviations to the mean tracking trajectory). Right: Mean standard deviations from average tracking line (n = 3 animals per group)

5.5 Discussion

5.5.1 A role for rILT in action sequence chaining

If the DLS is key for formation and recall of procedural memories, an important question is where the inputs, necessary for driving activity in this region, originate. I have suggested that striatal inputs from the rostral portion of intralaminar thalamus may be key in this respect. Further, I have hypothesized that neurons in this region may form part of a motor feedback loop, necessary for routing motor efference information from the midbrain and brainstem back into the striatum, facilitating action-to-action chaining.

Motivated by this hypothesis, I have first provided evidence that a thalamic region is required for motor sequence generation. I found that broad pharmacological inactivation of midline thalamus led to impaired task execution in expert animals through specific increase in one type of mistake: repeat poke errors. Since animals still had memory of the spatial locations of the correct ports as well as memory for how to engage with the task (poke into ports), it seems this deficit was specific to their ability to correctly transition between ports. One possible explanation for this result is that rILT is required for chaining motor sequence elements together. Hence, silencing this region prevents normal action-to-action haining leading to curtailed behaviour. However, claims from this result are limited. Though muscimol was targeted to rILT, it is inappropriate to claim that this effect was caused by specific silencing of a thalamostrital pathway – or even a specific thalamic nucleus – since infusions are unspecific and likely to have silenced a large area of thalamus. Using a more targeted optogenetic approach I found that transient inhibition of DLS projecting cells in rILT also caused a significant impairment to task performance. Strikingly, just as in the muscimol experiments, this impairment was largely driven by an increase in repeat poke errors. Once more, a hypothesis motivated explanation for this highly specific deficit is that silencing striatal projecting thalamic neurons breaks the thalamo-striatal portion of the efference feedback chain. Hence, the next action in the

sequence cannot be triggered and animals instead got stuck at the trigger port. That animals in the muscimol experiments could still perform an 'initiation' action (poking into port) is consistent with this hypothesis since the proposed action chaining architecture would still require an elsewhere generated initial movement as an ignition: triggering a domino effect of linked actions in the chain.

These results also reproduce findings from previous work looking at the rodent brain: Wolf *et al*. (2022) showed that silencing striatum projecting cells in intralaminar thalamus abolished stereotyped expression of a learned motor skill. While my results are broadly consistent with these findings it is worth noting that in the silencing experiments of Wolf and colleagues, their rats were still able to compete the task (2 lever presses) – albeit poorly. These findings somewhat contrast observations from my inhibition experiments where the higher-level structure of behaviour was altered. A possible explanation for this discrepancy is that the two lever presses required in the lever press task was too simple a sequence to observe any high order structural or syntactic errors. A comparable task would be two repeated nose pokes; a behaviour which mice in my experiments did often express during inhibition. Furthermore, while my experiments involved acute, short term inactivations of thalamus, Wolff and colleagues performed chronic silencing. It's unclear how each of these experimental strategies affect circuits involved. Certainly though, differences are highly likely so comparisons between these results should be framed in this context. A final distinction, besides possible differences between mice and rats, is that the work of Wolf and colleagues targeted the entire ILT and not just the rostral portion as in my experiments. It is possible that this could also underly the small differences between my observations and their findings.

My findings are also in line with other work in the field. I have shown evidence that a thalamic nucleus (rILT) has a role in action-to-action chaining, and I am proposing that this thalamic pathway functions by routing motor efference

information back to the DLS. Certainly, a similar role has previously been shown for the thalamus. For example in primates, efference copy or 'corollary distance' has been shown to be vital for normal saccadic eye movements. In particular for a sequence of saccadic movements or for self-generated eye movements, silencing thalamus via muscimol injections, causes delayed and inaccurate saccades (Guthrie *et al.*1983; Sommer and Wutz 2002, 2004; Tanaka, 2006). Hence, both the timing and kinematic specifics of future eye movements has been shown to depend to some extent on triggering by internal feedback about previous action. In my experiments I also find that inhibiting a thalamic region caused inaccurate and delayed movements. Therefore, it is possible that a similar mechanism – one based on thalamic efference copy feedback – controls the behaviour in my task. Moreover, the specific thalamic regions targeted in these saccade studies are close to, if not lying within rILT. Hence, these experiments may even provide direct evidence that rILT is essential for controlling behaviours which rely on efference copy information.

Another possibility is that rILT is not unique in this function and that a motor feedback role is generalised across multiple thalamic nuclei. The specific function of these individual nuclei likely relates to their connectivity. For instance, in the works discussed (Guthrie *et al.*1983; Sommer and Wutz 2002, 2004; Tanaka, 2006) medial-dorsal thalamus is proposed to have a role in controlling saccades due to its position connecting the ocular portion of superior colliculus (saccadic motor output control) to the frontal eye fields (motor planning). rILT also receives ascending motor input. These inputs come from a range of midbrain, brainstem, and cerebellar motor controllers (Yamasaki *et al.* 1986; Cornwall & Philipson 1988; Krout et al. 2002; Bostan & Strick 2018; Mandelbaum *et al* 2019). Since in this chapter I have provided evidence that rILT is involved in chaining fully body postural movements then my hypothesis predicts that the ascending inputs to rILT should provide efference copy information about these movements. Given this hypothesis, it should also be noted that rILT is not the only striatal projecting nucleus.

Though I have outlined that rILT may be best suited for this role, the importance of other striatal projecting thalamic nuclei (eg. PF, VT) for procedural behaviours may depend on the motor specifics of the behaviour being learned. Certainly, these nuclei have been implicated in sequential lever pressing (Dias-Hernandez *et al*. 2018; Hidalgo-Balbuena *et al*. 2019). Speculatively these other nuclei may carry more information for lever pressing actions – eg. hindlimb movements. This specific hypothesis could be tested by lesioning or silencing striatum projecting thalamic nuclei in turn and testing ability to learn different kinds of movement skills.

My experiments are also reminiscent of previous work in the songbird. For example, thalamic feedback from downstream motor regions in the song learning circuit has been shown to be necessary for normal song learning and production (Williams & Vicario 1993, Coleman & Vu 2005; Danish *et al*. 2017) Strikingly, it has recently been shown that that inhibition of thalamic projections in this circuit causes impaired chaining of sequential task movements (Moll *et al*. 2023). Inhibition only perturbed chaining when done at syllable boundaries and it was shown that thalamic inputs triggered the start of each syllable by activating starter cells (first neurons in a domino-like chain of neurons) in HVC. This suggests that in the songbird vocal circuit, thalamic inputs define the sequential or compositional structure of behaviour. In my experiments I found a remarkably similar effect when inhibiting thalamic inputs to striatum: the sequential order of behaviour was perturbed. This suggests that an equivalent mechanism (carried by rILT inputs to striatum) could underpin mammalian motor sequencing. This is an interesting avenue for future research. A future approach could extend the optogenetic silencing method I employed in this chapter by inhibiting thalamus at different times throughout task sequence progression. If, like in the songbird system, the thalamus is required for triggering behavioural syllables, then inhibition should only impair behaviour at boundaries between these syllables. A second approach could be to record from these rILT neurons during the sequence behaviour. In a sequential task, if thalamic

neurons define syllable boundaries, then we might expect that these cells should not be continuously active, but fire in timed intervals aligned to behavioural transitions. Recordings could also be used to investigate whether the role of the thalamus is to provide efference copy feedback. For this It would be necessary to identify striatal projecting cells and neurons which receives ascending inputs from motor regions. This could be achieved with combined antidromic and orthodromic opto-tagging approach. For orthodromic tagging, spikes triggered at upstream cell bodies could be recorded in recipient thalamic cell bodies to confirm input connectivity. For retrograde tagging, stimulation of DLS could produce backwards propagating spikes which, if recorded in upstream thalamic cell bodies would confirm connectivity (Lim *et al.* 2013). Having determined cells with this connectivity, according to the efference copy hypothesis, we might expect that they should specifically carry kinematic information about current ongoing movements.

Though my data are in line with previous work, conclusions drawn from my experiments should also be viewed in context of several experimental confounds. Firstly, JAWS was also expressed by cells in the caudal nucleus of ILT, the PF. Though fibreoptic ferules were implanted over rILT, it is possible that light which spread caudally, inhibiting the PF, could have been responsible for the effect seen. Further, though striatum projection cells were specifically targeted, it's unclear whether these neurons also send collaterals to other regions (for example to cortex). It cannot therefore be ruled out as a possibility that the behavioural effects of inhibition were mediated by suppression of a non-striatal pathway. Finally, though significant, the repeat poke inducing effect of inhibition was neither consistent across trials nor particularly strong compared to control mice. Notably the most dramatic effect of laser inhibition was to make the first transition (correct or incorrect) slower. This slowing effect was seen for both correct and repeat error trials individually. Although the sequence chain was not broken for correct trials, a possible explanation for this slowing effect is that inhibition was too weak to

fully arrest the behaviour, but strong enough to reduce thalamic drive to the striatum; sufficient to cause the next sequential movement to be less invigorated. If this is true, then it suggests that neural inhibition via the retrograde optogenetic approach used here was not strong enough, or was not able to influence a large enough proportion of the functional population to consistently disrupt activity in this circuit.

To more convincingly show that striatum projecting cells in rILT have a role in in action chaining, in future work these experiments must be taken further. Firstly, since I inhibited at the first port of the motor sequence, our results are arguably compatible with a role for thalamus in simply initiating motor sequences, rather than supporting chaining of the full behaviour. This seems unlikely if we recognise that this behaviour is circular: the point of movement initiation is more accurately described as at the reward port (port 5) as this is the only time mice pause (to drink water). Nevertheless, if our hypothesis is true, inhibition should break the movement chain at any point of the behaviour and so showing this is indeed true would provide far more compelling evidence. Secondly, future work must aim to distinguish between thalamo-striatal and other projection pathways from ILT. A potential way of doing this would be to target thalamic terminals locally in the striatum. This could be done with optogenetics by expressing an opsin in ILT neuron and implanting fiberoptic ferules into the DLS. Alternately, a pharmacogenetic strategy could be employed wherein a designer receptor such as hM4Di would be expressed in ILT neurons and then a ligand infused into the DLS via implanted cannulas (Zhu & Roth 2015).

5.5.2 A genetic marker, GRM2, provides useful leverage over a thalamic population which is involved in procedural memory formation

The experimental limitations which I have laid out motivated me to search for a genetic marker in ILT which could give better leverage for future experimental manipulations. In the final part of this chapter, I identified the gene GRM2 (which is highly expressed in rILT) as a candidate for this marker.

For this gene marker to be useful, then this population should have a functional role in procedural memory formation. Indeed, I find lesion to this population in rILT impairs task learning to the extent that mice were unable to reach the final, fully memory guided, version of the task. Though impaired, lesioned mice still engaged with the task well, learned the locations of task relevant ports and showed no signs of having any gross movement deficits. This suggests these animals specifically had impairments forming normal procedural memory for the task. Analysis of task expression after extensive training revealed that, while animals performed the first half of the task without issue, they had problems generating correct port-to-port transitions in the latter half of the task sequence. There are two explanations for this outcome. Firstly, this could be caused by degraded capacity to form, and store extended procedural memories. This is in line with the proposed role for these cells in the efference chaining loop wherein the rILT would be a key node in a circuit that chains actions together. Loss of function to this chaining circuit (and therefore reliance on compensatory circuits) could cause animals to have an impaired ability to learn elongated motor sequence chains. A second explanation, reminiscent of previous findings (Wolff *et al*. 2022), is that lesioned animals had impaired ability to learn how to generate stereotyped actions. If this is true, then this impairment must have been particularly catastrophic to the second half of the sequence. If these latter port-to-port transitions were already harder to complete and required more precise movements to properly poke into each port (and break the Infrared beam), then they would suffer more from increased movement variability. This explanation is also in line with the proposed efference chaining model, though if true may suggest this circuit is more important for specifying the specific low-level kinematics of movement rather than higher level structural ordering. Although it is also feasible that this ordering is still contingent on this circuit, but simply more robust to a partial lesion of the thalamic population.

These two possible explanations are made more complex by the fact that the automated training protocol guides animals with lights up to the final task. Hence, for the trials analysed, control mice were completing the full task while lesioned mice were at lower, easier, levels and therefore had light guidance for some ports. Guidance for the latter movements in the task sequence persist until higher levels than for the early parts of the sequence (see figure 4.2). This is implemented because during development of the task these transitions tended to be the least well performed by animals. Hence, the training protocol is designed specifically to support learning of these movements. Since these port-transition movements are the most difficult for non-lesioned animals to complete, it stands to reason that they might be most impaired by a lesion. Taking task guidance into account, the notion that animals had problems forming stereotyped movements is perhaps more likely: If animals were receiving guidance, then issues remembering which port to go to next should have been negated. Additionally, movement tracking analysis showed that lesioned animals were indeed unable to perform consistent movements between ports. In fact, lesioned mice had movement variability that was similar to that of non-lesioned animals during early learning. Notably, this result is very similar to previously published work, where chronic silencing of thalalmo-striatal neurons in ILT has been shown to impair the ability of rats to convert naïve lever presses into skilled movements (Wolff *et al.* 2022). Overall, there are two conclusions which can be drawn from the results I have discussed:

(1) Firstly, GRM2 appears to be a suitable genetic marker which targets cells in rILT that are functionally relevant for procedural memory formation. It should be noted that this is not evidence that GRM2 cells perform this function alone in rILT. Certainly though, cells with this genetic marker represent a large enough proportion of the functional population to be an interesting target group for future work.

(2) Secondly, this experiment provides more evidence that neurons in rILT are crucial for procedural memory formation. It is thought that in certain

behavioural contexts early (phase 1) procedural learning is contingent on the motor cortex, but a thalamo-striatal pathway takes over (phase 2) and drives movements to become more stereotyped (Kawai et al. 2015; Wolff *et al*. 2022). Since lesioned animals in my experiments could not convert their naïve, highly variable movements into precise stereotyped movements, an explanation is that the required thalamo-striatal pathway was impaired, and so procedural learning was stuck at phase 1.

While compelling, these conclusions are limited. Firstly, before any useful conclusions can be drawn, I must first characterise the GRM2 lesions in detail. GRM2 is not exclusively expressed in ILT but also has sparse expression in surrounding thalamic nuclei (**appendix figure3a**). Hence, to claim that a specific thalamic region is responsible, detailed characterisation of the lesions will be essential. The mouse line used was GRM2-Cre crossed with Ai75(flexed-tdtomato). This means GRM2 positive cells can be identified by the presence of this fluorescent protein in cell soma. These brains will be imaged with a serial section microscope and lesion sizes estimated by counting red cells in this region and comparing to baseline animals. Conclusions from the lesion experiments are also limited by the fact that no anatomical tracing of the GRM2 population has yet been done. Though it is high likely that this genetic populations do send inputs to striatum – due to their presence in ILT (Parent & Parent, 2005), careful anatomical tracing is needed to confirm this. Further, these cells may also send projections to other regions. Besides quantifying this with anatomical tracing, if this is the case, then the experiments suggested previously (based on terminal inactivation in striatum) will still be necessary to determine whether thalamo-striatal projections alone are responsible for procedural memory formation. From the experiments I have detailed here it is also not possible to conclude whether the GRM2 cell line is required for more than simply learning the task. Future, work should also aim to perturb these cells, or a striatal projecting subpopulation, during task production. Finally, I have proposed that a population in rILT forms part of a feedback loop, routing motor efference

information into the striatum. To support this hypothesis, future work must also investigate a role for inputs to thalamus from brainstem and midbrain motor regions. Preliminary rabies tracing has shown that cells in rILT do receive extensive inputs from motor nuclei, for example the superior colliculus (**appendix fig3b**). This region is a motor structure that receives direct modulation from the basal ganglia and is also known to be essential for orienting movements (Wheatcroft *et al*. 2022). Since, the trained motor sequence in my task is a series of orienting postural movements, it is a highly promising candidate region in our proposed motor loop. Testing whether perturbation of ascending projections from motor regions such as the superior colliculus impairs task performance in trained mice will be an essential experiment in pursuit of the circuit mechanism underlying motor sequence behaviours.

5.5.3 Summary

In summary, in this chapter I aimed to explore the role of thalamic inputs to striatum in both learning and recall of procedural memory. In doing so I ultimately aimed to test whether procedural behaviours could be supported by a motor feedback loop, with a thalamic population (rILT) forming a key intermediate node in this loop. I have shown evidence that procedural memory for my task is indeed contingent on striatal projecting cells in rILT. Further, I have argued that the behavioural deficits observed when this pathway is inhibited are consistent with a loss of the ability to chain actions together to form motor sequence behaviours. Finally, I have recognised that the questions posed in this chapter require careful and potent manipulation of the circuits involved. To address this, I have identified a genetic marker which distinguishes a population in rILT which is required for formation of stereotyped procedural memories. This marker has enormous experimental utility and I have described how it can be leveraged in pursuit of testing the guiding hypothesis for this chapter. Therefore, while no complete conclusion

can be drawn on the proposed hypothesis, the work in this chapter lays the foundations for future work which goes beyond the reach of this thesis.

This chapter, and the proceeding chapter have focused on the circuits which support awake learning and recall of procedural memories. In the next chapters, I aim to focus on processes that support the function of these circuits offline. As discussed, motor skill memory is contingent on consolidation during offline periods. Most notably during sleep. Hence, any investigation of the circuit mechanisms underlying procedural memory formation is incomplete without also examining the contribution of these circuits offline.

CHAPTER 6

6. SEARCHING FOR THE OFFLINE MECHANISM WHICH SUPPORTS PROCEDURAL MEMORY CONSOLIDATION

6.1 Introduction

So far, in chapters 4 and 5, this thesis has explored of the role of subcortical circuits involving the thalamus and striatum in procedural memory. I have provided evidence that these circuits are key for learning and executing procedural memories in awake animals. While procedural memory formation is contingent on awake experience, it is important to recognise that memory is also supported by offline processes. These processes, particularly during sleep, have fundamental influence over awake behaviours. Offline mechanisms have been show to both consolidate memory gains from previous awake practise, and even reshape and improve procedural abilities beyond levels previously observed during awake behaviour (Diekelmann & Born 2010; Rasch & Born 2013; Schmid *et al*. 2020). To gain a comprehensive understanding of the circuit mechanisms that underlie both the formation and execution of procedural memories, it is therefore essential to consider the contribution of these offline processes. At this stage, we have already established a strong foundation for investigating offline mechanisms for the formation of procedural memory. We possess a well-defined procedural memory task in which mice learn to generate a motor sequence from memory, and we have compelling evidence that this memory is both formed and stored in the DLS.

6.2 Blocking offline plasticity in the DLS impairs early learning and late stabilisation of procedural memory

If the DLS is the site of learning and long-term procedural memory storage, then this region should be involved in the offline process that acts to support procedural memory formation. To first test if this is true, we aimed to block offline plasticity in this region during consolidation. NMDA receptor activation

has been associated with skill learning both online and offline (Jin and Costa, 2015; Santos *et al.*, 2015; Lemke *et al.*, 2021) and so blocking these receptors should acutely impair normal synaptic plasticity mechanisms. Accordingly, I infused 2-amino-5-phosphonopentanoate (D-AP5), an NMDA receptor antagonist, into the dorsolateral striatum (**figure 6.2a,b**). Procedural offline consolidation is believed to be most effective and critical immediately after motor skill practise (Holz *et al.* 2012). Hence, I targeted this critical period by performing infusions immediately after mice had been trained on the task. AP5 and saline control infusions were performed alternately during training, and infusions were done intermittently; on training days that were applicable (see methods) (**figure 6.2c**). After infusion, mice were placed back into their home cage and training performance (change in training levels) was assessed in the next training session (24 hours later). Across all infusion sessions (for all animals) linear regression indicated a significant difference in the effects of infusion on the level change in the test session ($p>[t]$ intercept = 0.033, F-statistic = 3.464, Ordinary Least Squares regression) (**figure 6.2d**). Compared to saline infusions, AP5 more often impaired performance the next day. This was most pronounced for infusion in early training and no longer true in late training. To better quantify this interaction between training stage and AP5-induced learning deficits, changes in levels between infusions for early, middle, and late learning were compared. Early levels were defined as those which are entirely light guided (levels 13-20), middle stage levels were those from which one of the four port guidance lights began to dim up to when the second of these four lights began to shut off (level 21-35). Late levels were from this point up to the penultimate level before the final task (levels 36-49). Splitting training in this way we find that AP5 caused significant task deficits the next day when infused post training during early learning. In test sessions which resulted in an overall decrease in levels, the minimum level achieved was used to calculate the number of training days the performance of the animal had regressed (see methods). In all sessions with a decline, the

performance of animals regressed 1 day on average (**figure 6.2f**). This suggests that offline NMDA-dependent plasticity is a critical component of the mechanism which consolidates the improvements in performance gained the previous day.

Since the effect of AP5 on consolidation disappeared as the animals became better at the task, this could indicate that an offline mechanism only supports procedural consolidation in early learning. However, another explanation is that animals simply became more robust to perturbation as they improved at the task such that a single post session AP5 infusion was not sufficient to see impaired performance the next day. To test if this was the case, we took fully trained mice and carried out infusions of saline and AP5 for 4 consecutive days (**figure 6.2g**). Blocking post-task offline plasticity for 4 consecutive days led to a significant drop in levels (and hence, performance) as can be seen in the cumulative level change across the 4 test sessions compared to saline infusion (**figure 6.2h**). While conclusions from this experiment are partially limited (see section 6.8), together these data suggest that offline processes in the dorsolateral striatum are critical for both learning and maintaining procedural memory of a stereotyped action sequence.

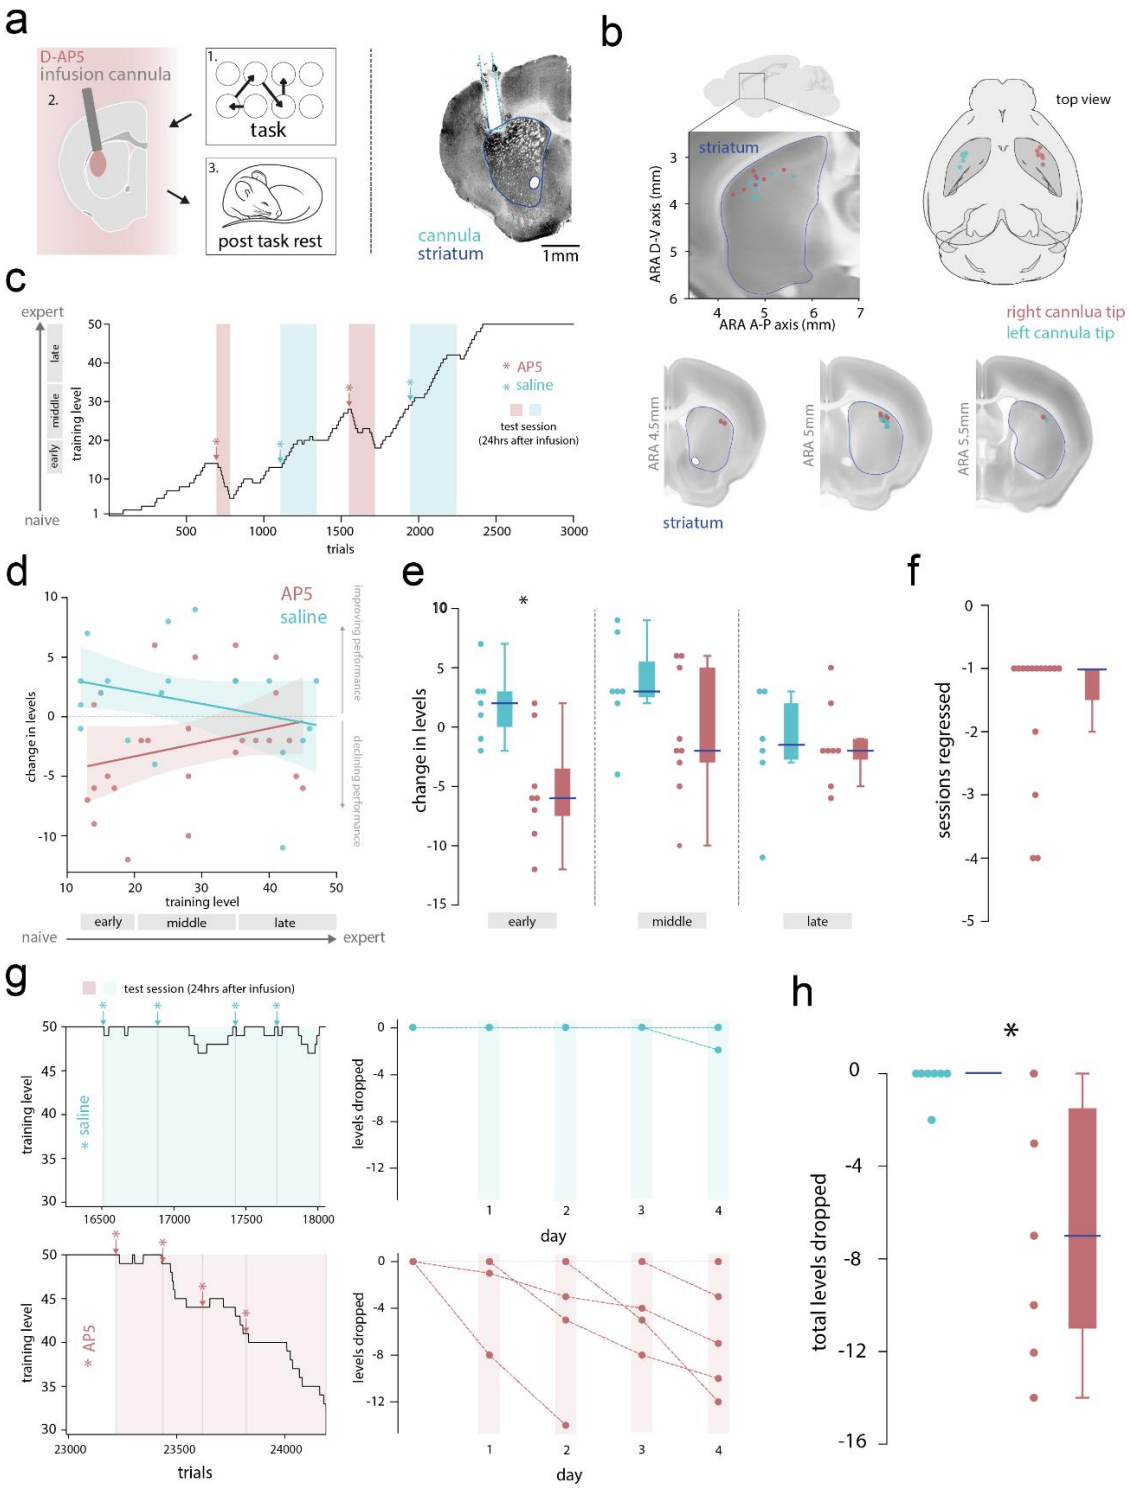

Figure 6.2: Blocking offline plasticity in the DLS impairs early learning and late stabilisation of procedural memory

Legend on next page

Figure 6.2: Blocking offline plasticity in the DLS impairs early learning and late stabilisation of procedural memory

(a) Left: schematic showing experimental design including cannula placement. Mice performed the task (1), drug or saline was infused bilaterally into the DLS (2) then animals were allowed to rest overnight (3), then tested the next day. Right: example slice histology showing cannula tract (light blue) and striatum (dark blue). **(b)** Left and right hemisphere cannula tip positions in Allen reference atlas coordinates. Sagital view (top left), top-down view (top right) and coronal sections (bottom). **(c)** Example animal learning curve showing AP5 and saline infusions (stars) and test sessions 24 hours after infusion (shaded) (n = 8 mice). **(d)** Training level change against start level for all post infusion test sessions (saline and AP5) for all animals. Lines show linear fit for each dataset with confidence interval (shaded region) **(e)** training level change for each infusion experiment grouped by initial training level for each test session. (p = 0.00448, independent t-test). **(f)** Number of sessions since minimum level was last seen for AP5 infusions experiments that resulted in negative trial changes**. (g)** Left: example training level progression for fully trained animals given consecutive infusions (infusions points marked by stars, sessions are marked by grey vertical lines, shaded regions are test sessions – 24hours post infusion). Right: cumulative levels dropped across the four consecutive infusion days for each animal. **(h)** Summary plot showing total levels lost for each infusion type (AP5 & saline) after 4 days of consecutive post session infusion (p = 0.024, Mann-Whitney U, n = 6 mice).

6.3 Adaptation of an existing point process model, PPseq, as a novel method for unsupervised, biomarker free replay detection

In line with previous work (Lemke *et al.*, 2021), I have so far shown evidence that offline processes, acting in the DLS during post learning rest periods, are essential for normal learning and maintenance of procedural memory. If an offline mechanism supports procedural memory formation in the DLS, what could the substrate of this process be? For episodic memories, neural replay of previously observed activity is thought to be the mechanism by which

memory is consolidated offline (Wilson & McNaughton 1994; Nadasy *et al*. 1999; Olafsdottir *et al*. 2018). These sequential reactivations are believed to strengthen neural engrams for experience by acting as a kind of offline simulation; allowing plasticity mechanisms to engrave experience into memory. Whether a similar mechanism – sequential replay in motor skill circuits – exists for procedural memory consolidation is not known. To search for this mechanism, I recorded from neural activity in DLS and motor cortex via chronic implantation of neuropixel probes in trained animals (**figure 6.3a&b**). However, to examine these recordings for replay I needed a replay detection method.

Episodic replay is thought to occur most frequently during short high frequency LFP oscillations named sharp wave ripples (SWR) (Roumis and Frank 2015; Buzsaki 2015) In previous studies, episodic replay has been identified by first identifying SWR events as candidate epochs in which to search for replay. Using these LFP biomarkers as a heuristic, replay is usually recognised and quantified by using linear Bayesian decoders trained on awake activity to search for reactivations during these epochs. This approach yields a set of probability distributions for the time bins associated with each ripple which are commonly assessed as either replay or noise by fitting a line to this data and testing whether the fit of this line is statistically significant compared to shuffled data (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009; Carr et al., 2012; Ólafsdóttir et al., 2018; Tang et al., 2017; Shin et al., 2019; Tingley and Buzsáki, 2020; Bhattarai et al., 2020). This approach relies on multiple assumptions about the nature of replay. Firstly, only linearly propagating replay can be detected – no stationary, highly fragmented replay, or replay that changes of direction are possible. Secondly, only replay that propagates at a constant speed can be detected. Finally, only replay that occurs during predefined biomarker epochs can be identified. Since striatal replay (neural progressions of task related activity) has not previously been described these assumptions are particularly problematic.

Firstly, unlike hippocampal place cell coding, single cell activity in the striatum is not so well understood. Even short bouts of offline hippocampal activity can be tied to external variables (place) and recognised more confidently as replay because of this fact. Activity in the DLS appears to correlate best with features of movement (Barbera et al. 2016; Markowitz et al. 2018; Jin et al. 2014), however, this coordinate framework is much harder to understand intuitively at the level of single spikes. Secondly, there is no prior understanding of what offline striatal dynamics may look like. Finally, and importantly, there is no known LFP biomarker for when to search for procedural replay. Simply, I don't know exactly what to look for, I don't know when to look for it, and it may be hard to recognise if I do find it.

Classical linear decoding approaches are not appropriate when considering these three issues and so other approaches must be considered. A recently developed method based on combining decoding with a state space model (Denovellis *et al.* 2021) has overcome some of the limiting assumptions. This approach, when applied to hippocampal data, identified far more replay events with greater diversity, further suggesting linear methods may constrain the ability to properly characterise replay. While a potential improvement, this method still relies on a biomarker to first identify candidate epochs. An ideal approach to satisfy my aims of characterising offline dynamics in the striatum would be unsupervised. This type of approach would allow search for replay without making assumptions that would bias towards certain replay types or certain brain states. Moreover, such a method must not use any biomarker heuristic, and should identify reactivated activity based purely on features of spiking content – irrespective of inferred position or speed. To satisfy these goals, we adapted an existing unsupervised method based on a point process model called PPseq (Williams *et al.* 2020).

PPseq aims to describe spiking events by the underlying latent structure of neural activity. Via an iterative Poisson sampling process, the model fits free parameters to determine descriptive latents and then assigns spikes to that structure based on features of activity. In simple terms, the model takes in raw spikes, fits descriptive latent events to these spikes, then labels spikes which contribute to each latent (**figure 6.3c**). The features of activity which determine this assignment are the offset of the peak of smoothed spiking activity to the median of the identified latent event, as well as the amplitude and width of smoothed spiking activity. Overall, PPseq takes raw spiking data, and then returns the same spikes, ordered and labelled as neural sequences defined by underlying neural structure. Hence, my use of this model is as a kind of pattern detector for neural data. We aimed to use this approach to first characterise patterns of repeating sequential neural activity in awake task related activity. Then I aim to use these patterns as templates to search for replay of this activity offline. To apply this method, I first filter spiking data to remove cells with highly regular and sparse firing rates (Fano factor 0.5 -12) for a 600s period of high task engagement. I then fit the PPseq model to this filtered period of task related activity. Next, we made a small adjustment to the original description of PPseq such that the free parameters of the model could be fixed. With this adjustment, I then apply this same fixed model (fitted on awake activity) to spikes from the same neurons recorded during post task sleep (**figure 6.3d**). If the same latent structures are detected recapitulated offline, PPseq will label the spikes which contribute to these structures and hence, identify sequential replay of task related activity. A final consideration is that episodic replay is often altered compared to awake activity: replay is often time compressed but can also be stretched and even reversed (Lee & Wilson 2002; Diba & Buzsaki 2007; Davidson *et al.* 2009). When applying PPseq to offline periods, a range of compressing and stretching time warping parameters are implemented including negative warps (reversed dynamics).

Hence, PPseq is also a suitable method to account for this possibility in striatal replay.

In sum, I aim to use an adapted version of PPSeq, an unsupervised method based on a point process model, to search for replay. As this method is unsupervised and requires no prior identification of a biomarker, it can be applied in an unbiased way to large epochs of offline spiking. Since little is known about striatal offline dynamics this novel approach appears highly suitable for searching for striatal replay.

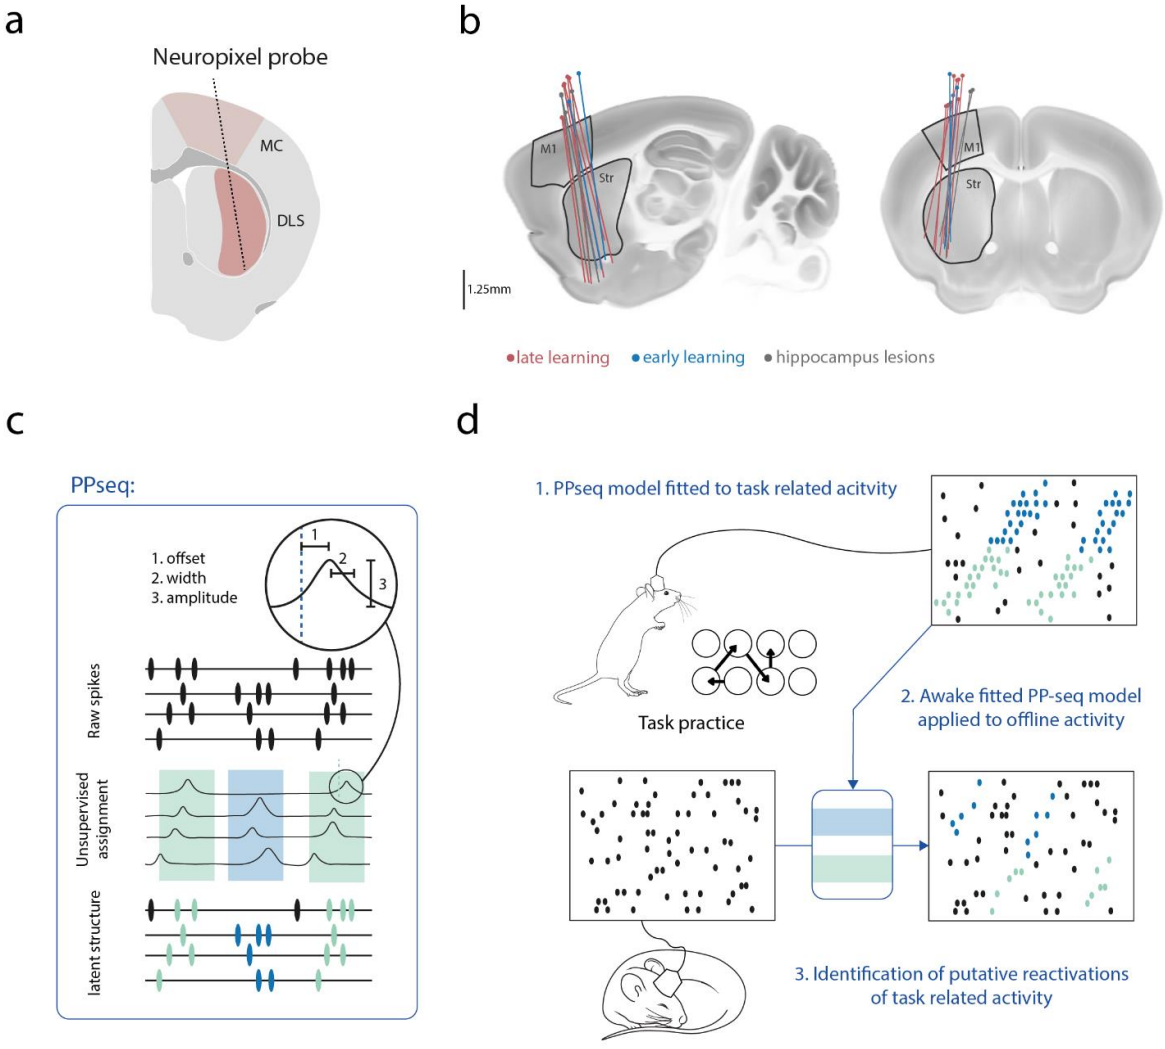

Figure 6.3: Adaptation of PPseq as a novel method for unsupervised, biomarker free replay detection.

Legend on next page

Figure 6.3: Adaptation of PPseq as a novel method for unsupervised, biomarker free replay detection.
(a) Schematic diagram showing recording method and intended probe placement for chronic recordings. **(b)** Traced implanted neuropixel probe locations projected onto standard Allen atlas sagittal (left) and coronal (right). The three recording groups (see chapter 7) are shown as 3 colours. **(c)** Schematic diagram outlining unsupervised detection and labelling of latent neural structure from raw spikes by PPseq. **(d)** Schematic diagram of proposed replay detection protocol. Top: PPseq model is trained on and detects repeating neural sequences within task related activity (1). Bottom: Recordings extended after training to include post task sleep, PPseq model fitted on awake task activity is applied to sleep recording (2). Recapitulated task related neural sequences are identified by PPseq during offline activity (3).

6.4 PPseq model hyperparameter selection

Before applying a PPseq model to my recordings, it is essential to first understand how model features may affect how PPseq describes spikes data and select model hyperparameters in a principled way. Running PPSeq required setting 12 hyperparameters; 1 corresponding to the number of sequence or motif types in the model, the other 11 controlling the various priors on variables within the model. Due to the unrealistic compute time required to search across all 12 parameters, instead, we made use of the fact that many of the hyperparameters are very interpretable, such that we can make reasonable guesses for their values a priori. Hence, we fixed the values of 7 of the hyperparameters for which we have reasonable guesses, as detailed below, and performed a grid search through the remaining 5 that were likely to have the largest effects on the results of the model (**see table 6.4**).

Hyperparameter	Meaning	Value	Chosen how
n	Number of seq/motif types	6	Grid search
ψ	Rate of events	$1 \cdot 0.2 \cdot n$	Grid search
γ	Event concentration	1	Fixed
μ_A	Mean event amplitude	$0.3 \cdot N \cdot P$	Grid search
σ^2_A	Event amplitude standard deviation	$10 \cdot \mu_A$	Fixed
Φ	Neuron response concentration	0.6	Grid search
ν	Neuron response width prior	0.5	Fixed
σ	Neuron response width scale	1	Fixed
κ	Neuron offset scale	0.5	Grid search
μ_0	Mean background spike rate	$0.3 \cdot N \cdot P$	Fixed
σ^2_0	Background spike rate variance	$(0.3 \cdot N \cdot P)^2$	Fixed
γ_0	Background firing rate concentration	0.3	Fixed

N = number of neurons, P = av. firing rate

Table 6.4: *PPseq hyperparameter details*

The hyperparameters can be split into 4 groups which control different aspects of the model. I will explain each of these in turn and the rationale behind each hyperparameter setting.

(1) Sequence characteristics:

Sequence events are drawn from a Poisson Process, occurring at an average rate ψ. There are n different types of sequence events, and γ controls the prior on the proportions of each of the sequence types. We swept through the number of sequence types in our grid search. We set the other two parameters as follows. γ sets the prior on sequence type proportions; if it is large all sequence types are likely to occur at the same rate, if it is small a few

sequence types likely occur much more often than others, and if it is 1 the prior distribution over proportions is uniform (**figure 6.4.1 a**). Since we do not want to bias the model one way or the other, we choose $\gamma = 1$, i.e. a uniform prior. ψ sets the sequence event rate. We expect that as the number of sequence types increases PPSeq will increasingly split spikes into smaller motifs. This rate value will likely be an important factor determining how PPseq labels the spiking data and so this value is one which we sweep across in the grid search. As we are interested in observing full structures and want to avoid overly splitting single neural sequences into multiple arbitrary chunks, we therefore scale this event rate as a starting point by multiplying our swept values by $0.2 \times n$.

(2) Spiking Amplitude:

In the generative model, when a sequence event occurs it triggers a number of spikes, called the amplitude of the event, and those spikes are then distributed between the neurons. We control the prior distribution on the amplitude through its mean, μ_A, and variance, σ^2_A, and the prior on the proportion of spikes each neuron receives through Φ. We scale the standard deviation and mean together, by setting $\sigma_A = 10 \times \mu_A$. We then specify μ_A up to a scale factor. We expect an event to trigger spiking over roughly 1 second, over some fraction of the neurons, f. We do not a priori know how many neurons will be involved in each sequence, and hence how many spikes we expect in each event, but likely this will play a large role in guiding PPSeq's choice of clustering. We therefore set $\mu_A = f \times N \times P$, where N is number of neurons and P is their average firing rate and sweep through possible values of f in the grid search. Φ, like γ is another concentration parameter. If it is high PPSeq will find sequences that include as many neurons as possible, if it is low, it will find sequences containing only a few neurons. This is clearly highly

influential over how PPseq will label the data and therefore we also sweep this parameter in the grid search.

(3) Firing Rate Modulation:

When a sequence event occurs, it causes a characteristic bump in the predicted firing rate of each neuron, governed by a neuron and sequence dependent pair of parameters, b_{ir} and c_{ir} (ν), where *i* denotes neuron index, and *r* a sequence type index. c_{ir} controls the width of the bump in firing rate, and b_{ir} controls the peak's offset from the event time (**figure 6.4.1 b**). The role of hyperparameters on the prior distributions of these variables are the most difficult to interpret, thanks partly to their coupling. We explored the parameter space and found that ν had a limited effect (**figure 6.4.1 c**) and so set it to 0.5 giving reasonable prior distributions for a range of values of σ^2 and κ. σ^2 causes large variations in the distribution, and hence leads to big changes in the priors on the length of sequence events in time. However, preliminary sweeps revealed that a relatively narrow range of sequence lengths is so strong in the data that changes to the sequence length priors have little effect. We thus conclude σ^2 can be fixed to a value that gives a reasonable distribution of spread lengths of 1, (**figure 6.4.1 c**). Finally, c_{ir} sets the variance of the prior on b_{ir}, up to a scaling factor κ. Hence κ leads to changes in scale between the distribution of b's and c's: this encourages sequences of different shapes (**figure 6.4.1 d**). To search through this range, we therefore included κ in our grid search.

(4) Background Firing Rate:

The final component of the model is the background firing – spikes that do not come from a sequence/motif event. These occur at a rate drawn from a prior, that we govern through its mean, μ_0 , and variance, σ^2_0 , and then are

distributed amongst the neurons. The prior on this distribution is another Dirichlet distribution on proportions controlled via a concentration parameter, γ_0, as in **figure 6.4.1 a**. We fix all of these parameters in the following way. Firstly, once again we set the standard deviation equal to the mean of the prior on background rates, to ensure the two scale together, $\sigma_0 = \mu_0$. We then roughly set the mean background firing rate to be 30% of the total, 0.3 × N × P. Finally, we expect a few neurons to be much more involved in the background spiking than others, but we do not want to bias the prior too much in this direction. Therefore, we set $\gamma_0 = 0.3$, i.e. a weak bias away from uniform towards splittings of the background rate that favour some neurons over others.

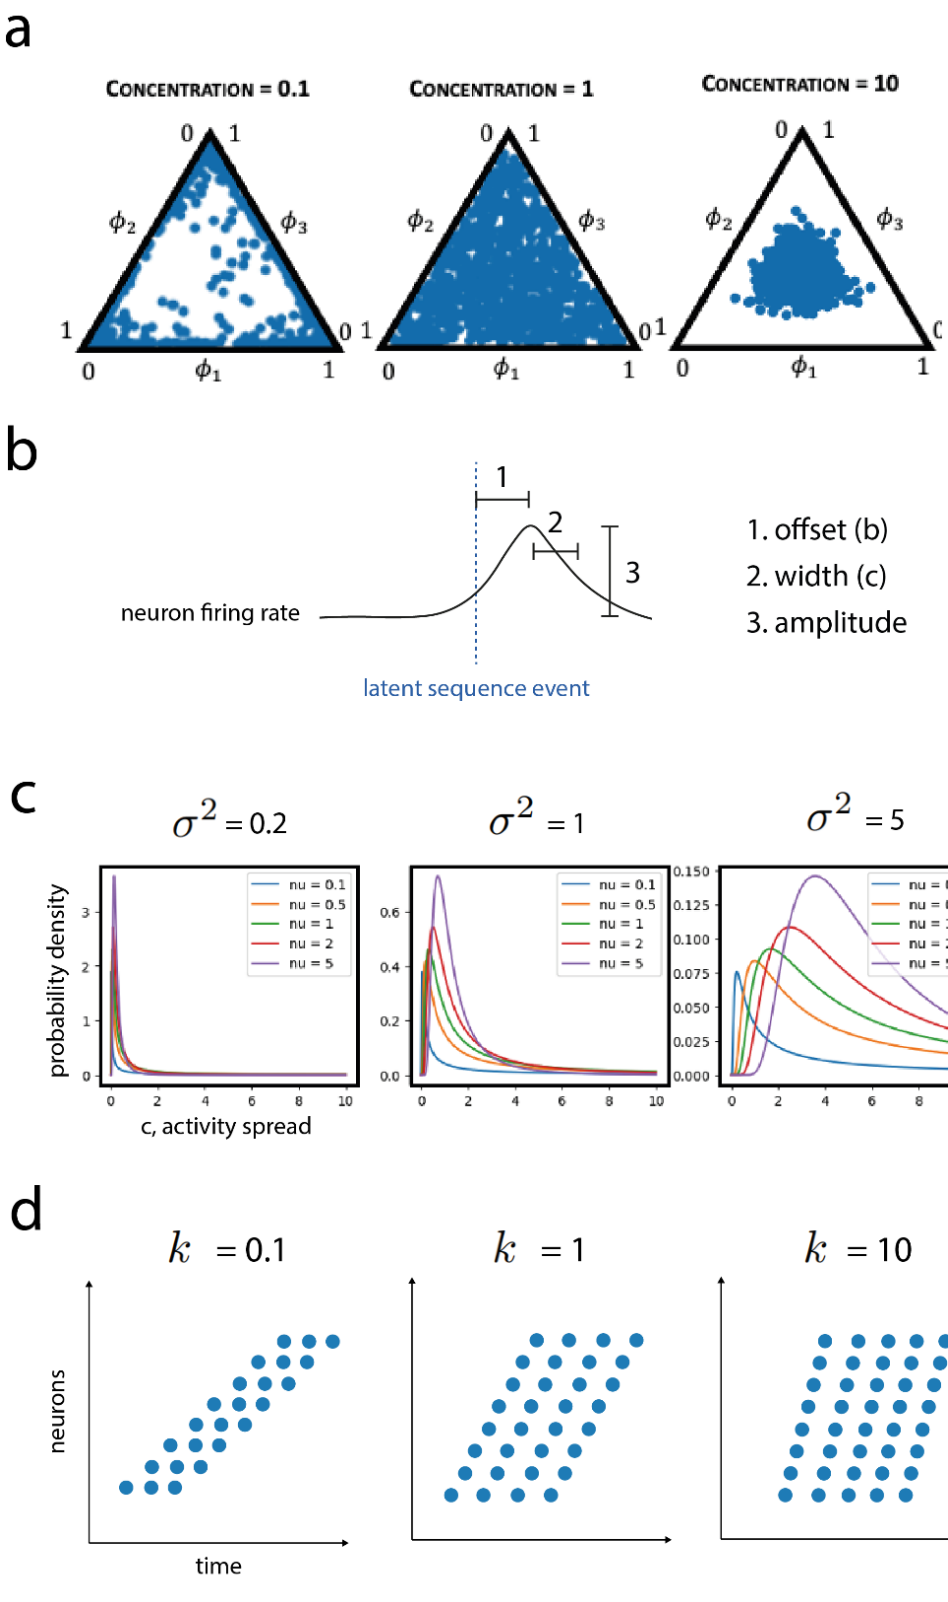

Figure 6.4.1: PPseq hyperparameter features

Legend on next page

Figure 6.4.1: PPseq hyperparameter features
(a) Examples of the Dirichlet distribution. This is a distribution on proportions, i.e. sets of variables that sum to 1. We sample three dimensional proportions from Dirichlet distributions with varying concentration parameters, low values (left) lead to highly concentrated distributions (only one dimension is large), high values (right) lead to a very uniform split. **(b)** Schematic showing how the offset, b, and width, c, parameters control features the firing rate response of a neuron to a sequence event. **(c)** Plots showing the effect of two parameters, v (nu) and σ^2 on the prior on activity widths, c. **(d)** Schematic illustrating the broad effects of parameter *k* in controlling the relative scales of the offsets and activity spreads: low values correspond to tight activity peaks drawn with large offsets (left), large k is the opposite (right).

For the non-fixed hyperparamters, we performed a grid search, training four models for each setting of the hyperparameters and choosing the model with the highest average log-likelihood at the end of training. We searched first at a coarse level, then we zoomed in on the best region of the coarse search to perform a fine search. We used cross-validation: a subset of spikes was held-out from the data, the rest used to train the model. Then the log-likelihood of the held-out spikes was measured under the trained model. Doing this for all ranges of combinations of all 5 swept parameters yielded loss values for 1535 possible combinations (**figure 6.4.2 a**). We then inspected the top 50 models finding they all had loss values that were within error (SEM) of each other (**figure 6.4.2 b**). Since these models had similar likelihood values, we then visually inspected the spike labelling output of the top 15 models and selected a single preferred model by eye. Though the top models all produced qualitatively similar results, we based our selection on output labelling that appeared consistent and sequential. To better understand the chosen model, I then plotted the loss as a function of each varied parameter (**figure 6.4.2 c**). These plots can be thought of as planar slices through the n-dimensional

hypercube of all possible hyperparameter combinations (where n was the number of all possible combinations). While these slices (dimensions for which each parameter value was fixed) should be interpreted with caution they can give some estimate of how varying each parameter impacted model fit. Plotting these we find that when viewed individually, our chosen parameter values tended to generate models which were close to the highest value possible when varying that parameter. The most interpretable, and possibly the most important hyperparameter, is the number of latent sequences our model could use to describe the data. We found spiking activity was best (and most simply) described by 6 latent sequences suggesting that in the data tested, there were 6 repeating latent events which best described the observed spiking.

In summary, I chose PPseq model hyperparameters via a grid search of the parameters which were least interpretable and most influential over the model outcome. For the remaining parameters, I aimed to choose values that would not cause bias in the resultant model. The chosen model was selected from the top 20 best models from the grid search. These models were all within error of each other and produced very similar results. The chosen model was then selected based on visual inspection of these results.

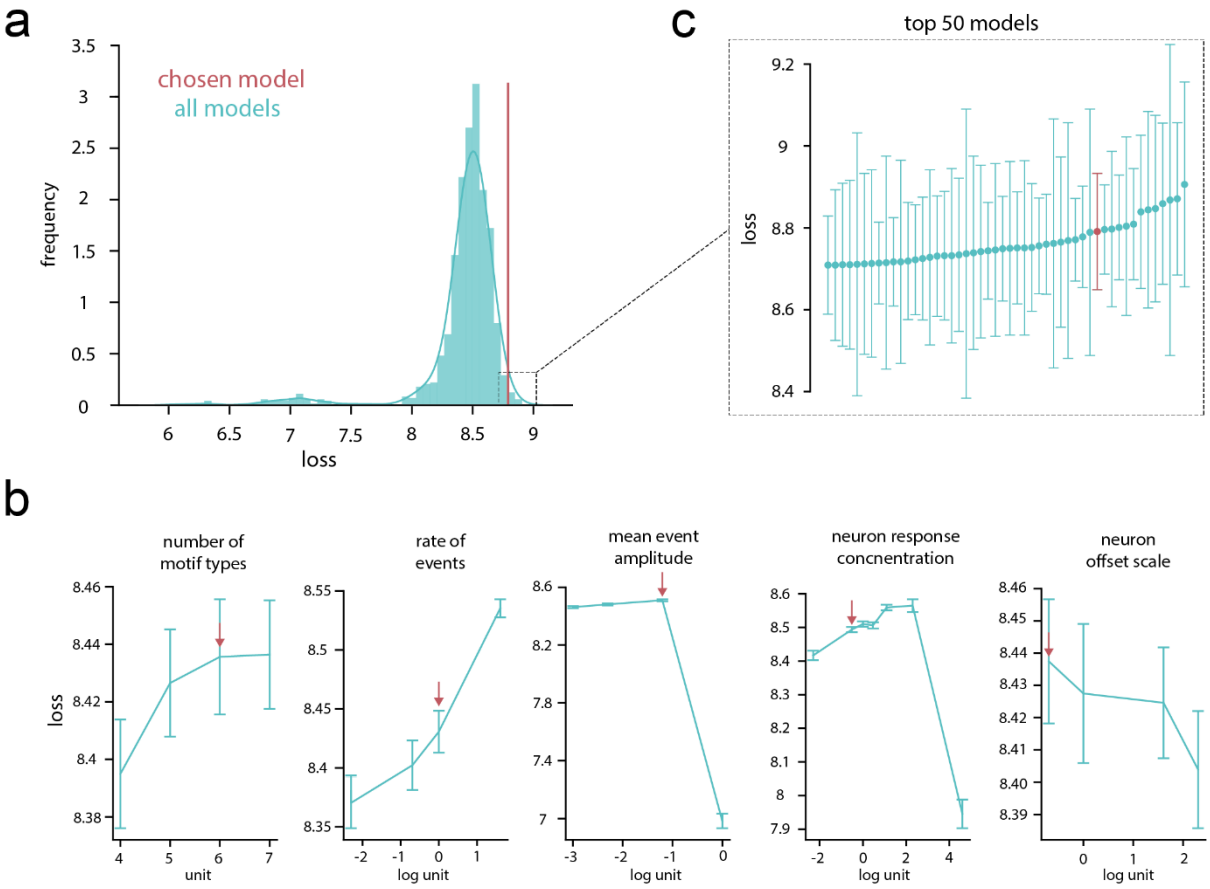

Figure 6.4.2: PPseq hyperparameter selection

(a) histogram showing loss values for every combination of hyperparameters (models) tested in the grid search. The chosen model is indicated by the red line. **(b)** loss values and error (SEM) for the top 50 models from the search, chosen model is showing in red). **(c)** Loss values as a function of hyperparameter value for each varied parameter. Error bars are SEM loss for models in which the parameter of interest was fixed and the other 5 variable hyperparameter were swept across the full range of tested values. Red arrows show the value chosen.

6.5 PPseq reveals latent sequential structure in striatal neural activity during awake task performance.

I aim to use PPseq as a replay detection method by fitting the model to task related striatal activity, then using this as a template to search for the same task related neural structures reactivated offline (during post task sleep). The first step, therefore, is to characterise how PPseq describes procedural activity in the striatum during awake task completion. Recording sessions from trained animals were used and for each recorded session the chosen PPseq model was applied to a 600s period of behaviour. These time periods were selected manually with the aim of choosing epochs with high task engagement and sequence accuracy. As described previously, PPseq is unsupervised, and aims to describe spiking data by underlying repeating latent structures within the neural activity. Such structure is not obvious to the human eye when looking at unordered spikes – before analysis with PPseq. However, after fitting PPseq to the data and rearranging the neuron order by their PPseq identified latent elements, repeating sequences of task related structural activity become clear (**figure 6.5a**). PPseq also assigns spikes to this structure based on features of activity. Hence, after fitting PPseq to each session, spikes which contribute to each identified latent event could be coloured by their respective latent motif. PPseq described task related activity as a series of repeating sequential motifs (**figure 6.5b**). These motifs had a stereotyped order suggesting they may align to distinct phases of task execution. Aligning the motif occurrences to video tracking data and colouring the tracking by the current dominant motif showed this was the case; the repeating elemental motifs robustly tiled the substructure of task execution (**figure 6.5c**). Task related motifs mostly aligned to distinct movements within the task, though motifs were also observed that correlated with reward consumption. For instance, in the example data shown in **figures 6.5 a-e** the purple motif was only apparent at the reward port and only during

rewarded trials when the mouse was drinking. Analysing motif-to-motif transitions revealed a consistent transition structure across the analysed epoch (**figure 6.5d**). To compare PPseq identified structures across multiple recordings, smoothed motif occurrences were projected onto a standard space defined by the average movement tracking trajectory during task execution warped between task port locations. The dominant (most often occurring) motif for each standard space bin was then reported and any task related motifs which were hidden by this analysis were shown with a star (**figure 6.5e**). Applying this analysis to all recording sessions (and circularising the standard spaces) I found that PPseq mostly described neural activity during behaviour with 4 motifs, though for some sessions only 3 or even all 6 latent events were aligned to task related behavioural events. For sessions with additional motifs which were not associated to any task event, motifs were manually assessed from aligned video data. Interesting, this scoring revealed that more than half of these extra motifs were robustly aligned to grooming events. This further illustrates the ability of PPseq to identify hidden repeating neural motifs and underlies its suitability for identifying neural structures related to procedural behaviours (**figure 6.5g**). Other motifs, which were not aligned to a task movement or to grooming tended to be variable, though mostly aligned to background movements that occurred when animals were not engaged with the task.

In summary, applying PPseq revealed that the model could robustly identity repeating task related neural structures. These repeating structural neural motifs titled task execution and tended to align to distinct movement phases within stereotyped expression of the task. My aim is to use PPseq as a replay detection method to search for these same neural motifs reactivated offline.

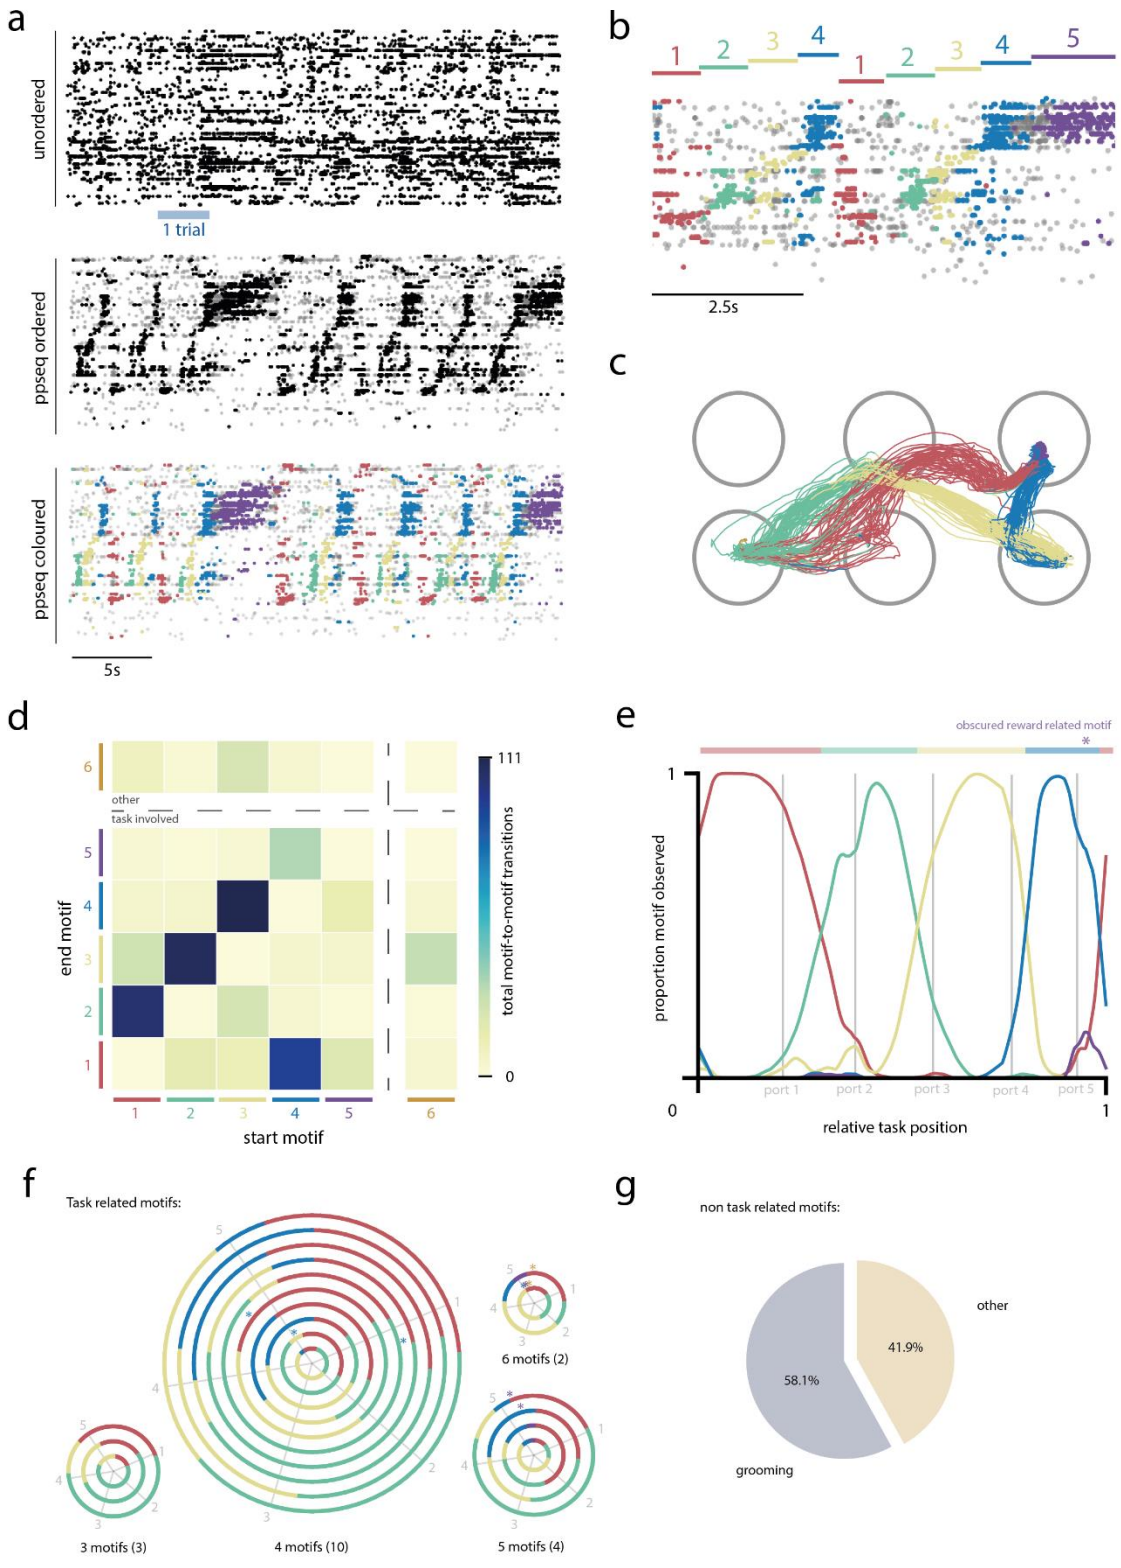

Figure 6.5: PPseq reveals latent sequential structure in striatal neural activity during awake task performance

Legend on next page

Figure 6.5: PPseq reveals latent sequential structure in striatal neural activity during awake task performance

(a) Top: unordered spike raster from an example animal showing spikes for neurons in striatum recorded during task execution (an example single trial duration is shown in blue). Middle: the same spike raster but neurons ordered based on PPseq identified neural structure, (neurons are ordered by latent event contribution and mean offset from latent event midpoint). Bottom: Same as above but spikes which contribute to a latent PPseq motif are coloured by motif. **(b)** Example spike raster (as in 'panel a' bottom) but zoomed in and showing two task trials with task relevant motifs labelled. **(c)** Example movement tracking from 400s of task engagement, coloured by current dominant PPseq motifs (tracking point was centre of the head). **(d)** Motif transition histogram showing numbers of motif-motif transitions during the analysed epoch. Motifs are ordered by preferred transition order. Task associated and other (non-task associated) motifs are separated by dotted line. **(e)** Relative motif incidence curves across standardised task space. Grey lines indicate respective task port locations across standardised space. Top: Flattened representation of incidence curves showing only dominant motifs. Hidden task relevant motif is represented by the star. **(f)** Circularised motif incidences (as in top of panel e) from all recorded sessions separated by total number of task relevant motifs. Grey lines indicate port locations in standardised space (n = 8 mice, n = 19 sessions). **(g)** Percentage of non-task related motifs which were identified as related to grooming or some other feature of behaviour.

6.6 Characterisation of the sequential composition of PPseq identified task associated motifs

PPseq is a recently developed method and application of this model for replay detection is novel. Novelty makes this approach opaque as less is understood about the workings of the method. Therefore, before applying PPseq to offline activity during post task sleep, it is useful to characterise the PPseq motifs identified during awake activity in greater detail. Such characterisation is important for understanding how the chosen PPseq model describes task related activity. Certainly, it is important to recognise that the composition of the observed motifs is highly influenced by the chosen PPseq parameters. Additionally, since I aim to use motif labelling by PPseq to find sequential replay at the level of single neurons, it is key to understand how individual neurons themselves contribute to each labelled motif. Firstly, given the specified hyperparameters of the chosen model, analysis was directed to understanding how frequently each neuron contributed to each motif. Plotting the contribution frequencies (how regularly each neuron contributed at least one spike to each motif of a given type) to different motifs types revealed that for each type there was a similar number of neurons involved and the distribution of neuronal contribution frequencies was similar. For each motif type, there was also a similar range of neuronal contributions – some only contributed spikes to 10-20% of motifs while others consistently contributed spikes to almost all observed instances of a motif (80-90% of the time) (**figure 6.6 a**). Most neurons also contributed spikes to more than one motif, though on average most neurons only appeared in one or two motifs in total (**figure 6.6 c**). Also, across all sessions, the mean relative contribution to the motif in which neurons most commonly contributed (the 'dominant' motif) was 70.3% (**figure 6.6 d**). In other words, for all motifs which neurons contributed spikes to, approximately 7 out of 10 motifs were of a single motif type (the dominant type for that neuron). When neurons did contribute spikes

to another type, analysis of these motifs revealed that that they were significantly more likely to be clustered together: adjacent to, rather than distal to the dominant motif when motifs were ordered by common motif-motif transition frequencies (**figure 6.6 e**). Finally, analysis was also directed to understand the composition of individual motifs. Are motifs simply defined by coactive neurons with jittered timing or is there a temporal substructure? Analysis of the mean relative offset to each latent event revealed that most neurons spiked with a tight temporal relationship (**figure 6.6 b**). Hence, if a neuron contributed to a motif, they tended to contribute spikes at precise timing during that motif. Potting mean timing offsets for all dominant neurons for each motif (ordered by relative motif-to-motif timings) revealed a temporal progression across the full task time. Thus, while full task trial execution can be described in terms of a sequence of stereotyped motifs, each motif can be described as a stereotyped sequential progression of temporally precise neural activity.

To sum up, while neurons in each motif type had similar properties, neurons within each motif were variable in terms of their motif-to-motif spiking consistency. Most neurons contributed to more than one motif in total but overall contributed spikes relatively monogamously to a single motif type. When neurons did contribute spikes to other motifs, usually this was confined to just one or two other types that tended to be neighbouring to the dominant type. Neural activity within each motif was a robust temporal progression of spikes suggesting each motif contained a stereotyped substructure.

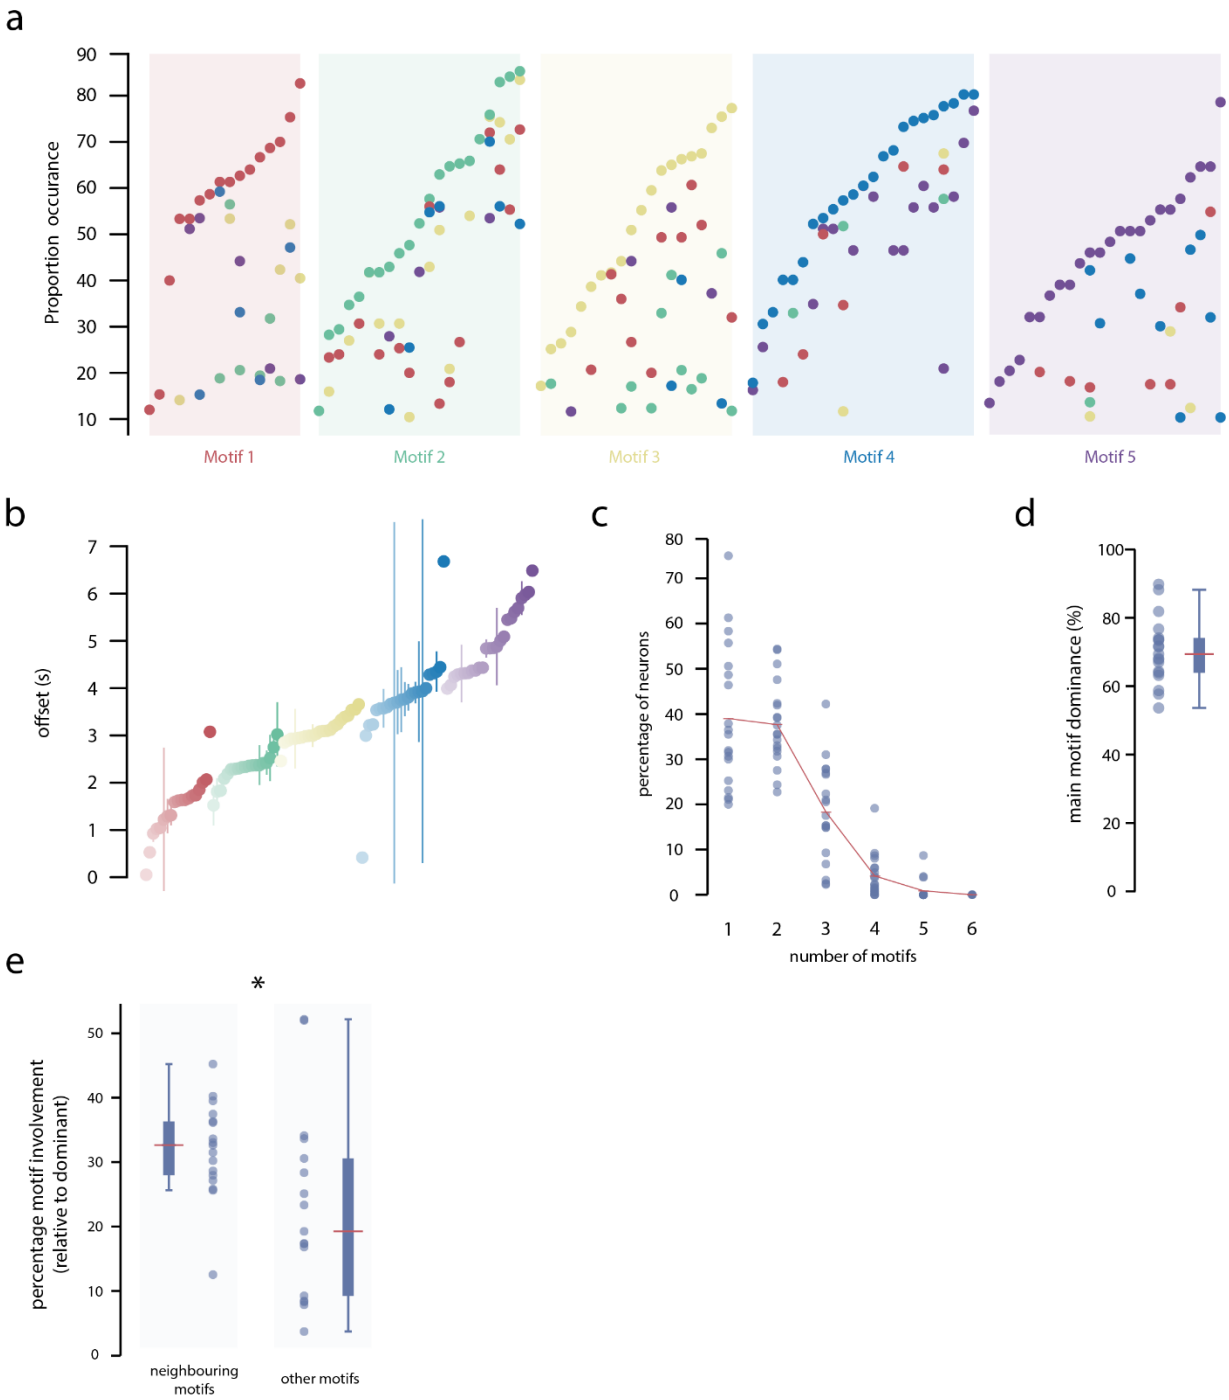

Figure 6.6: Sequential composition of PPseq identified task associated motifs

Legend on next page

Figure 6.6: Sequential composition of PPseq identified task associated motifs

(a) Proportion involvement in each task related motif for each neuron in example animal, ordered by dominant motif and by occurrence rate in dominant motif (motifs ordered by common task transition structure). **(b)** Mean spike times during instances of each neurons dominant motif in example animal (Error bars show standard deviation). Neurons with motifs are ordered by offset from earliest neuron in motif and each motif is ordered and offset by mean distance between motif midpoints. **(c)** Summary plot showing the number of motifs neurons appeared in. For each analysed recording session, the percentage of neurons that contribute (appearing at least 10% of the time) to different numbers of motifs is shown (n = 19 sessions, n = 7 mice, connected red markers indicate median) **(d)** Plot showing the mean relative percentage of spikes that all neurons for each recording session contributed to their most common (dominant) motif (n = 19 sessions, n = 7 mice, red markers indicate median). **(e)** Plot showing mean neuron occurrences between neighbouring and distal motifs for neurons in each analysed recording session. Mean proportions are calculated relative to the dominant motif proportion. Motif order was most commonly observed task order and neighbouring motifs were defined by those adjacent to the dominant motif ($p = 0.027$, paired t-test), (n = 19 sessions, n = 7 mice, red markers indicate median).

6.7 Synthetic data tests to benchmark and validate PPseq replay detection method

As a final step before applying PPseq to offline activity, I first aimed to test whether this novel method is at all feasible as a replay detection approach. There is no ground truth with replay detection so is essential to have prior expectations for how accurately the method will perform. To gauge the performance of PPseq I therefore aimed to test how well this approach works on synthetic data for which the ground truth was known. I also aimed to

benchmark my PPseq approach by comparing the results to that of a Bayesian decoder: a commonly applied method for replay detection (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009; Carr et al., 2012; Ólafsdóttir et al., 2018; Tang et al., 2017; Shin et al., 2019; Tingley and Buzsáki, 2020; Bhattarai et al., 2020). The Bayesian decoder I used was a state-of-the-art method based on a state space model (Denovelis *et al*. 2021). Compared to linear decoding, this method makes fewer assumptions about the speed and direction consistency of replay and has been shown to be more sensitive to a greater range of replay events. Hence, I chose to test PPseq against this decoding model as it is the most sophisticated decoding approach currently available and hence, should provide the strongest competition to PPseq.

In replay detection, the experimenter is searching for occurrences of awake-like neural sequences amongst irrelevant background activity. Hence, replay can be likened to fragments of awake neural sequences embedded into background noise. I generated synthetic data to test the approaches based on this principle. For PPseq tests, synthetic replay was generated by taking a section of PPseq labelled awake activity and finding PPseq identified awake neural motifs. 200 representative motifs split equally across multiple motif types were extracted for each test run. For PPseq only motif relevant neurons were selected and explanted motifs were implanted into shuffled background noise: ie. spikes from the same awake period but with neuron ID shuffled. To test the decoder, synthetic data was generated in a similar but slightly different way. Having first trained the decoder to predict 2D tracking position from the selected period of awake data, the same timepoints extracted from PPseq were chosen but rather than a chunk of neurons, a full portion (all neurons) was explanted (**figure 6.7a**). This was done as it's less clear which neurons the decoder relies on to predict spatial position and so cutting only a specific chunk of PPseq identified neurons could have biased the result.

Having implanted these motifs, regularly spaced into noise, I tested the two methods on their ability to find the hidden motifs (synthetic replay events). For PPseq, identified replay events were defined as clusters of labelled spikes (at least 5 spikes) of a single motif type (see methods for more detail). For the decoder, found replay events were defined by time bins passing a confidence threshold based on decoded spatial coherence (see methods for more detail).

To first test how often the two detection methods identified false positive events I took only shuffled background noise (no implanted motifs) and applied the trained PPseq model and decoder to this noise. As no replay events were implanted, zero events should be detected. Overall, both methods performed very well as their false positive rates were very low (**figure 6.7b**). However, the decoder was significantly more likely to falsely report replay instances in noise suggesting PPseq is better in this regard: predominantly less prone to labelling false positive events. Having established how often the two methods spuriously labelled events, the next tests aimed to determine how good they are at finding true events. In the first test motifs were implanted unmodified. In these tests, once again both methods performed well. On average PPseq identified 90.5% of all implanted synthetic replay events and the decoder identified 85.9% (**figure 6.7c**). Occasionally both methods correctly identified that an event had been implanted but labelled it as the wrong type. However, for unmodified implantations 'mislabelled' events were relatively rare and occurred at as similar rate for both detection methods - although PPseq was slightly (though not significantly) less likely to mislabel data. Since replay has often been identified in reverse temporal order (Diba & Buzsaki 2007) I also tested the detection methods for data in which the implanted motifs had been flipped. Once again, both methods performed well in this test and correctly detected events in a comparable way to the unmodified data: PPseq found 92.3% and

the decoder found 82.5%. As before, PPseq was less likely to mislabel data: significantly more reverse events that were detected by the decoder were wrongly labelled (**figure 6.7c**).

In previous studies, identified replay has also tended to be made up of fewer spikes and from fewer neurons than the same activity during awake experience (Lee & Wilson 2002). To mimic this sparsity, I next tested how the two detection methods faired when a percentage of the implanted spikes were randomly removed or dropped out prior to implantation (**figure 6.7d**). This was done for a range of percentage dropouts. As might be expected, dropping spikes from the implanted motifs caused fewer events to be correctly identified by both methods. Though PPseq performed marginally better initially, divergence between the methods was observed after dropping more than 70% of implanted spikes suggesting the decoder may be more robust to extreme sparsity. On the other hand, however, though the trend across both methods was similar, the decoder was often more likely to mislabel data. Hence, while the decoder was more sensitive, this method was overall slightly less accurate than PPseq.

To determine how important the temporal structure of data was, I next tested whether disordering implanted spikes would impair detection. This was done by swapping a certain percentage of spikes times (**figure 6.7e**). While disordering did impair both detection methods, they were relatively robust to this perturbation. However, for larger perturbations of the spiking order the decoder method was significantly more robust suggesting the temporal order of the spikes is less important for this method than simple coactivity. Interestingly, for both methods disordering did not seem to have much impact on the percentage of mislabelled events. However, as in previous tests, PPseq slightly outperformed the decoder in this respect.

I next aimed to determine to what extent background noise distracts and impairs the two methods (**figure 6.7f**). For both PPseq and the decoder noise was implanted in the same way: spike times were randomly added for all

neurons that were not labelled as involved in each event by PPseq. While both methods were impaired by increasing noise, they were reasonably robust to this perturbation. As in previous tests PPseq was marginally better at finding and labelling true events than the decoder. Additionally, while PPseq remained relatively accurate, adding noise caused the decoder method to mislabel data far more often. In fact, once the background rate had been more than doubled, the decoder was more likely to mislabelled events than correctly identify them. Hence, a strong possibility therefore is that even a proportion of the true labelled events in these circumstances may have only been correctly labelled through chance alone.

Finally, since neural activity is often found replayed at different speeds compared to the same awake dynamics (Lee & Wilson 2002; Davidson *et al*. 2009), I tested how different warping's of the implanted motifs effected the two methods (**figure 6.7g**). In these tests the decoder was more robust than the PPseq method when stretching the spikes. The same was true for compressed motifs, however, the mislabelled rate in these instances was also very high suggesting high inaccuracy. For PPseq, though correct labelling was highly variable, mislabelling remained relatively low across warps

In summary, I generated synthetic replay data to test whether the PPseq detection method is suitable as a replay detection method and to benchmark this approach against a state-of-the-art decoding method. Both methods were very good at finding replay, however there were some differences in their ability to resist perturbations to the implanted spikes motifs. In general, while the decoder method tended to be more robust to extreme perturbations, PPseq was slightly more sensitive and often more accurate than the decoder.

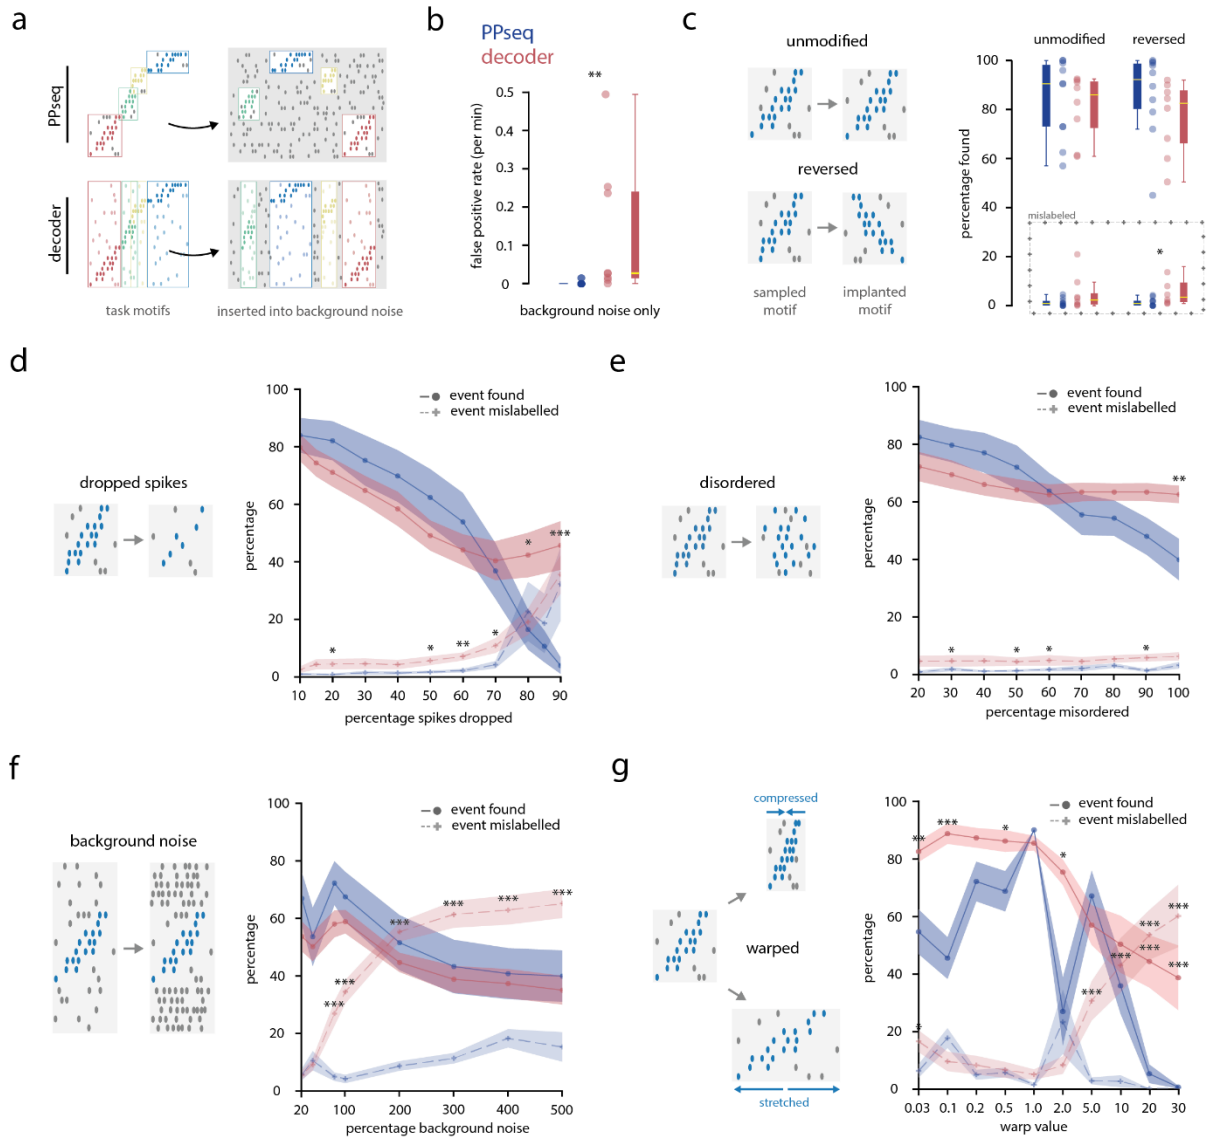

Figure 6.7: Synthetic data tests benchmark PPseq against state space decoding for replay detection

Legend on next page

Figure 6.7: Synthetic data tests benchmark PPseq against state space decoding for replay detection

(a) Schematic diagram illustrating the method for synthetic data generation for tests involving PPseq and the decoder. **(b)** False positive rates (events detected per minute in background noise) for PPseq and the decoder ($p = 0.0051$, Mann-Whitney U test). **(c)** Left: schematic diagram of the type of test done: the modification (or lack of modification) done to each implanted motif. Right, main: Percentage of normal and reverse implanted motifs that were correctly identified by PPseq and the decoder. Right, inset: Percentage of normal and reverse implanted motifs that were identified but mislabelled as the wrong motif type by PPseq and the decoder. ($p = 0.031$ independent t-test) **(d)** Left: schematic diagram of the type of test done: the modification done to each implanted motif. Right, circles: for different percentage spike drop out, the percentage of motifs that were correctly identified by PPseq and the decoder (starred p values left to right: $p = 0.018$ Mann Whitney U, $p = 0.00076$, Man Whitney U) Right, pluses: Motif percentages that were identified but mislabelled as the wrong motif type by PPseq and the decoder (starred p values left to right: $p = 0.018$ Mann Whitney U, $p = 0.015$ independent t-test, $p = 0.0028$ independent t-test, $p = 0.023$ independent t-test). **(e)** Same as d, but for percentage misordered spikes. (Circles, starred p values left to right: $p = 0.017$ independent t-test. Pluses: $p = 0.023$ Mann Whitney U, $p = 0.029$ Mann Whitney U, $p = 0.032$ independent t-test, $p = 0.0027$ independent t-test). **(f)** Same as d, but for percentage background noise added to motifs (Pluses, starred p values left to right: $p = 0.000064$ independent t-test, $p = 0.00045$ Mann Whitney U, $p = 4.65e^{-10}$ independent t-test , $p = 8.54e^{-9}$ independent t-test, $p = 4.61e^{-7}$ independent t-test , $p = 0.000066$ independent t-test). **(g)** Same as d, but for different motif spike warps. (Circles, starred p values left to right: 0.0084 independent t-test, 0.00061 Mann Whitney U, 0.033 Mann Whitney U, 0.014 Mann Whitney U, 0.00060 Mann Whitney U, 0.00018 Mann Whitney U. Plusses, starred p values left to right: 0.045 Mann Whitney U, 0.0006892548609739258 Mann Whitney U, 0.00081 Mann Whitney U, 0.00049 Mann Whitney U, 0.00091 Mann Whitney U). For all plots shown, for each test value and for each group (PPseq and Decoder), n = 10 sessions from n = 6 implanted animals.

6.8 Discussion

6.8.1 Offline plasticity in DLS is required for procedural consolidation and stabilisation

In earlier work (chapter 4), I have provided evidence that the DLS is the site of learning and long-term procedural memory storage. Since, previous evidence suggests that procedural learning includes an offline component (Diekelmann & Born 2010; Rasch & Born 2013; Schmid *et al.* 2020), it stands to reason that this offline mechanism may act in the DLS. Indeed, I found that pharmacologically blocking plasticity in the in this region during post task rest impaired early learning of procedural memory for the task. I also found that the behavioural effects of blocking stratal plasticity for a single rest period was strongest during early learning. While this effect went away in late learning it could be recovered by multiple consecutive days of infusion. What can explain this effect? Possibly, animals became more robust to perturbation as they improved at the task such that a single post session infusion was not sufficient to see impaired performance the next day. Since deficits in performance re-occurred after consecutive infusions, this suggests that for already formed procedural memories, an offline learning mechanism could act in late learning to maintain and stabilise memory – rather than being essential for consolidating day to day improvements in performance as in early learning. However, to some extent this result is surprising. Procedural memories, once formed, are thought to be very stable. Why then does blocking plasticity across days in expert animals destabilise learned behaviour? One possible explanation, which fits within the previous notion, is that the behaviour had not properly stabilised, and animals were still in the learning phase when the testing was done. Certainly, I have shown that even after reaching the final level in my task, mice undertake a period of slow refinements to movement accuracy over several 1000s of trials (**figure 4.3f**). During this period, it is possible that memory is more vulnerable to

perturbation and reliant on offline mechanisms. For instance, post sleep consolidation may be required to selectively enhance memory for accurate movements and suppress nosier movements. It is feasible that without this mechanism, if noisy movements are more prevalent in the behaviour, then memory would be biased towards less accurate movements. In the opposite case, for animals with highly stereotyped movements for which accurate movement already dominate, losing the offline component would be far less detrimental. Whether this explanation is true or not, it should also be noted that the task destabilisation observed in these mice was relatively subtle. Though animals did drop down task levels, indicating performance had dropped, this is a very sensitive description of behaviour based on small changes in task accuracy over tens of trials. A better description is that with consecutive plasticity blocking the behaviour becomes slightly less accurate. Certainly, this kind of motor drift might be expected after a week without proper memory consolidation – especially since this task is skilled and requires a lot of learning. As a comparative (though hyperbolic) example, consider a professional tennis player trying to compete at the highest level after missing several training sessions. Unfortunately, a limitation of these experiments is that these mice were not filmed during testing and so tracking analysis was not possible. Hence, more in depth quantification of changes in movement variability and accuracy could not be performed.

The result from this experiment corroborates the findings of previous work applying the same method in the DLS to block motor learning in a reach to grab task (Lemke *et al.*, 2021). That I find the same result in a different procedural behaviour suggests that this mechanism is likely generalizable across different motor learning paradigms. Furthermore, it is thought that he high firing rates observed in hippocampal replay promote spike time, NMDA dependent plasticity (reviewed in Goto & Hayashi 2023). Since I have shown evidence that procedural memory is supported by an NMDA dependent mechanism, it suggests similar high frequency reactivations might underpin

procedural consolidation. This does, however, raise further questions. Since striatal plasticity also is known to be contingent on local dopamine release (Reynolds & Wickens 2001; Fiorillo *et al.* 2003) it's unclear how this would be coordinated. Certainly, this represents an interesting area for future investigation.

While the conclusions I have outlined so far fit with previous work, there are other possibilities which point to limitations in this work. One caveat is that it is possible that the effects observed are due to side effects from the infused drug that indirectly effected animal performance. For instance, initially (during early learning) mice may have been more sensitive to these effects leading to deficits and developed tolerance through exposure. During the consecutive infusion experiments it is possible that over the course of several days the drug had a hangover effect, causing side effects which led to the observed task deficits. This explanation does seems unlikely for several reasons. Firstly, there was at least a 2-week gap between the last learning infusion and first post learning consecutive infusion. We might expect any built-up tolerance to subside by this time, yet mouse performance was not as perturbed by a single infusion as in early learning. Secondly, in the previous work which employed this method, results are not overly consistent with an explanation based on AP5 side effects (Lemke *et al.*, 2021). Nevertheless, this issue illustrates that my experiment requires better supporting evidence to make strong claims from the results. Future work should aim to better control for the effect of AP5. For example, a useful test would be to check whether a break in training for several days has a similar effect on performance as blocking offline plasticity for several days. If this is true, then it could be taken as evidence that the behavioural effect observed is due to a lack of task consolidation. Secondly, to prove that the drug itself is not impairing performance through some kind of side effect it may be useful to test this by infusing the drug into another brain region or by injecting an equivalent dose

of AP5 into the blood stream. A better aim for future work should be to extend and support these results with complimentary experiments. For example, consolidation has previously been perturbed via closed loop optogenetic methods (Kim *et al*. 2019). A similar approach could overcome some of the limitations of the current approach and allow targeting of specific brain states or LFP markers. Which markers could be useful to target in such an experiment are at this point not clear, however these ideas will be explored further in chapter 7.

6.8.2 PPseq as a novel unsupervised replay detection approach

Having found evidence that an offline mechanism acts in the DLS to support procedural memory I next aimed to search for the neural substrate of this mechanism. Since previous work has consistently shown neural replay underpins other forms of memory consolidation, I specifically aimed to try to detect procedural replay in the DLS. Commonly, Bayesian decoding has been used for replay detection. However, these detection approaches are not well suited to my use case. Firstly, as they have been developed for hippocampal data, decoding methods rely on using SWR events as biomarkers to identify candidate time periods. Secondly, these methods look for neural activity which can be linked to a predefined template – usually a spatially coherent trajectory. Since neural reactivations in striatum have not been previously identified, there is no such biomarker heuristic which I can leverage. Further, while using spatial templates works well for place cell coding, it is not clear how this method could be generalized to non-spatial neural coding. Even with good behavioural characterisation, since we do not have a strong grasp over the single cell code of striatal activity, it becomes difficult to know whether offline activity is just noise or is in fact replay but in a coding schema that I was unable to measure.

Like others before me (Tingley & Peyrache 2020), I have proposed that an unsupervised method could solve these problems. These methods do not rely

on biomarkers or predefined templates and are beginning to become more common in the literature. For example, recent work has used a hidden Markov model to decode position based on population burst events (Maboudi *et al* 2018). While an improvement this approach still relies on typically hippocampal features of the data (spatial coding and population bursting). More suited to my aims are, for example, methods based on quantification of population structure in low dimensional manifolds (Chaudhuri *et al* 2019) or proposed methods based on unsupervised clustering. Such clustering approaches find neural sequences in spike data using dissimilarity measures (Grossberger *et al*. 2018) or neuron-to-neuron spike timing (Van der Meij & Voytek 2018) and hence, could be powerfully applied replay detection. Rather than these methods, however, I have outlined a novel approach based on an unsupervised point process model called PPseq. PPseq is conceptually similar to the examples described above and can be roughly described as a mixture of these approaches: unsupervised clustering of population structure. I chose this method because it is highly suited to my data: PPseq is designed to identify repeating neural structures. Since, my animals repeatedly produce multiple sequential movements when completing the task these structures should be abundant in my neural recordings.

The PPseq model I employed was chosen through an extensive hyperparameter search. This search was limited to only 5 parameters; the others were set based on a small amount of exploratory testing and intuitive interpretation of their effects. This is arguably a slight limitation; a better model could have been found from a full grid search of all 12 hyperparameters. However, this was simply not possible given the compute resources available. Additionally, this limitation is not one that is particularly important in the context of this thesis. I do not aim to use the model to make claims about features of neural activity based on labelling by the model. Even if the model fits the data very well (with minimal loss) it's hard to interpret to what extent the structures that are found relate meaningfully to biology rather

than just being imposed onto the data by the model. Instead, my aim is to use PPseq as a replay detector. Hence, the PPseq labels are just used as a tool to find reactivated spikes. Additionally, the presence of the same labelled structures in sleep can to some extent justify that they are meaningful (or at least non-random) in awake activity.

Applying PPseq to periods of task related activity revealed repeating neural structures. Remarkably these motifs aligned to distinct movement epochs within the task, or even to structured non-task movements such as grooming. It's tempting to claim that this labelling could reflect compositional coding of different movement syllables within the DLS. Possibly this is the case. However as discussed, it's impossible to know whether this labelling reflects true features of neural structure or is simply imposed by the model. This does, however, represent an interesting avenue for future investigation. My own experiments (chapter 5) suggest that inhibition of thalamo-striatal circuits causes impaired chaining of sequential task movements. Hence, these inputs may define compositional 'chunks' of neural activity in the striatum. In the songbird brain this exact kind of compositional 'chunking' has been shown for thalamic inputs to HVC (Moll *et al*. 2023). These inputs triggered the start of each syllable by activating starter cells (first neurons in a domino-like chain of neurons) in HVC. Hence, inhibiting this input only perturbed syllable-syllable chaining when done at syllable boundaries. If PPseq labelled structures reflect true compositional coding, then future work could use these labels to make predictions about the behavioural effect of perturbing thalamic inputs.

Whether motifs are representative of functionally relevant compositionality or not, other questions arise from this result. Firstly, what causes PPseq to separate task phases into distinct motifs? Certainly, the selected model encouraged such splitting, however, a strong hypothesis is that labelling of motifs is driven by orthogonal properties of these neural sequences. In the example mouse shown in **figure 6.5** the motifs matched well to distinct

movement phases. Hence, it seems likely that in this case PPseq was labelling kinematic features of neural coding – such as movement direction or body posture. Secondly, what drives differences in labelling between animals and sessions? Though labelling by PPseq for different recordings was relatively similar it was not completely consistent. From the perspective of the model this is not that surprising since it is probabilistic. However, we might have expected that features of neural activity across animals would be similar enough to negate model stochasticity in this way.

I also analysed the microstructure of each PPseq labelled motif. I found that motifs were on average a tight temporal progression of activity suggesting that PPseq is specifically labelling this temporal order rather than simply finding coactive units – irrespective of ordering. I also found that single units were relatively monogamous to a single motif type and that those which did contribute to other motif types tended to appear in neighbouring motifs. Whether this is driven by overlapping edge neurons (at the start/end of the motif) or by shared coding (eg. shared movement kinematics) between motifs is not clear. In future work analysis could be extended to answer this question.

6.8.4 Validating and benchmarking the PPseq approach with synthetic data

Having established a PPseq model and characterised labelling of awake activity, the next step is to apply the trained models to sleep data to search of reactivations. However, before doing this I aimed to test this method and benchmark it against a decoding approach. I did this by testing these two methods on different kinds of synthetic, ground truth data. I found that PPseq was able to find hidden sequences – synthetic replay. Moreover, this method slightly outperformed the decoder in terms of sensitivity and was significantly more accurate – less prone to mislabelling. However, PPseq was also less robust than the decoder in some extreme cases when the hidden motifs were highly perturbed. Whether this is an issue depends on whether these extreme

cases exist in true replay. Future work should aim to address by cross referencing features of the replay I that find with the synthetic data. By doing this I will be able to better understand whether replay I find could have been biased by poor performance of the model.

A possible limitation of my analysis is that the synthetic data was generated differently for the two compared methods. This was done to avoid biasing the analysis in favour of PPseq, however the opposite may be true. The decoder was given access to more information (more spikes from more neurons) which could have helped it perform better. This may also have made this approach more robust to perturbation since there was a larger pool of background spikes. For instance, when randomly removing spikes, these could have provided an expendable buffer (think cannon fodder) saving the important spikes. Further evidence that this could have been the case comes form the one test (background noise implantation) in which the synthetic data was equal for both approaches. In this test PPseq consistently outperformed the decoder. Hence, the difference between the synthetic data in the other tests may have unfairly impinged PPseq. This should be taken into account when benchmarking PPseq against the decoder as a method for replay detection. Further testing could be done to confirm if this is true. As a final point, testing could also be expanded to analyse combinations of permutations. While tests were done on each perturbation individually (warping, sparsity, flipping etc.) it is not clear how the two approaches (PPseq and the decoder) would perform when faced with combinations. This is an important point considering true replay is likely to look like a combination of modifications – for example, time compressed and reversed.

6.8.5 Summary

In summary, in this chapter I have provided evidence that an offline mechanism acts on striatal plasticity to consolidate procedural memory for the sequence task. To search for the neural substrate of this mechanism I

have introduced a novel replay detection approach based on an unsupervised method. I have demonstrated that this approach is able to appropriately find neural structures in awake activity and that these motif structures relate to task movement phases. Testing this approach on synthetic data revealed it can find replay-like motifs and does so accurately and with reasonable sensitivity. When benchmarked against a classical replay detection approach based on a Bayesian decoder, I have shown that my detection method is comparable and often outperforms the decoding method. Now that I have characterised my approach, in the following chapter I aim to apply this method to search for striatal reactivations during post task sleep.

CHAPTER 7

7. UNSUPERVISED DETECTION OF PROCEDURAL REPLAY IN THE DLS

7.1 Introduction

In the previous chapter I provided evidence that a mechanism which involves the DLS acts during rest or sleep to consolidate procedural memories. For other types of memory, replay of previously observed neural activity is thought to underpin consolidation during offline periods (Wilson & McNaughton 1994, Skaggs & McNaughton 1996; Lee & Wilson 2002; O'Neill et al. 2010). A possibility is that a similar mechanism – neural replay of procedural activity – acts in the DLS during post task sleep. Motivated by this possibility, I have so far established the approach I will use to search for procedural replay; after recording neural activity in the striatum during task execution and subsequent post task sleep, I aimed to employ a novel replay detection method based on an unsupervised point process model called PPseq. PPseq is analogous to a pattern detector. When fitted to task related patterns during awake activity, in theory PPseq should be able to detect the same neural patterns (replay of task related neural activity) during post task sleep. Using ground truth data, I have validated that this approach should work well and should perform in a manner comparable to previously demonstrated replay detection approaches. Hence, now that I have characterised my approach, in the following chapter I aim to apply this method to search for striatal reactivations during post task sleep.

7.2 Unsupervised detection of reactivated task related activity during post task sleep during both late and early learning

To search for striatal reactivations, each recording session was extended to include several hours of post task home cage rest. Post-hoc for each session, a PPseq model was fitted to 600s of task related neural activity. If the model fitted to the data was suitable, then the free parameters of the model were fixed, and the model applied to selected periods of post task sleep (**figure 7. 2.1 a**). Generally, after a short period of home cage activity, animals went to

sleep quickly. Putative sleep epochs were selected for analysis based on movement tracking and average firing rate (**figure 7. 2.1 b**). These periods were then confirmed as reliable sleep periods by eye from video data. The majority of recordings (n = 8 mice, n = 19 sessions) were performed after the learning phase when mice had already reached the final task. However, to determine whether there are differences between offline activity patterns for early and late learning, recordings were also performed during the learning phase (before level 50), (n = 3 mice, n = 6 sessions) (**figure 7. 2.1 d (top)**). For both sets of recordings, applying PPseq to sleep epochs revealed patterns of awake activity replayed offline (**figure 7. 2.1 c**). Average detected event rate was variable across different recording sessions but was not found to be different between early and late learning recordings (**figure 7.2.1 d**). Tests with ground truth data showed that PPseq does not often spuriously label noise as replay (**figure 6.7b**). Equally, I found that for shuffled sleep data (neuron IDs permuted for spikes in each session) spurious labelling by the model was extremely rare: false labelling was only seen in 5 out 25 sessions (**figure 7. 2.1 d**) and when it was observed the rate was extremely low (median false positive rate was very zero, mean = 0.18).

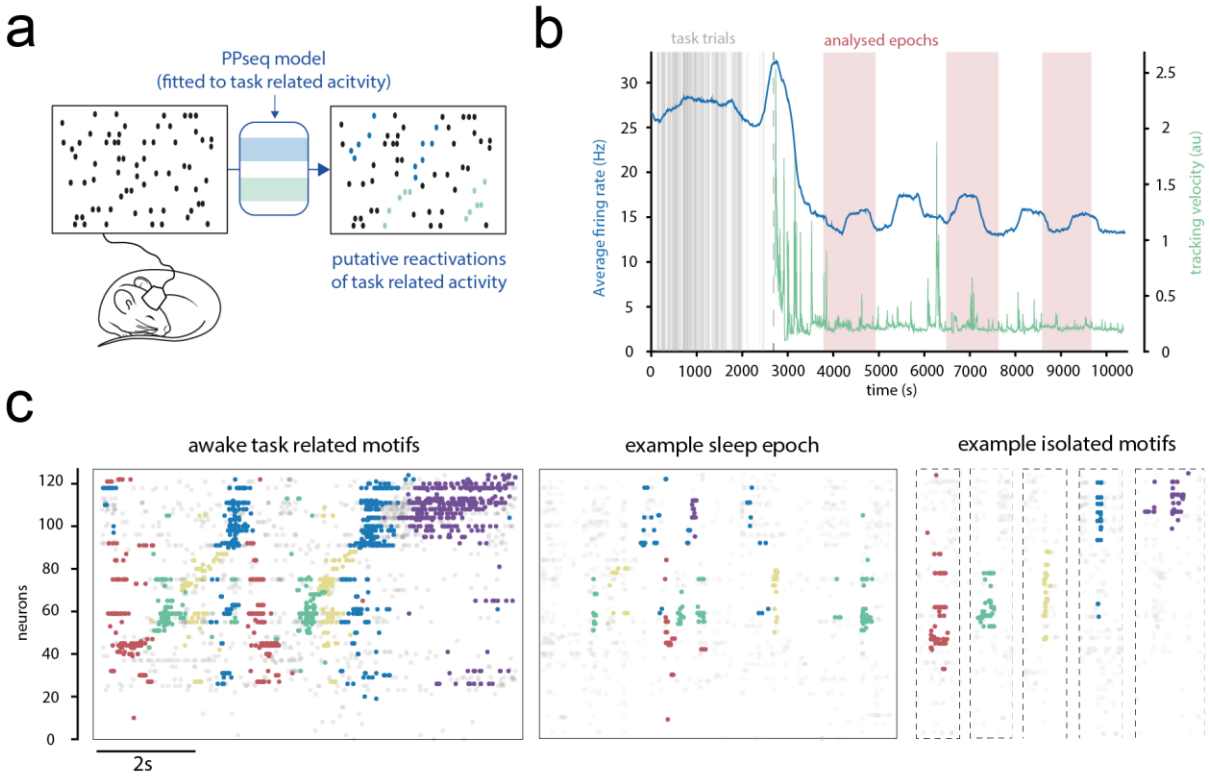

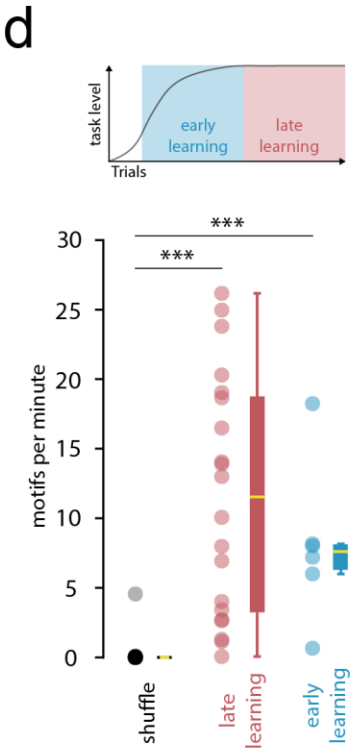

Figure 7.2.1: Unsupervised detection of reactivated task related activity during post task sleep

(a) Schematic showing PPseq replay detection method. Model fitted on awake task activity is applied to activity recorded during post task sleep. **(b)** Example selected sleep epochs (red highlighted). Task activity (trial onsets) are shown as grey lines, smoothed average firing rate (blue line) and home cage movement velocity (green line) are overlayed (video tracking point was the centre of the body). **(c)** Example PPseq labelled motifs for task related spikes (Left) and an example sleep period (centre). Individual examples for each task motif observed during the sleep periods (right). **(d)** Top: illustrative schematic showing early learning and late learning recording groups defined by learning curve progression. Bottom: mean PPseq labelled event rates during sleep periods for each recording group and for shuffled data.

As a final test, for a selection of arbitrarily chosen sessions (n = 6 sessions, n-6 mice), I evaluated time periods containing PPseq events with the previously tested decoding method (Denovellis *et al*. 2021). Time chunks were as small as possible, containing either single replay motifs or multiple overlapping coactive events. For each time period, compared to PPseq motifs, the decoder either found a match – a compatible spatial position (**figure 7.2.2 a**) or decoded a replay event but in a different mismatched location, or did not decode any event above the confidence threshold (**figure 7.2.2 b**). On average, 58% of all PPseq time periods were also decoded as replay containing (**figure 7.2.2 c**) and of these events 71% were spatially matched (**figure 7.2.2 d**). For the minority cases when PPseq and the decoder were not in agreement, analysis was extended to better understand what caused disagreement between the two methods. A possible cause is that one method is biased towards certain motif types: either PPseq could be over labelling this type, or the decoder could be under reporting it. If there is no bias, then the decoder should decode similar percentages of found proportions for all motif types. Overall, the mean variability in the number of events the decoder found for different event motif types was relatively high: on average greater than 10% (**figure 7.2.2 e**), suggesting some bias in one of both models. Besides motif-to-motif variability, the decoder also performed differently depending on whether the analysed time period contained a single isolated motif or contained multiple simultaneously coactive motifs. Single motif events were more likely to be mismatched between PPseq and the decoder and the decoder was more likely to entirely miss periods with coactive events (**figure 7.2.2 f**).

In summary, applying PPseq to post task sleep epochs revealed labelling of task related activity patterns during offline periods. These events occurred at variable rates across sessions, but no rate differences were found between early and late learning recordings. As before, PPseq had a very low false positive rate, labelling very few motifs in the same spikes after neuron IDs had been shuffled. Testing PPseq identified time periods with the Bayesian

decoder revealed that the two methods agreed for the majority of events. PPseq events that were spatially mismatched or missed entirely by the decoder are to some extent explained by motif type bias in one or both of the methods and by differences in the decodability of single and coactive replay events.

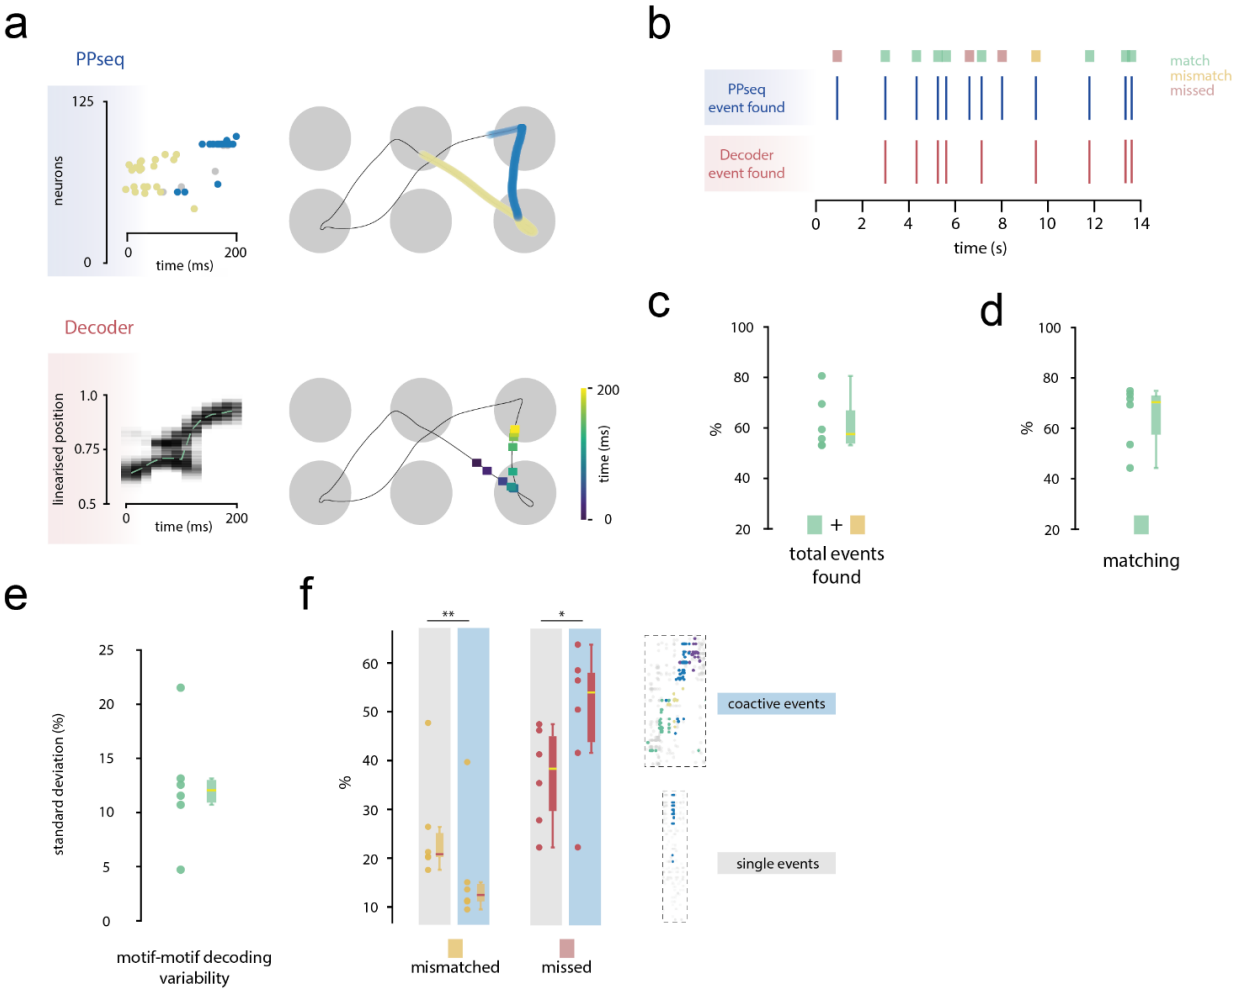

Figure 7.2.2: Testing PPseq identified events with the Bayesian decoder

Legend on next page

Figure 7.2.2: Testing PPseq identified events with the Bayesian decoder

(a) Top left: PPseq labelled spikes. Top right: corresponding PPseq motif positions observed during task activity projected onto average tracking trajectory. Bottom left: for the spikes shown above, the decoded 1D position across linearised task space (green dashed line). Shading shows the posterior likelihood across all spatial bins (white = 0, darker = more likely). Bottom right: maximum of the 1D decoded position (the most probable position) projected back onto 2D average tracking trajectory. **(b)** Example epoch showing PPseq identified events and events that were also found by the decoder. Matches, events that were found by both are show by green markers. Times when an event was found by both methods but did not match in terms of related spatial locations are shown by yellow markers. Events found by PPseq but not confirmed by the decoder are shown by red Markers. **(c)** Percentage of PPseq events that were also found by the decoder. **(d)** The percentage of found events which were spatially harmonious between the two methods. **(e)** Standard deviation of total found percentages between motif typea for each session **(f)** Percentage mismatched and missed motifs for decoded epochs which contained either a single PPseq motif (grey) or multiple coactive motifs (blue). Example decoded periods for single and coactive labelled motifs are shown on the right (n = 6 sessions, n = 6 mice. p = 0.001 & p = 0.01 paired t-test).

7.3 Characterisation of early and late learning replay features

I have established that PPseq was able to identify replay of neural patterns in sleep data and have found good evidence that this labelling was not spurious. I next aimed to characterise these reactivation events. Analysis of the structure of individual labelled motifs revealed that replay events could be split into 3 groups based on their temporal ordering: sequentially ordered motifs, stationary motifs, and fragmented motifs (**figure 7.3.1 a [i]**). To do this, neuron-to-neuron spike ordering for each event was compared to the expected awake ordering for that event type. Expected ordering was defined by each neurons average position across all awake motifs of that type. Plotting spike times coloured (**figure 7.3.1 a [ii]**) or ordered (**figure 7.3.1 a [iii]**) by these relative expected positions for each replay event revealed how much temporal ordering the reactivated motifs shared with their corresponding awake motif type. Events were split into the three groups by fitting regression lines. Those with a strong, sloped fit (see methods) were classified as sequential. Stationary events were classified by events with good fitting, but overall flat regression. Finally, events with incoherent order which could not be suitably fit by a linear regression line were labelled fragmented. By these classifications most replay events were categorized as either sequential or fragmented (**figure 7.3.1 b**).

Similar stationary and fragmented replay have been observed in hippocampal recordings (Denovellis *et al.* 2021). In fact, in general less than half of SWRs– sometimes as low as just 5% - contain spatially sequential replay (Wilson & McNaughton 1994, Skaggs & McNaughton 1996; Lee & Wilson 2002; O'Neill et al. 2010 Davidson *et al.*2009; Shin *et al.* 2019; Tingley & Peyrache 2020). However, linear decoding methods which rely on sequential templates, specifically select for sequential replays. Hence, the sequential ensemble reactivations that I find are more akin to canonical notions of replay. The primary aims of this section are to characterise replay. Since both my methodology and the reactivations I am aiming to understand are novel, this

characterisation is more usefully done when it can be put in context with previous replay findings. As a result, as a starting point I only directed analysis onto sequential replay events. For these data, I first analysed single replay event lengths - as before all analysis was done for both late and early learning recordings (**figure 7.3.1 c**). Hippocampal replay events tend to occur over periods of roughly 100-400ms (Olafsdottir *et al*. 2018). In both striatal recording groups, I found single replay events had similar lengths (first to last spike interval) across groups that were in line with hippocampal replay lengths at around 100 – 400ms on average (late learning median = 200.1ms, early learning median = 280.5ms) (**figure 7.3.1 d**). I next analysed how replay event rates change over time. Procedural consolidation is thought to occur during a critical period just after skill practice (Holz et al. 2012). If replay is the substrate of this consolidation, we might therefore expect reactivation rate to decay over time after sleep onset. Certainly, Hippocampal replay for a particular location has been shown to decay after exposure (Wilson& McNaughton 1994; Kudrimoti *et al*. 1999). Across all data I find a small but significant relationship between time from first sleep onset and replay rate (slope = 0.07, p = 0.035, ordinary least squares; no significant difference between regressions for each recording group). In other words, replay rate decayed by approximately one less event every 14 minutes (**figure 7.3.1 e**). Interestingly, however, for each session across both recording groups there was a strong linear relationship between decay rate and the current replay rate (late recordings, $p = 2.19e^{-8}$, $R^2 = 0.61$, early recordings, $p = 0.005$, $R^2 = 0.64$, no difference between recording groups) (**figure 7.3.1 e, inset**). Hence, event decay was directly related to the current observed rate suggesting there was nonlinear decay back to an equilibrium rate (approximately 5-10 events per minute). Another feature of hippocampal replay which directed my analysis is the propensity of these events to be time compressed, stretched and even reversed (Lee & Wilson 2002; Diba & Buzsaki 2007; Davidson et al. 2009; Denovellis *et al*. 2021). In my striatal recordings I also found that reactivations had all of these characteristics (**figure 7.3.1 f**). Across all

events, for both recording groups roughly equal proportions of forward and reversed replay were observed (late recordings, mean forward event percentage = 56.9%, early recordings, mean forward event percentage = 55.4%). Event speeds ranged from more than 5 times slower up to more than 20 times faster than awake activity (**figure 7.3.1 g**) and this distribution was not different between recording groups (p = 0.72, Kruskal-Wallis). Finally, analysis was directed to determine how the start and end points of labelled replay events related to awake activity. This analysis was done by taking of the first and last neuron in each replay motif and finding the relative average position of these neurons in each awake motif. For both forwards and reversed replay events (for both recording groups), activity tended to propagate through the centre of the motif, rarely reaching the motif boundaries (**figure 7.3.1 h**). In other words, regardless of the replay type, sequential reactivations tended to preferentially involve neurons which made up the central portion of each motif.

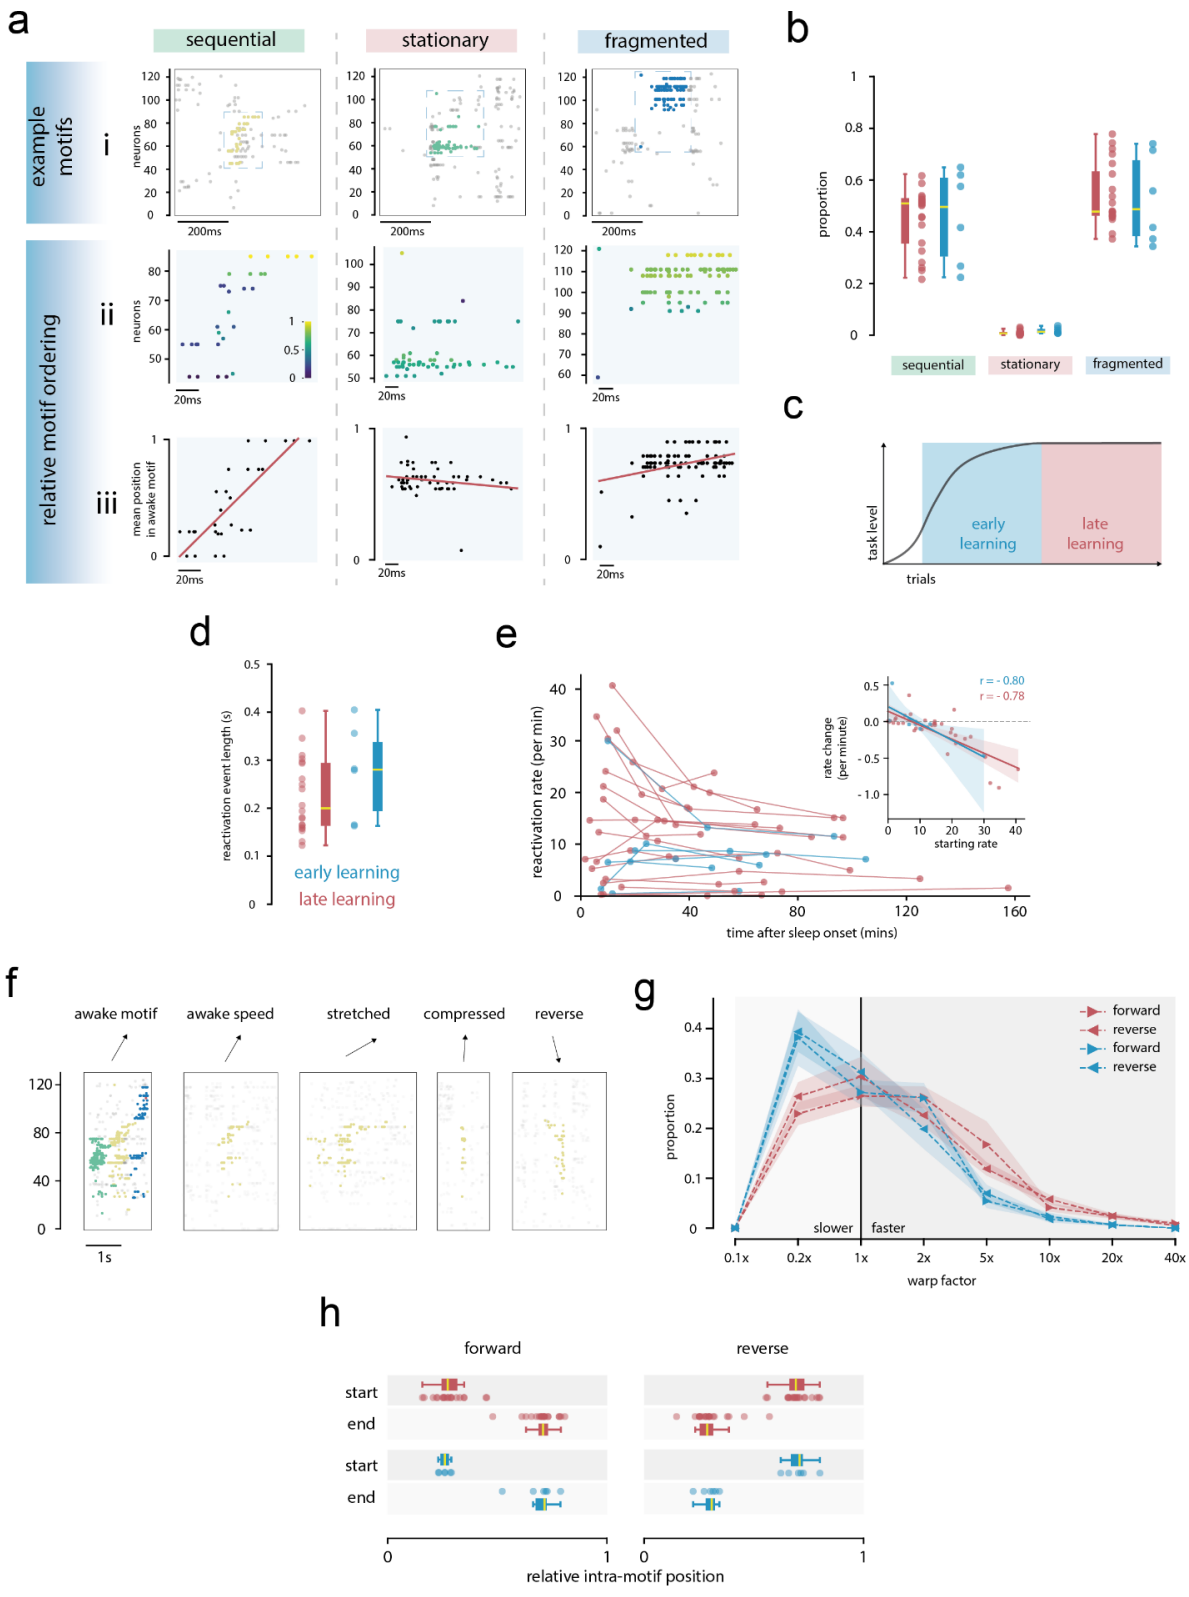

Figure 7.3.1: Characterisation of early and late learning replay features for individual replay motifs

Legend on next page

Figure 7.3.1: Characterisation of early and late learning replay features for individual replay motifs

(a) i: Example PPseq labelled repay motifs for each classification. **ii:** PPseq labelled spikes (from blue dashed boxes in spike raster above) coloured by each neurons relative position in average awake motif. **iii**: Same as above but spikes reordered by relative position in awake motif. Red line is linear regression. **(b)** Proportion of events of each classification for each recording group. **(c)** Illustrative schematic showing early learning and late learning recording groups defined by learning curve progression. **(d)** Single reactivation event lengths (duration from first to last spike) for each recording group. **(e)** Main: Reactivation rates for each analysed sleep epoch against time from first sleep onset. Inset: rate change against starting rate for each pair of analysed epochs per session. **(f)** Example single replay events with different observed characteristics. **(g)** Relative distribution of warp factors for forwards and backwards replay events. 1x represents real world speeds. **(h)** Mean start and end points for all forward (left) and reverse (right) replay events. Position is relative to corresponding average awake motif. Across figure late learning recordings are shown in red and early learning recordings in blue.

As mentioned previously, though most labelled replay events were isolated, PPseq also labelled motifs with overlapping time courses. Any replay which overlapped (occurred within 500ms of one another) were classified as coactive (**figure 7.3.2 a**). Once more, no differences were observed between early and late learning recording groups. Task related coactive replay could either be ordered or disordered with respect to task order (**figure 7.3.2 b**). For both recording groups most observed coactive replay were ordered (late learning recordings; mean = 82.8%, early learning recordings; mean = 73.2%) (**figure 7.3.2 c**). Ordered events could either be forward with respect to awake task order (eg. A,B,C), reversed (eg. C,B,A) or repeats (eg. A,A,A). Analysis of the percentage of each of these categories revealed no significant differences between category occurrences, however, on average across both recording groups most coactive replays tended to be forward (**figure 7.3.2 d**).

In summary, in this section I aimed to characterise the replay events labelled by my unsupervised method. Similar to recent findings from hippocampal data (Denovelis *et al*. 2021) I found that events could be categorised by their sequential relationship to awake activity. Replay was either ordered, stationary or fragmented with respect to average awake activity structure. Also similar to hippocampal replay, analysis of sequential events revealed that replay tended to last around 200ms on average. The rate of these events was variable but appeared to decay over the course of tens of minutes to a baseline rate. Like hippocampal replay, single identified replay events could be reversed and occurred at varying speeds and most replay motifs were slightly time compressed. I also find that on average replay tended to be slightly truncated compared to awake motifs with start and end points biases towards motif centre rather than at motif boundaries. I also analysed coactive events. These were less numerous than single events and could either be ordered or disordered with respect to awake task motif order though I find most coactive replay events were sequentially ordered.

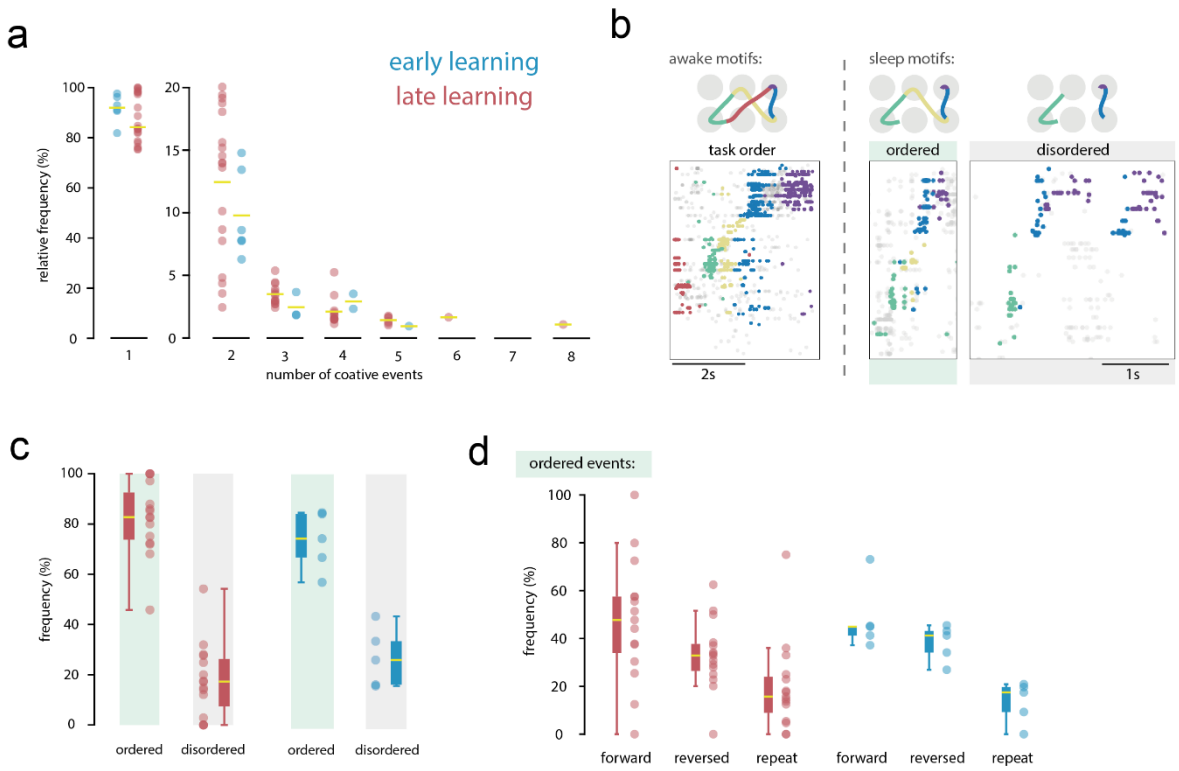

Figure 7.3.2: Characterisation of early and late learning replay features for coactive replay motifs

(a) Relative frequency of single (isolated) and coactive events for each recording session. **(b)** Example PPseq labelled spikes from an awake task related epoch (right) and two sleep epochs (left) containing ordered and disordered coactive replay motifs. Schematics (top) show approximate tracking position of each motif during awake task execution. **(c)** For coactive events, relative frequencies of task ordered and disordered motifs. **(d)** For all ordered motifs, breakdown of the different ordered categories observed. Across the figure, late learning recordings are shown in red and early learning recordings in blue.

7.4 Interactions between awake activity features and observed replay

Having characterised features of the identified striatal reactivations, an outstanding question is how these events related to features of awake activity. Previous work has suggested that procedural consolidation is to some extent directed. For example, consolidation can be prioritised for weaker memories or even prioritised based on future reward expectancy (Kuriyama *et al*. 2004; Diekelmann & Born 2010; Fischer & Born 2009). Since for some sessions my unsupervised method also labelled non -task related motifs such as grooming (**figure 6.5**), I next aimed to determine whether similar prioritisation for task related activity could be observed in the replay I have identified. Indeed, analysis of the relative proportion of task and non-task related events revealed task related activity patterns were more likely to appear reactivated in post task sleep that would be expected given chance (**figure 7.4a**). As before, analysis was performed for both late and early learning recording groups and this difference was true for both groups. I next aimed to quantify whether individual motif occurrence frequencies in awake activity defined their reactivation frequencies in sleep. For all task related motifs in late learning recordings there was a significant linear relationship between motif occurrence frequencies in awake and sleep (linear regression, $p = 6.14e^{-5}$, $r^2 = 0.05$). For early learning recordings this relationship was not apparent ($p = 0.92$, $r^2 = 0.00013$), though analysis of the interaction between regressions suggested no significant difference between groups ($p = 0.82$, combined linear regression and ANOVA) (**figure 7.4b**). Finally, for non-task related motifs during late earning recordings there was no interaction between awake and sleep motif frequency ($p = 0.82$, $r^2 = 0.0006$). During early learning, however, a significant inverse relationship between awake frequency and replay rate of non-task related motifs was observed ($p = 0.028$, $r^2 = 0.19$). However, as before interaction analysis between regressions

showed no significant difference between groups ($p = 0.073$, combined linear regression and ANOVA).

Besides interactions between awake and sleep motif frequencies, I also aimed to determine the relationship between awake and sleep neuron involvement frequencies. For all motif types, neurons that were more frequently involved in replay events tended to be neurons that had also been highly active for the same motif type during awake activity (**figure 7.4c**). As replay spiking was more sparse than awake activity this relationship was not proportionate: neurons were overall less likely to be involved in motifs during sleep than they were in awake activity. Additionally, this relationship was not linear; neurons which more frequently appeared in awake motifs were preferentially boosted in replay compared to less involved neurons. Since the majority of observed motifs were task related this relationship was very similar for when looking only at neurons involved in task related replay motifs. On the other hand, for neurons involved in non-task related motifs, far fewer highly active (frequently involved) task related neurons were observed. Additionally, in opposition to the small effect observed for task related motifs, there was a slight boosting of more infrequently active awake neurons in the observed replay.

In summary I find that observed replay motifs were more likely to be task related motifs than non-task motifs. For late learning recordings task related motifs that occurred more often in awake activity were slightly more likely to be observed as replay. For early learning recordings, individual motif frequency during awake activity was not related to replay frequencies for task related activity, however, more active non-task motifs tended to be supressed during post task sleep. Similar observations were observed for single neurons. For late learning recordings, neurons which had high awake involvement tended to be slightly boosted during sleep (more likely to appear in replay) compared to less involved neurons. For early learning recordings, this subtle correlation was reversed.

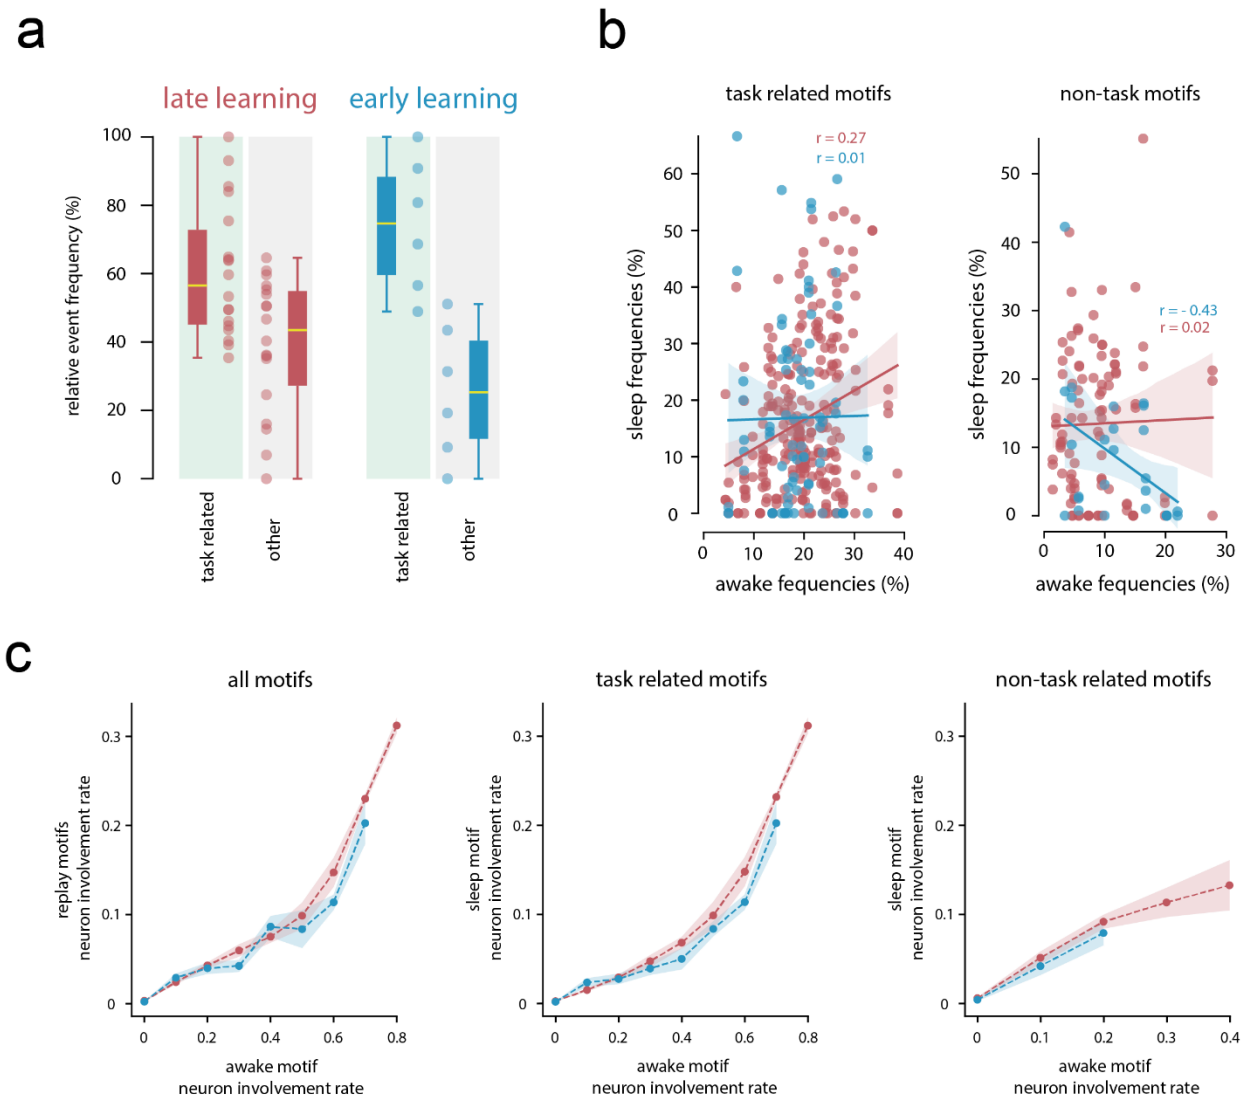

Figure 7.4: Interactions between awake activity features and observed replay

(a) Normalised relative percentages of task and non-task related motifs observed in replay (left; $p < 0.001$, right; $p < 0.001$, binomial test). **(b)** Relative individual motif frequencies during awake task activity and sleep periods. Data shown for all motifs (left), task involved motifs only (centre) and non-task related motifs (right). **(c)** Relative individual neuron involvement frequencies for each motif during awake task activity and sleep periods (frequency = 1, means a given neuron appeared in all observed motifs). Data shown for all neurons for al motifs (left), neurons from task involved motifs only (centre) and neurons from non-task related motifs (right). Data binned by awake involvement rate (bin size = 0.1). Across the figure, late learning recordings are shown in red and early learning recordings in blue.

7.5 Bilateral hippocampal ablation does not impair procedural consolidation of the sequence task

So far, I have established that task related neural activity is reactivated offline and shown that it shares many features with previously observed hippocampal replay. I have provided evidence that features of awake activity to some extent shape the relative neuron and motif frequencies observed during sleep and that this relationship may differ across the learning stage of the animal. Having characterised the identified striatal replay, I next aimed to investigate how local striatal reactivations might be influenced by extra-striatal dynamics. Importantly, I aimed to determine how striatal replay might fit with previous ideas about how brain wide mechanisms shape memory consolidation. Most prominently, consolidation has been considered to be a two-step process; whereby memory traces are transferred from temporary storage in hippocampal plasticity to stable representations elsewhere (usually cortex) (reviewed in Geva-Sagiv & Nir, 2019). For episodic consolidation there is good evidence that this is the case. Firstly, hippocampal lesions are known to cause anterograde amnesia – inability to retain new information without losing already consolidated memories. Further, memory consolidation correlates with both SWR events in hippocampus but also with nesting between SWRs and extra-hippocampal LFP events such as cortical slow wave and spindle events (Latchoumane *et al*. 2017, Maingret *et al*. 2016). These nesting events imply increased inter-area connectivity and hence, these coupling events are thought to be windows for memory transfer (consolidation). Does the same mechanism also support procedural consolidation? Could hippocampal SWR events be upstream triggers for the reactivations I have observed in striatum? Certainly, similar hippocampal-cortical LFP coupling has been showing to correlate with motor improvements in the early stages of procedural memory formation (Kim et al. 2022). However, unlike episodic memory, hippocampal damage does not seem to impair motor skill learning in either humans (Reber & Squire

1998; Hopkins *et al* 2004,Squire, 2009) or rodents (Jackson & Strong 1969; Eckart *et al*. 2011; Will *et al*. 2013; Busse & Schwarting 2016[A]; Busse & Schwarting 2016[B]; Schwarting & Busse 2017). Procedural memory formation may therefore rely on circuits which are part of an entirely separate memory consolidation system to episodic (White & McDonald, 2002).

To better understand whether this is the case, I aimed to determine whether the hippocampus is required for normal consolidation of the sequence task. To do this I performed large bilateral lesions via injection of viral caspase across the extent of the hippocampus (**figure 7.5a**). Histology revealed these lesions were large, relatively complete, and mostly contained to the hippocampal formation (**figure 7.5b**). Despite these large lesions, compared to saline injected control animals, hippocampus ablated mice showed no learning deficits for the task. Lesioned mice reached the final task level in an equivalent number of trials and their learning curves were indistinguishable from controls (**figure 7.5c**). In fact, the only significant difference between cohort learning curves was a brief period during early learning in which hippocampus lesioned mice outperformed control mice. Though learning was unimpaired by lesions, a possibility is that lesioned animals had some deficits in performing the task after learning. To investigate this, analysis was next directed to only expert performance of the task (trials 4000 to 5000). As before, no differences were found between controls and lesions. Mice completed the task with comparable movement speeds (**figure 7.5d**) and made a similar number of port-to-port transition errors. Also, like control animals, lesioned mice were highly task focused; rarely poking into task irrelevant ports (**figure 7.5f**)

In sum, I find that large bilateral lesions to hippocampus did not impair learning or expert execution of the sequence task. Lesioned mice were indistinguishable from controls across all measures of task performance.

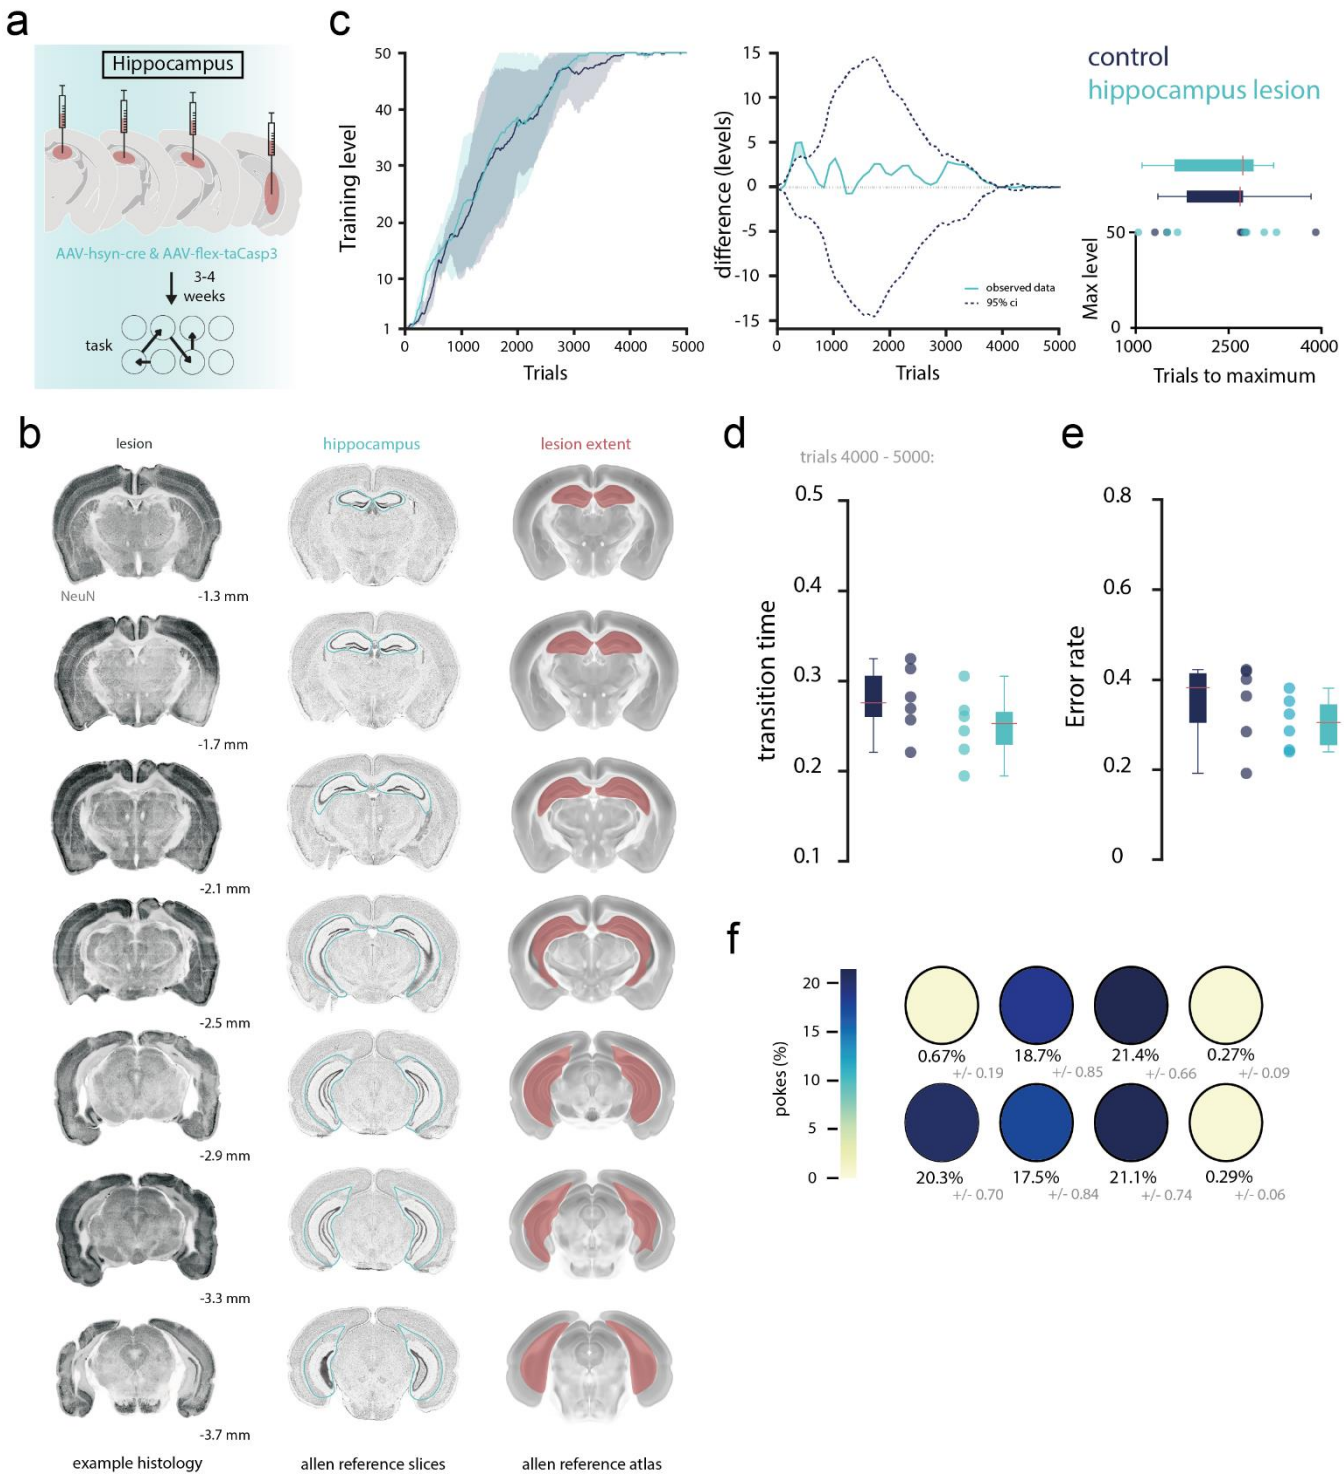

Figure 7.5: Bilateral hippocampal ablation does not impair procedural consolidation of the sequence task

Legend on next page

Figure 7.5: Bilateral hippocampal ablation does not impair procedural consolidation of the sequence task

(a) Top: schematic diagram showing experimental design for Hippocampus lesion experiment. Injection surgery was performed in naïve mice, 3-4 weeks prior to training onset. **(b)** Left: example slice histology showing lesion extent (NeuN immunostaining). Middle: example complete (non-lesioned) slices with hippocampus extent labelled (cyan). Left: approximate lesion extent (red) shown on Allen reference atlas. **(c**) Left: Average training level progression curves for control and lesion animal groups (shaded area denotes standard deviation). Middle: differences in performance between the groups. Dotted lines indicate the 95% confidence interval for shuffled data (see methods). Right: maximum training level obtained within 4000 trials against trials taken to reach maximum. **(d)** Mean transitions intervals between ports (poke out to poke in) for lesion and control group (for trials 4000:5000) **(e)** Port-to-port error rate for lesion and control mice (trials 4000:5000). **(f)** Percentage port poke occurrences, average (mean) across all lesioned animals for trials 4000-5000 (grey numbers are SEM). Control group: n = 6 mice. Lesion group, n = 6 mice.

7.6 Characterisation of striatal replay in hippocampus lesion mice

Since bilateral hippocampal ablation did not impair procedural memory formation, then a compelling hypothesis is that the hippocampus is not involved in shaping striatal reactivations. I next aimed to determine if this was true by characterising striatal replay in animals with bilateral lesions to hippocampus.

As in the preceding experiment, I lesioned the entire bilateral extent of the hippocampus using viral caspase. After training the mice to late learning stage, I then recorded from the striatum during post task sleep (**figure 7.6.1 a**). For each recorded session a PPseq model was fitted to the awake data and then applied to selected periods of post task sleep. As observed in baseline mice (late learning recordings from previous experiments: sections

7.1 to 7.4), PPseq identified motifs in offline activity indicating reactivations of awake neural patterns (**figure 7.6.1 b**). These events occurred at varying rates with a similar distribution to baseline recordings (**figure 7.6.1 c**). As before, replay events could be classified into three categories based on their sequential structure compared to awake activity. No differences were observed between the proportions of each of these classifications: sequential, stationary, and fragmented events (**figure 7.6.1 d**). Single motif lengths for all sequentially classified events were also not different to those seen in baseline recordings (**figure 7.6.1 e**). Again, similar to baseline recordings, replay events in hippocampus lesioned animals could also be forward or reversed (lesion group: 59.3% forward. Baseline: 56.9% forward) as well as stretched or time compressed compared to awake activity. For both forward and reverse motifs the distribution of stretched and compressed events was also not different to baseline ($p = 0.87$, Kruskal-Wallis) (**figure 7.6.1 f**). There were also very slight (though non-significant) differences between event rate decay rates for the two groups (**figure 7.6.1 g**) but no differences were observed when comparing the average start and end points of replay within each motif (**figure 7.6.1 i**).

As observed previously in baseline recordings, replay in lesioned animals could also be overlapping (coactive). I found similar proportions of coactive motifs between lesion recordings and baseline animals (**figure 7.6.1 j**). As before, these events tended to be temporally ordered with respect to awake task order (**figure 7.6.1 k**). Finally, detailed analysis of ordered events revealed no differences between baseline and lesion mice. As observed previously most coactive events tended to be sequentially forward with respect to task order (though differences from chance level were non-significant).

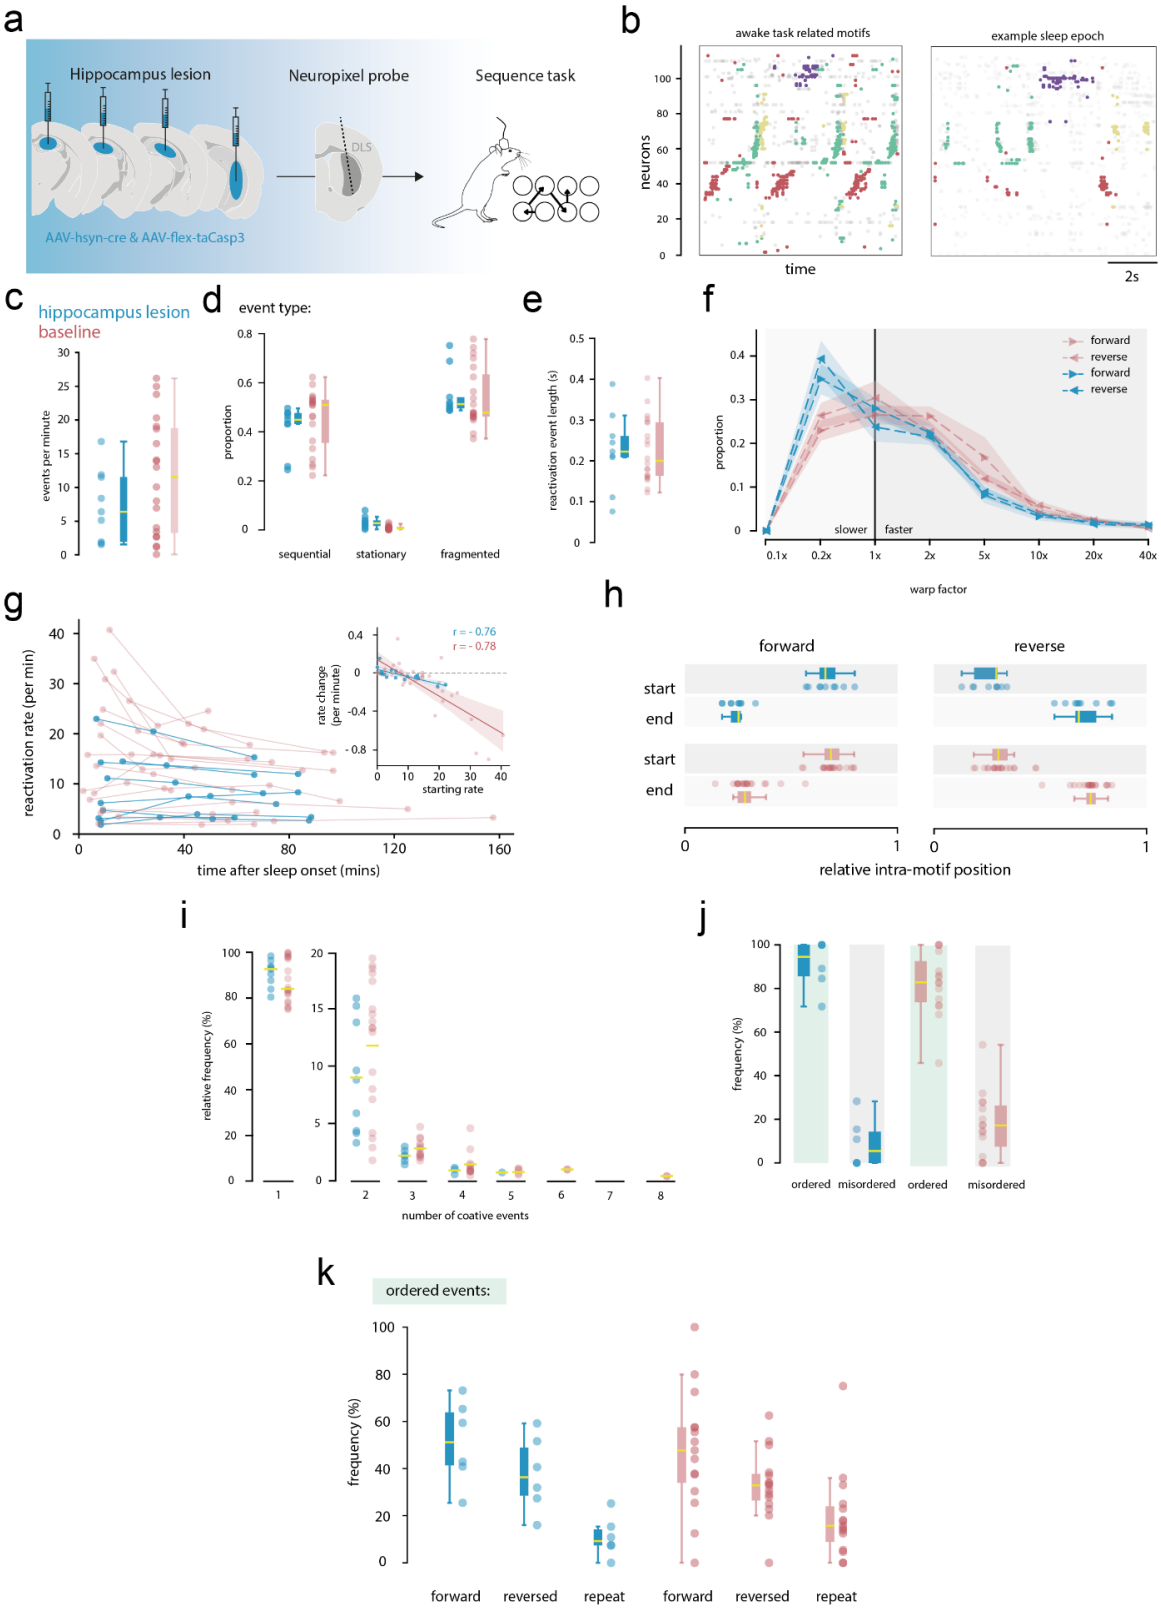

Figure 7.6.1: Characterisation of striatal replay in Hippocampus lesioned mice

Legend on next page

Figure 7.6.1: Characterisation of striatal replay in Hippocampus lesioned mice

(a) Schematic showing experiment design. Injection surgery was performed in naïve mice, 3-4 weeks prior to training onset. Neuropixel probe implantation was performed after training to expert level. **(b)** Example PPseq labelled motifs for task related spikes (Left) and an example sleep period (centre) from a hippocampus lesion animal. **(c)** Proportion of events of each classification for each recording group. **(d)** Mean PPseq labelled event rates during sleep periods for each recording group **(e)** Single reactivation event lengths (duration from first to last spike) for each recording group. **(f)** Relative distribution of warp factors for forwards and backwards replay events. 1x represents real world speeds. **(g)** Main: reactivation rates for each analysed sleep epoch against time from first sleep onset. Inset: rate change against starting rate for each pair of analysed epochs per session. **(h)** Mean start and end points for all forward (left) and reverse (right) replay events. Position is relative to corresponding average awake motif. **(i)** Relative frequency of single (isolated) and coactive events for each recording session. **(j)** For coactive events, relative frequencies of task ordered and disordered motifs. **(k)** For all ordered coactive motifs, breakdown of the different ordered categories observed. Across the figure, late learning recordings are shown in red and early learning recordings in blue. (Lesion group: n = 3 mice, n = 6 sessions. Baseline, n = 7 mice, n = 19 sessions). Across the figure baseline recordings are shown in red and hippocampus lesions recordings in blue.

Though the characteristics of single and coactive replay motifs were very similar between baseline and lesioned animals, a possibility is that the hippocampus does not define specific characteristics, but instead has a role shaping broad features of replay. To test whether this was the case, I completed analysis investigating awake and sleep interactions for lesioned animals (same as in **figure 7.4**). Similar to observations for baseline recordings, lesioned mice also tended to have slightly more task related replay than chance (**figure 7.6.2 a**). Quantifying the relationship between motif occurrence frequencies in awake and sleep activity for task related motifs also revealed no differences between recording groups (Lesions: $p = 1.8e^{-3}$, $r^2 = 0.1$, Baseline: $p = 6.14e^{-5}$, $r^2 = 0.05$) (**figure 7.6.2 b**). However, while for non-task related activity in baseline recordings the awake motif frequency did not correlate with sleep frequency, this was not entirely true for lesioned mice where a slight positive correlation was observed. However, for each recording group neither linear relationship was significant (Lesions: $p = 0.157$, $r^2 = 0.038$, Baseline: $p = 817$, $r^2 = 0.0006$) and analysis of the interaction between regressions suggested no significant difference between the groups ($p = 0.177$, combined linear regression and ANOVA).

Finally, analysis was performed to determine whether correlations between neuron occurrence frequencies in awake and sleep recordings were different between lesion and baseline mice. Non-task related motifs tended to have fewer highly involved neurons during sleep compared to baseline, however, no clear differences were discernible between groups (**figure 7.6.2 c**). Comparisons between groups were difficult to make since lesion data was noisier; due to there being fewer included neurons overall (hippocampus lesions n = 283 neurons, baseline n = 1248 neurons).

In summary, I find that procedural task related replay in the striatum is mostly unaltered by bilateral hippocampal ablation. Replay was similar to baseline recordings across almost all features of replay.

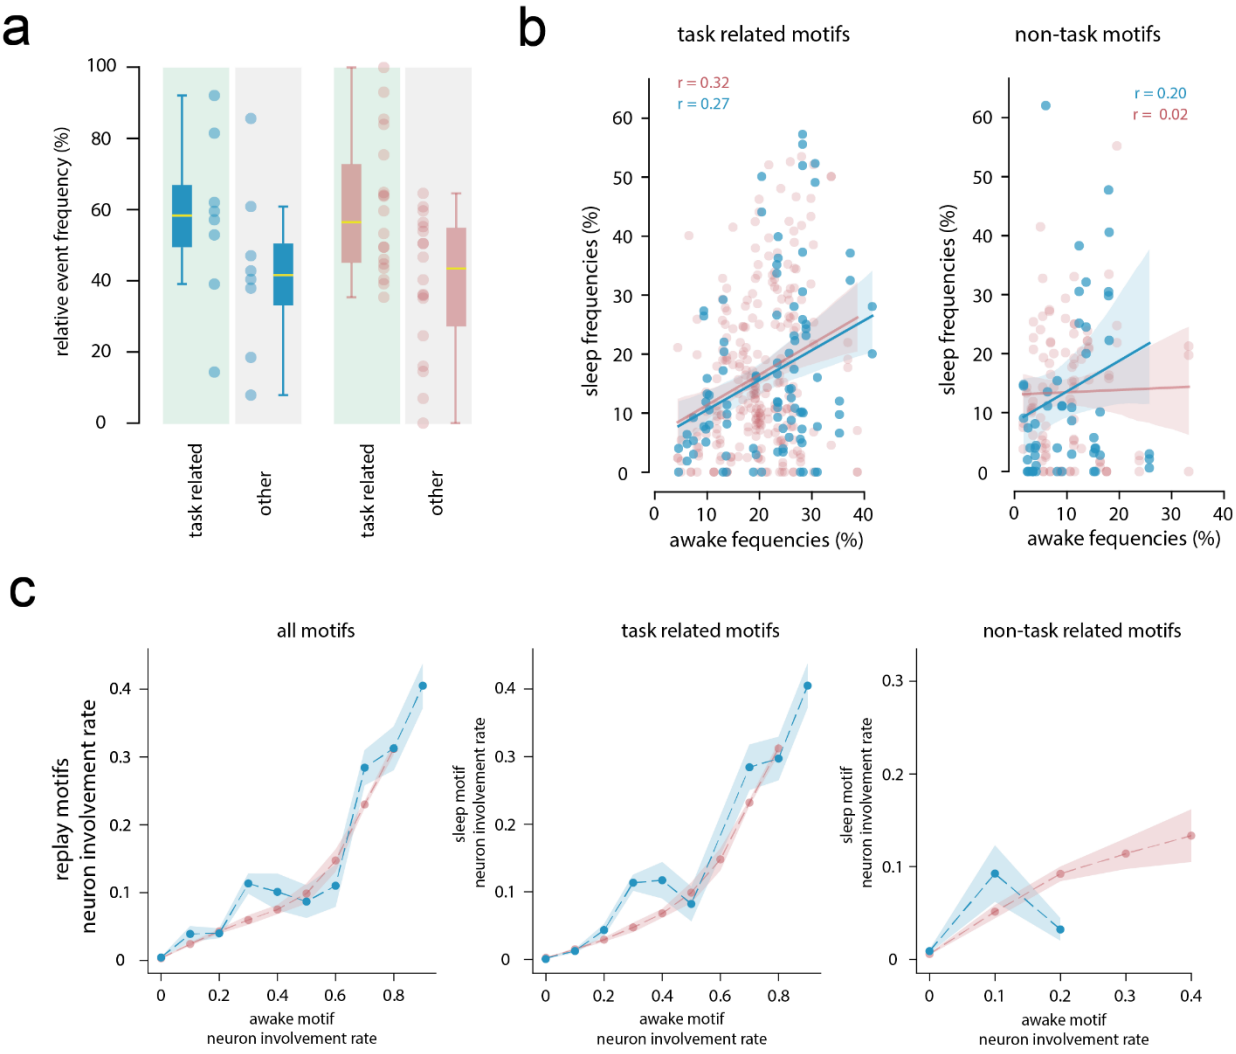

Figure 7.6B: Interactions between awake activity features and observed replay.2in hippocampus lesion mice

Legend on next page

Figure 7.6.2: Interactions between awake activity features and observed replay in hippocampus lesion mice

(a) Normalised relative percentages of task and non-task related motifs observed in replay (left; $p < 0.001$, right; $p < 0.001$, binomial test) **(b)** Relative individual motif frequencies during awake task activity and sleep periods. Data shown for task involved motifs only (left) and non-task related motifs only (right). **(c)** Relative individual neuron involvement frequencies for each motif during awake task activity and sleep periods (frequency = 1, means a given neuron appeared in all observed motifs). Data shown for all neurons for al motifs (left), neurons from task involved motifs only (centre) and neurons from non-task related motifs (right). Data binned by awake involvement rate (bin size = 0.1). (Lesion group: n = 3 mice, n = 6 sessions. Baseline, n = 7 mice, n = 19 sessions). Across the figure baseline recordings are shown in red and hippocampus lesions recordings in blue.

7.7 Discussion

7.7.1 Unsupervised discovery of task related patterns reactivated offline

Having established an unsupervised approach to search for reactivated task related activity in the striatum, in this chapter I applied this method to periods of post task sleep. I found that my method consistently labelled reactivated task related activity: procedural replay. Having found these events, an important first question was to what extent they could be trusted as true reactivations. Certainly, previous testing with ground truth data (chapter 6) lends the method credibility. Also, in line with the previous tests done on background noise, I found that shuffling neuron IDs in sleep activity did not lead to spurious labelling by the model. Hence the method labelled few false positive events. This is good evidence that the replay observed was not often fictitious. Furthermore, directly comparing PPseq with the decoder I found that for most events the two methods agreed that replay had indeed occurred. Moreover, the specific labelling (in terms of relation to awake activity) matched between methods for the majority of instances. However, there were a reasonable number of discrepancies between the two approaches - more than might have been expected given the results of the ground truth data testing done in chapter 6.

What could explain these differences? It seems unlikely that sensitivity differences (number of events found) were driven by PPseq over-labelling the data considering the false positive rate for PPseq on shuffled data was very low. This suggests that the decoder may have been less sensitive when applied to real sleep data. One possible explanation for this comes from the finding that the decoder tended to miss coactive events more often. This contingency (multiple overlapping motifs) was not tested for with synthetic data but there are reasons why overlapping events might cause decoding issues. The decoder expects to find a single spatial location, hence, if two motifs are highly intersecting then the posterior probability would be split

between two locations. This would result in smearing of the decoded position and probable rejection of the replay due to low confidence. In hippocampal data this does not tend to be an issue since this circuit is an attractor network (Tsdoyks 2005); replay here should be reasonably clean – you can only be in one place at a time. However, the striatum can be considered as multiple parallel streams or channels (Alexander & Crutcher 1990 ; Haber 2016). Hence, it possible that neural replay of disparate activity could occur simultaneously.

The other differences observed between PPseq and the decoder were in accuracy (mismatched labelling). This is more difficult to interpret, however, since in ground truth testing PPseq tended to be more accurate, this difference might once again have been caused by fallibility in the decoder. However, ultimately it is unclear what caused the observed differences. Since these differences were not well predicted by the synthetic testing, a clear outcome is that these tests were not extensive enough, nor sufficiently representative of true replay. Now that I have better understanding about which features of activity could be important to test, this limitation could be overcome by extending the ground truth testing. As proposed in the previous chapter, testing should also be done across multiple contingencies (for example, both sparse and temporally compressed activity). Additionally, since there were differences in the decodability of coactive and single events, this feature should be tested as well. Further, since the decoder was relatively variable across motifs within the same session, this suggests there was bias for or against certain motif types in one or both methods. Hence, besides sensitivity and accuracy, future synthetic data testing should also aim to account for and measure these biases. As a final point, it should also be noted that interpretations between methods are also limited by the fact that the decoder was not tested on shuffled data and only applied to time periods in which PPseq had already identified replay. Hence, I do not know whether the decoder might have found additional events which PPseq missed, and it is unclear how many decoded events might have been false positives.

Performing these additional tests would improve the comparative understanding of the differences between the methods.

Having identified replay events, I then sought to characterise them. I first found that across different sessions and animals replay rates were relatively variable. What could explain this variability? Some of this may come from the method. For ground truth data there was some variability in the ability of PPseq to find hidden events. Likely this comes from nuances in model fitting – possibly sessions with more trials or sessions with less variable movements (or perhaps more variable) lead to better, more generalisable models. Some of the variability observed could also be explained by biological features. For example, sleep epochs were chosen fairly arbitrarily, however consolidation is thought to correlate with certain brain states during sleep – most prominently NREM stage 2 (Kuriyama *et al*. 2004; Smith & MacNeill 1994; Forest & Godbout 2000). A possibility that was not explored is that this variability comes from discrepancies in sleep states for the epochs chosen. This should be investigated in future analysis. Additionally, the observed rate could relate directly to features of the preceding behaviour. Careful analysis of whether behavioural features – for example task accuracy or reward rate, correlate with replay rate should be done to investigate this. Finally, a remote but noteworthy possibility is that rate differences could be explained by some other modulatory feature related to circadian mechanisms. Global neural excitability is thought to cyclically depend on circadian factors (Ly *et al*. 2016) and such differences in network excitability could explain observed replay rate differences. Further, mice were housed on a diurnal cycle. During training and experimentation animals were frequently, and quite unpredictably disturbed during their natural sleeping period. Also, recordings were not always done at the exact same time of day for all sessions. Hence, across mice and recording sessions, circadian mechanisms were not well controlled for. Very little work previous has addressed this issue for replay so it's unclear

how important this is; however, this represents an important avenue for future investigation.

I next analysed the sequential structure of each replay motif with respect to the average neuron-to-neuron structure of corresponding awake motifs. I found that events were either sequential reactivations of awake activity, fragmented (nonlinear ordering) or stationary (a single position). This is in contrast to canonical descriptions of hippocampal replay which, constrained by biased linear decoding, is only of sequential activity. This is good evidence that the sensitivity of my unsupervised approach allows for the detection of a far greater extent of all neural replay. Moreover, these findings are in line with the results of recent hippocampal work using a more sensitive (and less biased) decoding approach (Denovellis *et al*. 2021). Together, this is a clear demonstration of the advantages of avoiding simple linear methods, and since my work reproduces findings from hippocampal data, it also suggests the replay categories I observed might be generalised across all neural replays. However, my categorisation is somewhat constrained by the fact that it was based on regression analysis and hence, makes limiting assumptions about the data. Most prominently, this method assumes that sequential events must be unidirectional and at constant speed. Thus, it is possible that many of the fragmented replays were in fact sequential replays that changed direction or rate. Certainly, in the hippocampal work mentioned (Denovellis *et al*. 2021) these changeable events have been shown to occur frequently. It's likely that events like this could be a feature of striatal replay also. Characterisation analysis was specifically performed for sequential events. This was done because the vast majority of previous replay work has focused on linear, sequential replay. Since both my method (unsupervised detection based on a point process model) and the replay I am characterising (striatal procedural reactivations) are novel, it is useful to be able to compare my data directly with previous findings. Hence, I specifically selected for sequential events in my analysis. While this was helpful for comparing my data, this is a

notable limitation in my analysis. Moreover, this undermines one of the main strengths my method; while linear decoding methods rely on sequential templates a major benefit of PPseq based detection is that it is unbiased. By selecting post-hoc for sequential replays this benefit is somewhat negated. Future work should aim to address this limitation by extending analysis to stationary and fragmented replay motifs.

Though some fragmented events might be made up of sequential replay which changes direction and or speed, what can account for truly fragmented reactivations? Possibly these should not be considered valid replay and are just background noise labelled by the method. However, this interpretation is doubtful considering the false positive rate of PPseq was low. Hence, these events must non randomly share enough structure with awake events for the model to consider them equivalent neural patterns. A different explanation could be that these events represent a different functional aspect of replay; perhaps exploration from the current population manifold, or a means for preventing overfitting by the network via disordering spike patterns to weaken Hebbian connections. Similar functions have been suggested previously for hippocampal replay which does not reflect past experience (Gupta *et al*. 2010). Along the same lines, it has been shown that when the whole distribution of replay events are examined, hippocampal reactivations do not necessarily reinstate previous experience but instead match the statistics expected from a fluid undergoing Brownian diffusion (Stella *et al*. 2019). A final speculative idea is that canonical sequential reactivations might not capture the true nature of replay. Possibly consolidation is a generative process, whereby noisy reactivations steadily push plasticity in a desired direction based on the statistics of awake activity. Like the fragmented replay I observed, such processes would be hard to recognise as replay when viewing events individually. However, such events should in some way capture the broad structure of awake activity patterns: leading PPseq to identify them. This is a highly speculative idea but could be an interesting

avenue for future theoretical work modelling whether this kind of fragmented replay could indeed support learning via a generative process.

Due to the uncertainty about fragmented replay, and because sequential events allow for more direct comparisons between my reactivations and previously identified replay, further analysis was directed onto sequential events only. Overall, I find these events were comparable to previously described replay. For example, striatal replay events tended to last for around 200ms – the same length reported for hippocampal replay (reviewed in Olafsdottir *et al*. 2018) suggesting events of this timeframe are a general property of offline activity. Also, I found that replay was reversed in nearly equally proportion to forward events. This is very similar to reported proportions found for hippocampal replay (Davidson *et al*. 2009 Ambrose *et al*. 2016). Also, like hippocampal replay, I found that striatal replay could be compressed in time (Lee & Wilson 2002; Diba & Buzsaki 2007; Davidson et al. 2009; Denovellis *et al*. 2021). However, I also found stretched replay and events at real world speeds. These kinds of replay events have been less often reported for hippocampal recordings, however, this may not be due to differences between the brain regions. In fact, the range of warp factors found for striatal replay were similar to those found in recent hippocampal work using sensitive decoding methods (Denovellis *et al*. 2021). Hence, it seems more likely that previous studies have tended to under report slower events. This once again points to the advantages of my unbiased detection approach. Nevertheless, even compared to sensitive hippocampal detection, in general I found that striatal replay tended to be slower (less time compressed), than hippocampal replay. When compressed, replay tended to be around 2 – 5 times faster than awake activity whereas hippocampal replay tends to be around 20 times faster. This is perhaps not surprising considering hippocampal replay is uniquely driven by fast SWR events (140-200Hz) (Wilson & McNaughton 1994; Olafsdottir *et al*. 2018). A hypothesis from this conclusion is that a different (and slower) oscillation may drive the

compression observed in striatal replay. Good candidates are sleep spindle events (9-16Hz) which occur in the striatum (Lemke *et al*. 2021) and are known to correlate with procedural memory consolidation (Laventure et al. 2016; Cousins et al. 2014).

Analysis of single events revealed that replay of each motif tended to be slightly truncated with respect to motif boundaries; on average replay tended to start around 20% or the way through each motif, transect the midpoint and then end at around 80% motif completion. If PPseq motifs represent meaningful compositional elements this could be a way that replay reinforces compositional structure in neural activity. Replay has been theorised to perform this kind of function (Kurth-Nelson *et al*. 2016, 2023; Liu *et al*. 2019). Also, since in my recordings mostly single isolated replay were observed, this suggests replay could be biased towards separating elemental chunks. However, to some extent this is in contrast to previous findings. For example, in the songbird HVC, procedural replay of song syllables has been shown to frequently cross syllable boundaries (Elmaleh *et al*. 2021). Hence, replay does not seem to be inclined towards promoting syllable separation. Though this is a different circuit, we might have expected procedural replay to look similar even between different species. Furthermore, assigning functional relevance to PPseq motifs is a big assumption. It's not at all clear whether PPseq motifs are at all meaningful in term of compositional coding. The truncated, chunk-like replay I observe could simply be the result of the hyperparameter settings of the PPseq method. PPseq aims to find patterns of sequential activity but does so by labelling short chunk like motifs that occur around derived latent events. Hence, the method may be biased towards finding neurons in the centre of motifs – closer to these latent time points. This could be tested in future work by extending the analysis comparing PPseq labelling with decoded positions. If, for example, PPseq and the decoder agree that replay is mostly made of single, truncated, element-like positions then this could be good evidence for replay having a role in reinforcing compositional structure. This might also suggest that PPseq

motifs are indeed meaningful labelling of true neural structures. Equally, however, decoding could show that PPseq events are artificially discrete and that reactivations actually regularly cross motif boundaries as defined by PPseq.

Times when these boundaries were certainly crossed during replay were for coactive, overlapping motifs. Interestingly I found these events tended to be ordered with respect to task structure. A speculative hypothesis is that these ordered coactivations could represent consolidation of the task structure. Future work could aim to investigate to what extent task structure shapes replay. In hippocampal studies it has been shown that the topology of exploration defines the structure of reactivation. For example, decision points in bifurcating mazes tend to segment replay (Davidson *et al*. 2009). An interesting avenue for future study could investigate how the structure of awake procedural behaviours shape the structure of activity offline. For example, how might training multiple overlapping sequences (eg. ABC and ABD), or suddenly changing the order of a pretrained sequence (eg. ABC to ACB) alter the structure of procedural replay? Such experiments could help shed light on the functional role of these coactive replay events. Related to this, since I find PPseq motifs which align to reward (or at least drinking), a final interesting avenue for future analysis (which I did not explore here) would also be to investigate how these reward events relate to the structure of replay.

7.7.2 Early and late learning reactivations

While most of my recordings were performed in mice that had already learned the task, a subset of recordings were done throughout task learning. I found striatal replay in both of these cases and found that the characteristics of replay events were similar across both recording groups. Replay is canonically discussed as a mechanism that supports learning, or memory formation (reviewed in Olafsdottir *et al*. 2018). Is it then surprising to find

replay after learning has taken place? My AP5 experiments from chapter 6 suggest that the consolation mechanism which supports procedural memory for my task persists after learning. Hence, if neural replay is the substrate of this mechanism, then it is perhaps not surprising to find replay persisting after task learning. Additionally, though replay is often presumed to be a learning mechanism, replay like activity has previously been shown to persist in the brain long after initial learning (Ribeiro *et al*. 2004). Very little previous research has focused on a role for replay in long term storage or consolidation of memory. This disparity perhaps originates from differences between the roles the hippocampus and striatum are thought to have in memory. Hippocampal (or episodic) memory is considered to be a rapid and continuous learning process – storing novel information (new experiences). After initial learning, hippocampal circuits are thought to no longer be required for memory storage (Scoville & Milner 1957; Buzsaki 1989, Siapas & Wilson 1998; Klinzing *et al*. 2019; Spens &Burgess 2023). Procedural memory on the other hand tends to be a gradual learning process, and the circuits that form these memories (such as the striatum) also store memory after learning. Hence, the notion that replay might persist in the striatum after learning is reasonable and far more relevant to procedural memory than for episodic.

Another question that arises from these results is whether or not one might have expected more differences between the characteristics of early and late learning replay. For example, temporal compression, possibly during sleep spindles, is thought to be a way of promoting plasticity (Dickey *et al*. 2021). Hence, one might have predicted that early learning replay would have a greater number of temporally compressed events – to better facilitate learning. Clearly however, this was not the case. This highlights that caution is required when trying to fit functional roles to observed features of replay as so little is understood about these mechanisms.

Though the characteristics of replay were similar between early and late learning there were possibly some differences observed in the way that reactivated activity was prioritised – or least in the way it related to awake

activity. In this analysis, for both recording groups, I aimed to determine whether features replay correlated with preceding activity during awake task execution. Evidence suggests that procedural consolidation can be directed or prioritised for certain memories – or even parts of memories (Kuriyama *et al*. 2004; Diekelmann & Born 2010; Fischer & Born 2009; Schonauer *et al*. 2014). Thus, I aimed to search for mechanisms which could underpin such prioritisation and also investigate whether prioritisation differed depending on learning stage. Firstly, for both recordings, if reactivations represent a consolidation mechanism, then it seems likely that replay should preferentially reflect task activity. Indeed, for both early and late learning recordings I found that a slightly larger proportion of replayed motifs were task related suggesting post task replay might preferentially support task memory consolidation. This is in line with similar findings from hippocampal replay (Igata et al. 2020). A potential limitation in this finding was that each PPseq model was deliberately fitted onto periods of task related activity. Hence, although PPseq did identify non task related motifs in these periods, the dominant patterns in the training set tended to be task related activity. Consequently, there could be some bias towards finding task motifs in sleep data. A better future experiment could be to add a second PPseq model, specifically fitted onto non-task related activity. With my current data set this could be achieved by fitting a model to the period of open behaviour when animals were first placed back in their home cage. Comparing the number of replay motifs found by each model in the same sleep epoch could overcome any potential bias and reveal whether task related activity was preferentially replayed.

Across both recording groups I also found that for task related neural activity, neurons which were highly involved in awake motifs also tended to be slightly more involved in replay than those with lower awake involvements. Though this was a small effect, the opposite was true for neurons in non-task related motifs. A possible conclusion is that for task related activity, neurons which carried information most strongly were preferentially boosted to facilitate

consolidation of these memories. If the opposite effect was indeed true for non-task related activity, then it suggests that the replay observed in post task sleep might have weakened these memories. This is, however, highly speculative, and based on a very small effect. Nevertheless, this does represent an interesting potential mechanism which should be explored in future research and analysis.

Finally, I also found differences between the relationship between task and non-task related motif frequencies in awake and sleep activity. Interestingly there were also differences between early and late learning recordings in this regard. For late learning recordings I found that task related replay correlated with how often each motif was observed in awake activity. The volume of non-task related motifs on the other hand, did not correlate between awake and sleep activity. Put simply, in late learning, replay of task related activity could to some extent be predicted by how often a neural pattern occurred in awake activity while non-task activity could not. For early learning recordings however, no correlation was observed for task related motif frequencies between awake and sleep activity. Furthermore, non-task activity was in fact anticorrelated across awake and sleep such that motifs tended to be more supressed during sleep if they appeared frequently during awake activity. This implies that different prioritisation mechanisms might underpin early and late learning replay. In early learning replay, circuits might aim to ignore irrelevant background details from memory (do less irrelevant movements). In late learning, however, impetus might be for reinforcing important (or at least frequently observed) memories (do more relevant movements). Once again though, these correlations were not particularly strong so a great deal of caution should be taken when making conclusions about this data.

7.7.3 Hippocampus lesions don't impair normal learning or alter replay

Having characterised striatal replay, I next aimed to investigate how procedural reactivations in the striatum might be influenced by dynamics in

the other brain regions. In particular I aimed to investigate whether the hippocampus has a role in shaping procedural replay. Certainly, this region is known the be essential for other forms of memory consolidation and there is evidence that SWR events in the hippocampus could be triggers which choregraph offline activity across the brain – even for motor learning (Latchoumane *et al*. 2017, Maingret *et al*. 2016; Kim et al. 2022). Hence, I aimed to determine whether hippocampal SWR events could be upstream triggers for the reactivations I have observed in striatum. To do this I first investigated whether the hippocampus was required for normal procedural memory formation in my task. If procedural memory is contingent on hippocampal dynamics, animals with bilateral hippocampal ablation should have had impaired task learning. However, I found that lesioned animals had no discernible impairments in learning or expert expression of procedural memory for the sequence task. This result is in line with multiple previous lesion studies (Reber & Squire 1998; Hopkins *et al* 2004,Squire, 2009; Jackson & Strong 1969; Eckart *et al*. 2011; Will *et al*. 2013; Busse & Schwarting 2016[A]; Busse & Schwarting 2016[B]; Schwarting & Busse 2017). A clear conclusion from these findings is that procedural memory formation is entirely independent from the hippocampus. Hence, as proposed previously (White & McDonald, 2002), procedural and episodic memory systems appear to be entirely distinct from one another.

If striatal replay is the substrate of procedural consolidation, a clear prediction from these findings is that hippocampal lesions should not alter this replay. I aimed to test whether this was true by recording from the striatum in mice with bilateral hippocampal ablations. In line with the predictions from the behavioural findings, procedural replay was indeed observed in hippocampus lesion mice. Furthermore, the characteristic of these reactivations were extremely similar to those found in baseline recordings suggesting the hippocampus does not have influence over the content or form of striatal reactivations.

If procedural replay is entirely independent of the hippocampus, what could be the mechanism that triggers and choregraphs striatal replay? One possibility is that this mechanism involves certain hippocampus independent LFP events. For instance, procedural memory enhancements are known to correlate with the number and density of observed sleep spindles (9-16Hz) (Fogel et al. 2006, Nishida & Walker 2007; Rasch *et al*. 2009; Barakat *et al*. 2012). Procedural consolidation also correlates strongly with NREM sleep periods; when delta band (0.1-5Hz) oscillations are largest. Spindle density also peaks during NREM times (Laventure et al. 2016) and it has been suggested that temporal nesting between spindles and delta band oscillations could define windows in which procedural consolidation occurs (Lemke *et al*. 2021; Schreiner *et al*. 2021). Further, coupling between spindles and two distinct types of delta band oscillation; δ wavs (isolated up states) and slow oscillations (a down state followed by an up state) has been shown to have different effects on consolidation – leading to either memory weakening or strengthening (Kim *et al*. 2019). Considering this evidence, and that I find no role for the hippocampus in procedural replay, a clear direction for future work is to investigate whether these LFP phenomena correlate with the striatal replay I have identified. Since nesting of these evens might have different effects on consolidation, an interesting avenue for future analysis would be to determine whether LFP nesting between these bands correlates with different features of replay. If this is the case, it could imply a functional significance for different replay characteristics.

7.7.4 Summary

In sum, in this chapter I aimed to apply the unsupervised replay detection method I have adapted (outlined in chapter 6) to neural activity recorded during post task sleep. In doing so I aimed to determine whether neural replay of task related activity exists in the striatum. I have demonstrated that my method finds reactivations of task related activity and validated these replay

events by testing the method on shuffled data and by comparing these results with those of a Bayesian decoder. I characterised the identified replay and found it shared many features with previously reported hippocampal replay implying reactivation across brain regions may share common features. Replay characteristics did not vary greatly between early and late learning recordings though weak evidence suggests that across these learning stages the brain might prioritise striatal replay in different ways. Finally, I investigated whether the hippocampus has a role in shaping procedural memory. I find bilateral ablation of this region has no effect on procedural memory formation and accordingly these lesions did not alter features of striatal replay. This suggests that procedural consolidation is independent of the hippocampus and that procedural memory systems are entirely distinct from systems related to episodic memory.

CHAPTER 8

8. GENERAL DISCUSSION

This thesis has investigated the role of subcortical circuits for procedural memory formation. My work on this topic has focused predominantly on how striatal and thalamic circuits function for this process in two separate (but highly overlapping) themes. Firstly, I have investigated the role these circuits play in orchestrating procedural behaviours in real time – during wakeful periods. The second part of my work has focused on how offline activity in these circuits supports their awake function. Throughout this thesis I have provided discussion for the findings in each chapter. I have outlined specific findings and their meaning in context with previously published work. I have also discussed the specific limitations of my experiments and have suggested experiments that would complement and extend the work I have show here. In this chapter I aim to discuss my results broadly in context with current literature. In doing so I will assess how contributions in this thesis fit within a larger conceptual framework and suggest where future investigation should be directed.

8.1 Mice learn a novel procedural behaviour based on sequential nose pokes

In chapter 4 I outlined a novel sequential behavioural task which I have developed. I demonstrated that mice were able learn this task well, developing stereotyped behaviours that had characteristics of procedural motor skills. This task was suited to the work shown in this thesis but also has utility for future investigation. Similar to a recently developed piano-like lever press task (Mizes et al. 2022) a strength of my task is that it allows the experimenter to define the exact movement-to-movement sequence an animal must perform. Hence, the task allows training of sequences which can easily be remixed. An outstanding question which this thesis has not addressed is how the brain is able to use and reuse motor elements between different motor sequences. In future work, my behaviour paradigm could be used to train overlapping sequences (eg. A,B,C and A,B,D) or even sequences

with repeating elements (eg. A,B,A,C). Tasks with these features will be essential for understanding how the brain hierarchically organises motor elements in this way.

While my task, and others like it are useful for studying procedural behaviours it is important to recognise that studying a single highly shaped behaviour, such as the one I have studied in this thesis, is limited in its scope. Such paradigms are always susceptible being overly specific. A true mechanistic understanding of how the brain learns and produces procedural behaviours should be generalisable across multiple paradigms. On one hand, this is a strength of my work. Many of my findings corroborate previous work done in rodents performing lever presses (Wolf *et al.* 2022). Hence, my work provides some evidence for generalised principles. However, future work should also aim to test these ideas in more ethological settings. With recent developments in behavioural tracking and segmentation it is becoming possible to better quantify the structure of naturalistic behaviour (Mathis *et al.* 2018; Wiltschko et al. 2015; Marshall *et al.* 2021). Consequently, at this time there is enormous experimental potential to contextualise our understanding of how procedural circuits function by studying them in rich naturalistic behavioural settings.

8.2 The striatum as the site of learning and storage of procedural memory

In chapter 4 I also showed that both learning and long-term storage of procedural memory are dependent on the dorsolateral (or more precisely, motor cortex recipient) portion of striatum. This finding adds to a large and growing body of literature which suggests that the DLS is a key node in a circuit which controls procedural behaviours (Miyachi *et al.,* 1997; Berridge & Whishaw, 1992; Yin 2010; Andersen *et al.* 2020; Wolff *et. al.* 2022*)*. While a role for the DLS in motor sequence learning is clear, an unresolved question is what the exact contribution of the DLS is to these behaviours? It has been proposed that the striatum is required for action selection (Park *et al.* 2020),

or even for invigorating movements (Rueda-Orozco & Robbe 2015). I find evidence that both of these are true: in my work, DLS lesions impaired both high-level task structuring and caused movements to be slower. However, I also found that the degree to which task movements were spatio-temporally stereotyped was to some extent dependent on the striatum. This finding corroborates previous work in rodents (Dhawale et al. 2021) and also in songbirds where area X (songbird BG) is known to modulate song production at the level of kinematic structure during learning (Aronov *et al.* 2008; Fee and Goldberg *et al.* 2011). Considering this, my findings fit best with the role of the striatum proposed by Dhawale and colleagues: that the exact form of striatal control depends on the particular challenges or features of motor learning. In this model, if a downstream midbrain or brainstem motor controller for specific movement elements already exists evolutionarily, then the striatum simply learns to sequence these downstream elements. For example, striatal lesions lead to out of sequence grooming, but do not impair the individual grooming motor elements themselves (Berridge & Fentress 1987). If, however, an appropriate controller does not exist for the required behaviour then the striatum learns to exert a richer, kinematic control signal onto downstream circuits. Since I see kinematic changes after striatal lesions my findings support this notion. However, if this is case then this suggests that for movements in my task there were no pre-evolved control circuits which the striatum could direct. This seems unlikely considering the task broadly required postural reorienting movements for which controllers downstream of the BG have been shown to exist; for example, in the superior colliculus (Masullo *et al.* 2019). This paradox is resolved if we consider that a black and white view of the striatum, either specifying kinematic control or action specification, is an oversimplification. Likely both can be true at the same time. Like in the grooming experiment (Berridge & Fentress 1987), my animals with DLS lesions could still produce subsequence movements (port-port transitions). Hence, to some extent the reorienting movements they produced were not kinematically defined by the striatum. However, since

after lesion, these movements were far less stereotyped, it is clear that the striatum also exerted some degree of fine level control. Possibly this control was in reshaping and specifying basic movement blueprints downstream. Another possibility is that this kinematic shaping was independent of these blueprints. The stereotyped behaviour animals produced when completing the task is poorly approximated by spatial reorienting. Movements were in fact extremely rich – down to individual arm hand and tail positioning throughout the task. Striatal kinematic control may instead have been for these auxiliary movements which could have acted to stabilise elementary orienting movements leading to stereotyped trajectories.

In summary, the work in this thesis contributes to our growing understanding of the DLS as a central node in the circuits which define motor sequence learning and control. My work provides further evidence that the contribution of striatum is pliable, such that it is configured in context with the requirements of behaviour. A weakness of this thesis is that I have only considered the striatum in a general sense. Future work must also aim to integrate these ideas with specific features of this circuit. For example, understanding how the striatum contributes to procedural learning will be incomplete without also considering the influence of the striatal microcircuit, specific pathways and nuclei within the BG and the role of neuromodulators such as dopamine. Similarly, the BG do not function alone, future work must also put these ideas in context with wider brain circuitry.

8.3 Thalamic contributions to procedural memory formation

In chapter 5 I investigated the role of thalamus in motor sequence learning. I found evidence that rILT inputs to striatum are required for action-to-action chaining. I also identified a genetic marker for a population which is mostly localised within rILT and found that lesions to this population in rILT impaired task learning by preventing the formation of stereotyped movements.

As discussed earlier, my results are in line with previous findings. Firstly, in rodents it has been shown that permanently silencing striatum projecting cells in the thalamus leads impaired stereotyped movement formation (Wolf *et al.* 2022). Similarly, acute inhibition of thalamo-striatal projections has been shown to impair motor performance in terms of both motor sequence structure and timing (Dias-Hernandez et al. 2018; Hidalgo-Balbuena *et al.* 2019; Melief *et al.* 2018) Secondly, my results are reminiscent of findings in the songbird brain. In this system, thalamic feedback is known to be essential for normal song learning and production (Williams & Vicario 1993, Coleman & Vu 2005; Danish *et al.* 2017) and in recent work the specific role of thalamus has been elucidated (Moll *et al.* 2023). In this study they showed that thalamic activity defined syllables during singing. Thalamic projections tiggered syllables by activating starter cells in HVC and hence, inhibition of these inputs only impaired singing at syllable boundaries. Consequently, similar to implications of my silencing experiments in mice, in the songbird vocal circuit thalamic inputs define the sequential or compositional structure of behaviour. Based on previous evidence that motor cortex disengages across learning (Kawai *et al.* 2015, Wolff *et al.* 2022), I have proposed that as cortex disengages, thalamus takes over as the main driver of striatal activity and that this circuit may function as a feedback loop, routing motor efference information to the striatum to chain actions to actions. These songbird studies are in line with this notion, however, a key difference between the songbird circuit and mammalian circuit mechanism which I have proposed, are the brain regions involved. In the songbird system, cortex is always required for singing while in fact the BG (area X) disengages across learning. More confusingly even some rodent behaviours – such as skilled reaching (Guo *et al.* 2015) – appear to always be dependent on motor cortex. An obvious question arises from this discrepancy: what defines whether a behaviour is always dependent on motor cortex or can be transferred to a subcortical (thalamo-striatal) circuit? A possible answer is that the degree to which cortex can disengage is defined by the ability of the BG to exert full

control over the behaviour. This is likely defined by the availability of suitable downstream motor controllers. A second possibility is that the degree to which behaviour is stereotyped defines the amount of cortical and thalamic control. Certainly, the transfer from cortico-striatal to thalamo-striatal control occurs across learning and so is marked by increasingly stereotyped movements. Further, in both my experiments and previous work (Wolff *et al*. 2022) lesion or silencing of cells in ILT did not impair task completion but caused animals to make movements that were highly variable.

If the transfer from cortical to thalamic control is associated with increasingly stereotyped motor control, then it suggests two further points. Firstly, it's probable that most behaviours sit somewhere on a spectrum between fully cortical and fully thalamic control, this likely varies between animals even within the same task. Secondly this suggests that the functional control thalamus exerts is very different to that of motor cortex. More variable behaviours imply motor cortex may have greater behavioural flexibility. Meanwhile, if thalamic control produces tightly stereotyped movements, this suggests the state information it sends to striatum must also be very stereotyped. The best, and most stereotyped predictors of future movements in an action sequence are the preceding movements. Hence, a strong hypothesis is that thalamus is routing efference copy information into the striatum – feedback about the current action from midbrain and brainstem motor controllers.

As discussed previously, the idea that thalamus might provide efference copy information for behavioural control is not new. Such efference feedback, or 'corollary distance' has been shown previously for saccadic eye movement control (Guthrie *et al.*1983; Sommer and Wutz 2002; Tanaka, 2006). Anatomically, rILT is ideally situated to perform this function as it receives dense ascending innervation from a range of midbrain, brainstem, and cerebellar motor controllers (Yamasaki *et al*. 1986; Cornwall & Philipson 1988; Krout et al. 2002; Bostan & Strick 2018; Mandelbaum *et al* 2019). These

controllers themselves receive direct or near direct input from BG output nuclei, completing this feedback loop. In motor sequence control this loop would function in serial order, by chaining action-to-action. Breaking any part of this loop, should break the chain, arresting behaviour. In my experiments, I found that inhibiting rILT inputs to striatum had this exact effect. If this model is true, then to what extent could this loop control the ongoing temporal dynamics of striatal activity? Possibly it could function continuously, with feedback about fine scale kinematics defining the evolution of striatal activity. Alternately, as has been found in the songbird (Moll *et al*. 2023), thalamo-striatal inputs may exert periodic control, defining the substructure of behaviour. If this later hypothesis is true its not clear how this substructure is defined by thalamus, however, this may come from dense innervation from PFC (Akert & Hartmann 1980; Saalmann, 2014). PFC is known to encode the structure of ongoing behaviour (Marton et al. 2019). Hence, one possibility is that continuous efference feedback information is gated by PFC projections to thalamus. However, if thalamus does define the substructure of behaviour, it's not clear what fills in the gaps in-between. Unlike songbird HVC, striatum is an inhibitory structure and not capable of sustaining chains of internal activity (Lanciego *et al*. 2012). Therefore, the striatum requires continuous glutamatergic drive. Further experimentation is required to answer these questions, though as a final speculation, it is possible that thalamic inputs are a mixture of both. Some inputs may carry continuous state information about movement kinematics while others could fire periodically to reshape the striatal microcircuit and define the syntactic structure of behaviour.

If rILT takes over from motor cortex as the main source of glutamatergic input to DLS this raises a final key question; how could this transition occur at a circuit level? Lesion to the same region of DLS impairs both motor sequence learning and execution (Berridge & Whishaw, 1992; Yin 2010; Wolff *et. al*. 2021) which suggests that this transition is not mediated by a remapping in the striatum. If the same MSN population which learns a motor sequence is

also responsible for its subsequent skilled execution, then this cortical to thalamic transition must be a reorganisation of synapses onto striatal dendrites. Certainly, cortico-striatal and thalamo-striatal inputs are highly convergent in the striatum (Huerta-Ocampo *et al*. 2014). One proposed mechanism is that motor cortical inputs guide thalamo-striatal plasticity in what has been termed synaptic tutoring (Wolff *et. al*. 2022). This mechanism may be analogous to what has been observed in songbird vocal learning, albeit in a rearranged way: in the avian brain striatal outputs (via the lateral magnocellular nucleus of the anterior nidopallium - LMAN) are thought to be the tutor, guiding motor cortical synapses (Andalman & Fee 2009). Plasticity in these motor cortical synapses is defined by concurrent inputs from LMAN via spike-timing dependent heterosynaptic plasticity (Mehaffey *et al*. 2015). In the mammalian brain a similar mechanism is feasible since cortico-striatal and thalamo-striatal inputs are highly overlapping. Hence, precisely timed co-activation of these inputs has been proposed as a way of inducing thalamo-striatal plasticity (Fino & Venance 2011; Perrin & Venance 2019; Wolff *et. al*. 2022). A mechanism based on precisely timed co-activation is highly feasible in this circuit considering ILT and motor cortex are reciprocally connected and have highly synchronised activity (Saalmann 2014).

Another proposed mechanism for how cortical and thalamic inputs may interact to define thalamo-striatal plasticity comes from modelling work by Murray & Escorla, (2020). In this model motor cortical inputs to DLS initially drive motor output via fast, dopamine-dependent plasticity (Kreitzer & Malenka, 2008). At first, thalamo-striatal inputs are silent and do not influence motor output but are increasingly strengthened via slow Hebbian transfer of plasticity. This is mediated by co-activation of small thalamic inputs with their target SPNs which will have been potentiated by stronger cortical input. With repeated expression of the motor behaviour, this gradual Hebbian transfer would strengthen thalamic inputs until they could assume control over the learned behaviour themselves. At this point cortical input would be redundant and cortico-striatal synapses would be weakened.

Possibly thalamic inputs drive earlier SPN activation than their cortical counterparts and so this final synaptic weakening may be via a spike time dependent Hebbian mechanism. This model is particularly compelling since it fits with previously discussed ideas (based on behavioural observations) about thalamic control generating more stereotyped movements. As behaviour becomes increasingly stereotyped, thalamic inputs should become better and better predictors of striatal activity (and cortical inputs worse). Hence, a positive feedback loop should drive synaptic transfer here.
A final possibility is that transfer of synaptic connectivity to thalamic inputs could be via local dopamine modulation by thalamo-striatal interactions with cholinergic interneurons (Goldberg & Reynolds 2011). Since dopamine is required for reinforcement learning and broad dopamine fluctuations in DLS correlate with changes in action-to-action structure (Markowitz *et al*. 2023), thalamic driven dopamine release could facilitate a plasticity bias in the favour of thalamic synapses on MSNs.

In summary, the work I have shown in this thesis contributes to our growing understanding of how thalamo-striatal circuits support procedural memory formation. I find evidence for a pathway via rILT and have suggested this circuit may function, at least in part, by providing efference copy feedback to striatum. I argue that a fruitful avenue for future investigation would be to test the specifics of these ideas, focusing on rILT as a key nexus for motor sequence control. A clear weakness in this work is that I have neglected cortical contributions in favour of investigating subcortical circuits. An important step in our understanding of how procedural memory is formed will be to comprehend how the brain balances between cortical and subcortical motor control. Most likely this balance is dependent on behaviour and circumstance and therefore understanding this would give important insights into the function each system serves to the brain. This is a clear target for future investigation.

8.3 An unsupervised method for replay detection

In chapter 6 I first provided evidence that an offline mechanism involving striatal NDMA dependent plasticity supports procedural memory consolidation. This result replicates recent findings in rats (Lemke *et al.* 2021) and is in line the findings of with multiple previous studies across species (Robertson *et al.* 2004; Fischer et al. 2002; Walker et al. 2002; Korman et al. 2007; Nishida & Walker 2007; Walker et al. 2003-A;Yang et al. 2014;Nagai *et al.* 2017;Ramanathan & Ganguly, 2015). Having established this for my task, I then aimed characterise the mechanism which underpins procedural consolidation. For other kinds of memory, a large body of literature has established that neural reactivations of previously observed activity underpin offline consolidation (Wilson & McNaughton 1994, Skaggs & McNaughton 1996; Lee & Wilson 2002; O'Neill et al. 2010;Lee & Wilson 2002;Foster & Wilson 2006) Hence, my specific aim was to determine whether similar procedural replay occurs within the striatum. As discussed previously, current replay detection methods – which have been developed mostly for hippocampal replay detection – are not well suited to this aim. Instead, I have proposed a novel method based on an unsupervised point process model called PPseq. The use of an unsupervised method for replay detection is a relatively new concept, however, PPseq is not entirely alone. As discussed, several other unsupervised methods have been proposed for, or even applied to replay detection (Maboudi *et al* 2018; Chaudhuri *et al* 2019; Grossberger *et al.* 2018 ;Van der Meij & Voytek 2018). A clear benefit of these approaches is the ability to relate two sets of data in a template free manner. This removes a lot of bias and allows for detection of reactivated neural patterns without making assumptions about (or having prior knowledge of) the coding schema of the population being analysed. Hence, while template-based methods can give a partial snapshot of offline activity, unsupervised methods are better suited to capture the true extent of these dynamics. However, because detected patterns are not predefined by the experimenter, a major problem

with unsupervised detection is reduced interpretability. This issue comes in two forms. Firstly, because there is no predefined template, it is harder to know what behaviour or awake state to ground replay to. Secondly, because there is no ground truth for replay it can be harder to validate whether detected events are meaningful. In other words, though these methods allow us to see a much greater extent of offline dynamics, we are more vulnerable to misinterpreting this activity and have less capacity to scrutinise it for its legitimacy. Still, these issues can be overcome. In this thesis I applied my method onto a dataset with focused behavioural dimensions. Animals performed the same task over multiple repeating trials and so the free parameters in each PPseq model fitted to various aspects of this behaviour. As a result, I could more confidently interpret the specific task relevance of these patterns when observed reactivated offline. Additionally, I tested the method using ground truth data and benchmarked its performance – comparing to a decoder-based approach. This kind of testing is invaluable for assessing the validity of detected replay and for making the abilities, properties, and shortcoming of unsupervised detection less opaque. Future work should aim to test this method even more thoroughly to better understand its capabilities but also to set a standard of methodological rigour. As proposed in a recent review (Tingley & Peyrache 2020), replay detection methods in general suffer from a lack of comparative testing, benchmarking, and standardisation. Since there is no ground truth when observing offline dynamics, in order to make meaningful observations, the methodology we employ must be rigorously tested.

To summarise, in this thesis I have introduced a novel replay detection approach based on an unsupervised method. Unsupervised replay detection can be extremely powerful but requires careful experimental design, caution when drawing conclusions and, like any replay detection approach, warrants rigorous methodological testing.

8.3 Procedural replay in the striatum

In chapter 7, I applied my unsupervised detection approach to periods of post task sleep. I identified striatal replay, and I validated these events: showing my method had a low false positive rate and high correspondence with decoder identified events. Striatal replay shared many features in common with previously identified hippocampal replay suggesting neural reactivations may share common characteristics across brain regions and memory systems. A clear outstanding question from this work is whether or not the replay I have identified is indeed the mechanism which underpins procedural consolidation. This certainly seems likely considering there is strong causal evidence that hippocampal replay has this role for episodic memory (Girardeau et al. 2009; Ego-Stengel & Wilson 2010; Roux et al. 2017). However, though I found evidence that plasticity in the striatum in required for consolidation, I have not provided causal evidence directly linking replay to this consolidation. A clear avenue for future research is to try to establish this link. Taking inspiration from the studies mentioned above, a suitable way to do this would be first identify a biomarker heuristic and then use a closed loop method to inhibit the striatum during reactivation events. If shutting the striatum down specifically during replay impairs consolidation, this will be strong evidence that replay is the mechanistic basis for offline learning.

An interesting finding was that replay rate and characteristics appeared unchanged throughout learning. This suggests that consolidation is a continuous process in the striatum and is not upregulated during learning. Certainly though, too little is known about this mechanism to make any strong conclusions. For example, possibly this observation is a reflection of my behavioural paradigm – not all skills are equal. In less complex motor learning task replay might not be so necessary after the learning stage. For instance, consider the practise required to maintain the ability to play the violin compared to opening a door handle. Another possibility is that my analysis and measurements may have missed crucial details that make early and late

learning replay qualitatively different. Previous evidence suggests nesting between LFP biomarkers might define windows in which consolidation occurs (Lemke *et al*. 2021; Schreiner *et al*. 2021; Kim *et al*. 2019). A possibility which future analysis should be directed towards is that these nestings might change across learning. If this is true though, quite how these LFP nestings define consolidation would still be unclear. Further, this would mean many background replay events occur besides those needed for consolidation. What could this background rate accomplish? Such events could be a homeostatic maintenance of memory (rather than boosting it) or indeed maintenance of circuit properties unrelated to memory but essential for normal circuit function. Ultimately though, the novelty of this research area means that for now these questions remain highly speculative. However, a clear outcome from this work is that this is a topic for which there is enormous scope for future inquiry.

One feature of the striatal replay which I identified, which agrees with very recent hippocampal findings (Denovellis *et al*. 2021) but is in contrast to more canonical notions of replay, is that many identified events were not simply sequential reactivations. Replay was often fragmented. A key question is how such events might contribute to consolidation? An interesting possibility that arises from this finding is that the current conventional understanding of replay is incomplete. Contemporary theories posit that sequential reactivations engrave awake activity into plasticity as like-for-like offline practice. However, with more sophisticated detection methods evidence is mounting that replay is far less clear cut than this model posits. Consolidation via replay may in fact be the result of a highly fragmented generative process. In this kind of process plasticity would be steadily shaped by noisy, irregular replay events that follow the statistics of awake activity without nessisairly recapitulating them. Similar ideas have also been proposed in hippocampal replay (Stella *et al*. 2019).

Another possibility is that the diversity observed in replay reflects a spectrum of events with different functional roles for consolidation. Replay might

strengthen or weaken memories or even represent exploration or generalisation across different memories. A key avenue for future research will be to try to assign a functional syntax to the characteristics of replay. This is a difficult challenge. However, a promising recent study suggests that different LFP nesting types contains replay which either strengthens or weakens memories (Kim *et al*. 2019). If the characteristics of replay do define the function of that replay, then analysis of these characteristics during different LFP nestings could reveal this relationship.

A final remaining question is where the glutamatergic drive for striatal replay originates. As discussed in previous chapters, the striatum is a non-spontaneously active inhibitory structure, (Lanciego *et al*. 2012). Hence, striatal activity is the result of excitatory input. Just as discussed for awake activity, there are two main candidates for this: cortical and thalamic inputs (Hunnicutt *et al*. 2016). Given the evidence discussed earlier a clear hypothesis is that the degree to which each area contributes may depend on the learning stage of the animal – becoming increasingly thalamic across learning. Other scenarios are also feasible. For example, cortical inputs may still actively drive replay during late stages of learning, possibly though this activity could have a more disruptive role – driving increased flexibility and preventing overfitting. Once more, these speculative ideas highlight that this is an area which requires future investigation. From my own recordings I have access to motor cortical activity. Analysis correlating cortical and thalamic reactivations could provide some insights for this purpose. Additionally, useful future experiments could aim to silence each input independently during sleep and observe whether striatal replay, or certain replay characteristics are perturbed.

Overall, this thesis has contributed the novel finding that striatal reactivations of procedural activity occur during sleep. These findings parallel a great deal of previous work investigating hippocampal and cortical replay and provide a possible mechanistic basis for offline consolidation observed during procedural memory. My findings support the notion that more sophisticated

detection methods such as the unsupervised approach I have employed here are required to properly understand the extent of offline dynamics. Another clear outcome from this work is there is large scope for future investigation in this area. One key area for future research will be to link observations about procedural replay to the functional roles these events have in memory consolidation.

8.4 No role for hippocampus in procedural consolidation or replay

Finally, in chapter 7 I demonstrated that procedural memory consolidation is not dependent on the hippocampus. This finding is in line with multiple previous studies (Reber & Squire 1998; Hopkins *et al* 2004,Squire, 2009; Jackson & Strong 1969; Eckart *et al*. 2011; Will *et al*. 2013; Busse & Schwarting 2016[A]; Busse & Schwarting 2016[B]; Schwarting & Busse 2017) and supports the notion that parallel – and non-overlapping – memory systems may exist in the brain (White & McDonald, 2002). However, if these systems are entirely independent, then a key question is how global memory systems function without interaction. Certainly, these systems share features which seem to link them together. For example, both procedural (Fang *et al*. 2020; Dickey *et al*. 2021; Lemke *et al*. 2021; Schreiner *et al*. 2021) and episodic (Latchoumane *et al*. 2017; Maingret *et al*. 2016) consolidation appears to be choreographed in relation to slow wave and spindle LFP oscillations. Previously, it been thought that the hippocampus might be the driver of consolidation, triggering and orchestrating replay across the brain in order to transfer memory into long term storage (Buzsaki 1989, Siapas & Wilson 1998; Klinzing *et al*. 2019; Spens & Burgess 2023). However, since procedural consolidation persists without the hippocampus, this notion cannot be entirely correct. One way to reconcile these findings, in line with multiple proposed roles of the hippocampus (reviewed in Ekstrom & Yonelinas 2020), is to reconsider the hippocampus as a conductor. Its role during consolidation may be to bind disparate brain regions (like musicians in an

orchestra) together into harmonious and synchronised ensembles. Such binding may be the foundation of rich multifaceted memories of experience (episodic memory). However, in the absence of this conductor, individual players still play their own instruments but do so at their own tempo. Hence, local dynamics in procedural circuits would act normally from the perspective of procedural memory formation and only episodic like memories would be impaired. If this is true, this generates predictions which can be tested in future work. A strong hypothesis from these ideas is that while simple procedural memory should be unimpaired, animals with hippocampus lesions should have deficits linking motor skills to other memory modalities such as a contextual trigger. For example, without the hippocampus to bind disparate circuits together via global synchronisation, animals should struggle to link different motor sequences to different odours or tones.

8.5 Conclusions

In this thesis I set out to investigate roles for subcortical cortical circuits in supporting procedural memory formation. I approached this topic by considering two distinct but highly overlapping themes. Firstly, I have investigated circuits that support skill learning and production online: during wakeful behaviour. I have demonstrated that learning and execution of a novel motor skill task are contingent on the DLS and provided evidence that supports the notion that this region supports procedural motor control in a multifaceted and behaviour dependent manner. I have also provided evidence that thalamic projections from the rostral portion of intralaminar thalamus support this function of striatum. I have shown evidence that this thalamic pathway is involved in both high-level behavioural structuring as well as being required to form stereotyped movements. I argue that rILT has a privileged position integrating motor efference feedback from midbrain and brainstem controllers with higher order structural information from frontal

cortical regions. This anatomy might make this region a key nexus in the circuit that learns and controls procedural memories.

In the second part of this thesis, I have examined the processes which support the function of these circuits offline: during rest or sleep. I have shown evidence that an offline process supports procedural memory and that this is dependent on striatal plasticity. To search for the content of this offline process I have outlined a novel replay detection method based on an unsupervised point process model. I have validated this method with ground truth data and shown that when applied to sleep, it can identify neural replay of task related activity. I have argued that extensive methodological testing and unbiased approaches such as the one described, are essential for understanding the true extent of offline dynamics. From analysis of detected striatal replay, I find this activity shared many features in common with previously identified hippocampal replay and persisted throughout all stages of motor skill learning. Procedural replay represents a novel research direction and I have argued that this is an area with enormous scope for future investigation. Finally, I have shown that both procedural consolidation and striatal replay are independent of the hippocampus. This suggests that procedural memory consolidation is an entirely distinct process from other kinds of neural consolidation.

REFERENCES

1. Adams, J.A. **(1984).** 'Learning of movement sequences.'. *Psychological Bulletin*, 96 (1), p.3.
2. Agostino, R., Berardelli, A., Formica, A., Accornero, N. and Manfredi, M. **(1992).** 'Sequential arm movements in patients with parkinson's disease, huntington's disease and dystonia'. *Brain*, 115 (5), pp.1481–1495.
3. Akert, K. and Hartmann-von Monakow, K. **(1980).** 'Relationships of precentral premotor and prefrontal cortex to the mediodorsal and intralaminar nuclei of the monkey thalamus'. *Acta Neurobiologiae Experimentalis*, 40 (1), pp.7–25.
4. Alexander, G.E. and Crutcher, M.D. **(1990).** 'Functional architecture of basal ganglia circuits: neural substrates of parallel processing'. *Trends in Neurosciences*, 13 (7), pp.266–271.
5. Alonso, R.G., Trevisan, M.A., Amador, A., Goller, F. and Mindlin, G.B. (**2015**). 'A circular model for song motor control in Serinus canaria'. *Frontiers in Computational Neuroscience*, 9.
6. Ambrose, R.E., Pfeiffer, B.E. and Foster, D.J. (**2016**). 'Reverse replay of hippocampal place cells is uniquely modulated by changing reward'. *Neuron*, 91 (5), pp.1124–1136.
7. Andalman, A.S. and Fee, M.S. (**2009**). 'A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors'. *Proceedings of the National Academy of Sciences*, 106 (30), pp.12518–12523.
8. Andersen, K.W., Madsen, K.H. and Siebner, H.R. (**2020**). 'Discrete fnger sequences are widely represented in human striatum'. *Scientific Reports*, 10 (13189).
9. Igata, H. Ikegaya, Y. and Sasaki, T. (**2020**) 'Prioritized experience replays on a hippocampal predictive map for learning | PNAS'
10. Sutton, R. S. & Barto, A. G. Reinforcement Learning: An Introduction (Bradford Books, **2018**).
11. Aronov, D., Andalman, A.S. and Fee, M.S. (**2008**). 'A specialized forebrain circuit for vocal babbling in the juvenile songbird'. *Science (New York, N.Y.)*, 320 (5876), pp.630–634.
12. Ashmore, R.C., Wild, J.M. and Schmidt, M.F. (**2005**). 'Brainstem and Forebrain Contributions to the Generation of Learned Motor Behaviors for Song'. *Journal of Neuroscience*, 25 (37), pp.8543–8554.
13. Averbeck, B.B., Sohn, J.-W. and Lee, D. (**2006**). 'Activity in prefrontal cortex during dynamic selection of action sequences'. *Nature Neuroscience*, 9 (2), pp.276–282.
14. Badre, D. and Nee, D.E. (**2018**). 'Frontal cortex and the hierarchical control of behavior'. *Trends in cognitive sciences*, 22 (2), pp.170–188.
15. Bapi, R.S., Doya, K. and Harner, A.M. (**2000**). 'Evidence for effector independent and dependent representations and their differential time course of acquisition during motor sequence learning'. *Experimental Brain Research*, 132 (2), pp.149–162.
16. Barakat, M., Carrier, J., Debas, K., Lungu, O., Fogel, S., Vandewalle, G., Hoge, R.D., Bellec, P., Karni, A., Ungerleider, L.G., Benali, H. and Doyon, J. (**2012**). 'Sleep spindles predict neural and behavioral changes in motor sequence consolidation: Sleep Spindles Predict Motor Consolidation'. *Human Brain Mapping*, 34 (11), pp.2918–2928.
17. Baraldi, M., Grandison, L. and Guidotti, A. (**1979**). 'Distribution and metabolism of muscimol in the brain and other tissues of the rat'. *Neuropharmacology*, 18 (1), pp.57–62.
18. Barbera, G., Liang, B., Zhang, L., Gerfen, C.R., Culurciello, E., Chen, R., Li, Y. and Lin, D.-T. (**2016**). 'Spatially Compact Neural Clusters in the Dorsal Striatum Encode Locomotion Relevant Information'. *Neuron*, 92 (1), pp.202–213.
19. Berridge, K.C. and Fentress, J.C. (**1987**). 'Disruption of natural grooming chains after striatopallidal lesions'. *Psychobiology*, 15, pp.336–342.
20. Berridge, K.C., Fentress, J.C. and Parr, H. (**1987**). 'Natural syntax rules control action sequence of rats'. *Behavioural Brain Research*, 23 (1), pp.59–68.

21. Berridge, K.C. and Whishaw, I.Q. (**1992**). 'Cortex, striatum and cerebellum: control of serial order in a grooming sequence'. *Experimental Brain Research*, 90 (2), pp.275–290.
22. Bhatia, K.P. and Marsden, C.D. (**1994**). 'The behavioural and motor consequences of focal lesions of the basal ganglia in man'. *Brain: A Journal of Neurology*, 117, pp.859–876.
23. Bhattarai, B., Lee, J.W. and Jung, M.W. (**2020**). 'Distinct effects of reward and navigation history on hippocampal forward and reverse replays'. *Proceedings of the National Academy of Sciences*, 117 (1), pp.689–697.
24. Bostan, A.C. and Strick, P.L. (**2018**). 'The basal ganglia and the cerebellum: nodes in an integrated network'. *Nature Reviews. Neuroscience*, 19 (6), pp.338–350.
25. Busse, S. and Schwarting, R.K.W. (**2016a**). 'Decoupling Actions from Consequences: Dorsal Hippocampal Lesions Facilitate Instrumental Performance, but Impair Behavioral Flexibility in Rats'. *Frontiers in Behavioral Neuroscience*, 10.
26. Busse, S. and Schwarting, R.K.W. (**2016b**). 'Procedural Performance Benefits after Excitotoxic Hippocampal Lesions in the Rat Sequential Reaction Time Task'. *Neurotoxicity Research*, 29 (1), pp.54–68.
27. Buysse, D.J., Reynolds, C.F., Monk, T.H., Berman, S.R. and Kupfer, D.J. (**1989**). 'The Pittsburgh sleep quality index: A new instrument for psychiatric practice and research'. *Psychiatry Research*, 28 (2), pp.193–213.
28. Buzsáki, G. (**2015**). 'Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning'. *Hippocampus*, 25 (10), pp.1073–1188.
29. Buzsáki, G. (**1989**). 'Two-stage model of memory trace formation: A role for "noisy" brain states'. *Neuroscience*, 31 (3), pp.551–570.
30. Campbell, S.S. and Tobler, I. (**1984**). 'Animal sleep: a review of sleep duration across phylogeny'. *Neuroscience and Biobehavioral Reviews*, 8 (3), pp.269–300.
31. Carr, M.F., Karlsson, M.P. and Frank, L.M. (**2012**). 'Transient slow gamma synchrony underlies hippocampal memory replay'. *Neuron*, 75 (4), pp.700–713.
32. Catsman-Berrevoets, C.E. and Kuypers, H.G.J.M. (**1978**). 'Differential laminar distribution of corticothalamic neurons projecting to the VL and the center median. An HRP study in the cynomolgus monkey'. *Brain Research*, 154 (2), pp.359–365.
33. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. and Fiete, I. (**2019**). 'The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep'. *Nature Neuroscience*, 22 (9), pp.1512–1520.
34. Christie, M.A. (**2004**). 'A New Rat Model of the Human Serial Reaction Time Task: Contrasting Effects of Caudate and Hippocampal Lesions'. *Journal of Neuroscience*, 24 (5), pp.1034–1039.
35. Coleman, M.J. and Vu, E.T. (**2005**). 'Recovery of impaired songs following unilateral but not bilateral lesions of nucleus uvaeformis of adult zebra finches'. *Journal of Neurobiology*, 63 (1), pp.70–89.
36. Cornwall, J. and Phillipson, O.T. (**1988**). 'Afferent projections to the parafascicular thalamic nucleus of the rat, as shown by the retrograde transport of wheat germ agglutinin'. *Brain Research Bulletin*, 20 (2), pp.139–150.
37. Costa, R.M., Cohen, D. and Nicolelis, M.A.L. (**2004**). 'Differential corticostriatal plasticity during fast and slow motor skill learning in mice'. *Current biology: CB*, 14 (13), pp.1124–1134.
38. Cousins, J.N., El-Deredy, W., Parkes, L.M., Hennies, N. and Lewis, P.A. (**2014**). 'Cued Memory Reactivation during Slow-Wave Sleep Promotes Explicit Knowledge of a Motor Sequence'. *Journal of Neuroscience*, 34 (48), pp.15870–15876.
39. Cox, J. and Witten, I.B. (**2019**). 'Striatal circuits for reward learning and decision-making'. *Nature reviews. Neuroscience*, 20 (8), pp.482–494.
40. Danish, H.H., Aronov, D. and Fee, M.S. (**2017**). 'Rhythmic syllable-related activity in a songbird motor thalamic nucleus necessary for learned vocalizations'. *PLOS ONE*, 12 (6), p.e0169568.
41. Davidson, T.J., Kloosterman, F. and Wilson, M.A. (**2009**). 'Hippocampal Replay of Extended Experience'. *Neuron*, 63 (4), pp.497–507.
42. Daw, N.D., Niv, Y. and Dayan, P. (**2005**). 'Actions, Policies, Values, and the Basal Ganglia'. , p.21.

43. Denovellis, E.L., Gillespie, A.K., Coulter, M.E., Sosa, M., Chung, J.E., Eden, U.T. and Frank, L.M. (**2021**). 'Hippocampal replay of experience at real-world speeds' A. Peyrache, T. E. Behrens, Y. Liu, and H. F. Ólafsdóttir (Eds.),. *eLife*, 10, p.e64505.
44. Desmurget, M. and Turner, R.S. (**2010**). 'Motor Sequences and the Basal Ganglia: Kinematics, Not Habits'. *Journal of Neuroscience*, 30 (22), pp.7685–7690.
45. Dezfouli, A. and Balleine, B.W. (**2012**). 'Habits, action sequences and reinforcement learning'. *The European Journal of Neuroscience*, 35 (7), pp.1036–1051.
46. Dhawale, A.K., Wolff, S.B.E., Ko, R. and Ölveczky, B.P. (**2021**). 'The basal ganglia control the detailed kinematics of learned motor skills'. *Nature Neuroscience*, 24 (9), pp.1256–1269.
47. Díaz-Hernández, E., Contreras-López, R., Sánchez-Fuentes, A., Rodríguez-Sibrían, L., Ramírez-Jarquín, J.O. and Tecuapetla, F. (**2018**). 'The Thalamostriatal Projections Contribute to the Initiation and Execution of a Sequence of Movements'. *Neuron*, 100 (3), pp.739-752.e5.
48. Diba, K. and Buzsáki, G. (**2007**). 'Forward and reverse hippocampal place-cell sequences during ripples'. *Nature Neuroscience*, 10 (10), pp.1241–1242.
49. Dickey, C.W., Sargsyan, A., Madsen, J.R., Eskandar, E.N., Cash, S.S. and Halgren, E. (**2021**). 'Travelling spindles create necessary conditions for spike-timing-dependent plasticity in humans'. *Nature Communications*, 12, p.1027.
50. Diekelmann, S. and Born, J. (**2010**). 'The memory function of sleep'. *Nature Reviews Neuroscience*, 11 (2), pp.114–126.
51. Ding, J., Peterson, J.D. and Surmeier, D.J. (**2008**). 'Corticostriatal and Thalamostriatal Synapses Have Distinctive Properties'. *The Journal of Neuroscience*, 28 (25), pp.6483–6492.
52. Douglas, R.J. and Pribram, K.H. (**1969**). 'Distraction and habituation in monkeys with limbic lesions'. *Journal of Comparative and Physiological Psychology*, 69, pp.473–480.
53. Dupret, D., O'Neill, J., Pleydell-Bouverie, B. and Csicsvari, J. (**2010**). 'The reorganization and reactivation of hippocampal maps predict spatial memory performance'. *Nature neuroscience*, 13 (8), pp.995–1002.
54. Durán, E., Oyanedel, C.N., Niethard, N., Inostroza, M. and Born, J. (**2018**). 'Sleep stage dynamics in neocortex and hippocampus'. *Sleep*, 41 (6), p.zsy060.
55. Eckart, M. t., Huelse-Matia, M. c. and Schwarting, R. k. w. (**2012**). 'Dorsal hippocampal lesions boost performance in the rat sequential reaction time task'. *Hippocampus*, 22 (5), pp.1202–1214.
56. Ego-Stengel, V. and Wilson, M.A. (**2010**). 'Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat'. *Hippocampus*, 20 (1), pp.1–10.
57. Eichenlaub, J.-B., Jarosiewicz, B., Saab, J., Franco, B., Kelemen, J., Halgren, E., Hochberg, L.R. and Cash, S.S. (**2020**). 'Replay of Learned Neural Firing Sequences during Rest in Human Motor Cortex'. *Cell Reports*, 31 (5), p.107581.
58. Ekstrom, A.D. and Yonelinas, A.P. (**2020**). 'Precision, binding, and the hippocampus: Precisely what are we talking about?'. *Neuropsychologia*, 138, p.107341.
59. Ellender, T.J., Harwood, J., Kosillo, P., Capogna, M. and Bolam, J.P. (**2013**). 'Heterogeneous properties of central lateral and parafascicular thalamic synapses in the striatum'. *The Journal of Physiology*, 591 (1), pp.257–272.
60. Elmaleh, M., Kranz, D., Asensio, A.C., Moll, F.W. and Long, M.A. (**2021a**). 'Sleep replay reveals premotor circuit structure for a skilled behavior'. *Neuron*, 109 (23), pp.3851-3861.e4.
61. Elmaleh, M., Kranz, D., Asensio, A.C., Moll, F.W. and Long, M.A. (**2021b**). 'Sleep replay reveals premotor circuit structure for a skilled behavior'. *Neuron*, 109 (23), pp.3851-3861.e4.
62. Ericsson, K.A., Krampe, R.T. and Tesch-Romer, C. (**1993**). 'The Role of Deliberate Practice in the Acquisition of Expert Performance'. , p.44.

63. Fang, Z., Ray, L.B., Houldin, E., Smith, D., Owen, A.M. and Fogel, S.M. (**2020**). 'Sleep Spindle-dependent Functional Connectivity Correlates with Cognitive Abilities'. *Journal of Cognitive Neuroscience*, 32 (3), pp.446–466.
64. Fee, M.S. and Goldberg, J.H. (**2011**). 'A hypothesis for basal ganglia-dependent reinforcement learning in the songbird'. *Neuroscience*, 198, pp.152–170.
65. Fino, E. and Venance, L. (**2011**). 'Spike-timing dependent plasticity in striatal interneurons'. *Neuropharmacology*, 60 (5), pp.780–788.
66. Fiorillo, C.D., Tobler, P.N. and Schultz, W. (**2003**). 'Discrete coding of reward probability and uncertainty by dopamine neurons'. *Science (New York, N.Y.)*, 299 (5614), pp.1898–1902.
67. Fischer, S. and Born, J. (**2009**). 'Anticipated reward enhances offline learning during sleep'. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 35 (6), pp.1586–1593.
68. Fischer, S., Hallschmid, M., Elsner, A.L. and Born, J. (**2002**). 'Sleep forms memory for finger skills'. *Proceedings of the National Academy of Sciences of the United States of America*, 99 (18), pp.11987–11991.
69. Fogel, S., Albouy, G., King, B.R., Lungu, O., Vien, C., Bore, A., Pinsard, B., Benali, H., Carrier, J. and Doyon, J. (**2017**). 'Reactivation or transformation? Motor memory consolidation associated with cerebral activation time-locked to sleep spindles'. *PLOS ONE*, 12 (4), p.e0174755.
70. Fogel, S.M. and Smith, C.T. (**2006**). 'Learning-dependent changes in sleep spindles and Stage 2 sleep'. *Journal of Sleep Research*, 15 (3), pp.250–255.
71. Forest, G. and Godbout, R. (**2000**). 'Effects of sleep deprivation on performance and EEG spectral analysis in young adults'. *Brain and Cognition*, 43 (1–3), pp.195–200.
72. Foster, D.J. and Wilson, M.A. (**2006**). 'Reverse replay of behavioural sequences in hippocampal place cells during the awake state'. *Nature*, 440 (7084), pp.680–683.
73. Fox, K.C.R., Nijeboer, S., Solomonova, E., Domhoff, G.W. and Christoff, K. (**2013**). 'Dreaming as mind wandering: evidence from functional neuroimaging and first-person content reports'. *Frontiers in Human Neuroscience*, 7, p.412.
74. Frank, M.J. (**2011**). 'Computational models of motivated action selection in corticostriatal circuits'. *Current Opinion in Neurobiology*, 21 (3), pp.381–386.
75. Freeman, J.S., Cody, F.W., O'Boyle, D.J., Craufurd, D., Neary, D. and Snowden, J.S. (**1996**). 'Abnormalities of motor timing in Huntington's disease'. *Parkinsonism & Related Disorders*, 2 (2), pp.81–93.
76. Geddes, C.E., Li, H. and Jin, X. (**2018**). 'Optogenetic Editing Reveals the Hierarchical Organization of Learned Action Sequences'. *Cell*, 174 (1), pp.32-43.e15.
77. Geva-Sagiv, M. and Nir, Y. (**2019**). 'Local Sleep Oscillations: Implications for Memory Consolidation'. *Frontiers in Neuroscience*, 13, p.813.
78. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G. and Zugaro, M.B. (**2009**). 'Selective suppression of hippocampal ripples impairs spatial memory'. *Nature Neuroscience*, 12 (10), pp.1222–1223.
79. Goldberg, J.A. and Reynolds, J.N.J. (**2011**). 'Spontaneous firing and evoked pauses in the tonically active cholinergic interneurons of the striatum'. *Neuroscience*, 198, pp.27–43.
80. Goto, A. and Hayashi, Y. (**2023**). 'Offline neuronal activity and synaptic plasticity during sleep and memory consolidation'. *Neuroscience Research*, 189, pp.29–36.
81. Graybiel, A.M. (**1998**). 'The basal ganglia and chunking of action repertoires'. *Neurobiology of Learning and Memory*, 70 (1–2), pp.119–136.
82. Grossberger, L., Battaglia, F.P. and Vinck, M. (**2018**). 'Unsupervised clustering of temporal patterns in high-dimensional neuronal ensembles using a novel dissimilarity measure'. *PLoS computational biology*, 14 (7), p.e1006283.
83. Gulati, T., Guo, L., Ramanathan, D.S., Bodepudi, A. and Ganguly, K. (**2017**). 'Neural reactivations during sleep determine network credit assignment'. *Nature Neuroscience*, 20 (9), pp.1277–1284.

84. Guo, J.-Z., Graves, A.R., Guo, W.W., Zheng, J., Lee, A., Rodríguez-González, J., Li, N., Macklin, J.J., Phillips, J.W., Mensh, B.D., Branson, K. and Hantman, A.W. (**2015**). 'Cortex commands the performance of skilled movement' M. Hausser (Ed.),. *eLife*, 4, p.e10774.
85. Gupta, A.S., Van Der Meer, M.A.A., Touretzky, D.S. and Redish, A.D. (**2010**). 'Hippocampal Replay Is Not a Simple Function of Experience'. *Neuron*, 65 (5), pp.695–705.
86. Guthrie, B.L., Porter, J.D. and Sparks, D.L. (**1983**). 'Corollary discharge provides accurate eye position information to the oculomotor system'. *Science (New York, N.Y.)*, 221 (4616), pp.1193–1195.
87. Haber, S.N. (**2016**). 'Corticostriatal circuitry'. *Dialogues in Clinical Neuroscience*, 18 (1), pp.7–21.
88. Haenraets, K., Foster, E., Johannssen, H., Kandra, V., Frezel, N., Steffen, T., Jaramillo, V., Paterna, J.-C., Zeilhofer, H.U. and Wildner, H. (**2017**). 'Spinal nociceptive circuit analysis with recombinant adeno-associated viruses: the impact of serotypes and promoters'. *Journal of Neurochemistry*, 142 (5), pp.721–733.
89. Hidalgo-Balbuena, A.E., Luma, A.Y., Pimentel-Farfan, A.K., Peña-Rangel, T. and Rueda-Orozco, P.E. (**2019**). 'Sensory representations in the striatum provide a temporal reference for learning and executing motor habits'. *Nature Communications*, 10 (1), p.4074.
90. Hikosaka, O., Nakahara, H., Rand, M.K., Sakai, K., Lu, X., Nakamura, K., Miyachi, S. and Doya, K. (**1999**). 'Parallel neural networks for learning sequential procedures'. *Trends in Neurosciences*, 22 (10), pp.464–471.
91. Hikosaka, O., Rand, M.K., Miyachi, S. and Miyashita, K. (**1995**). 'Learning of sequential movements in the monkey: process of learning and retention of memory'. *Journal of Neurophysiology*, 74 (4), pp.1652–1661.
92. Hikosaka, O. and Wurtz, R.H. (**1983**). 'Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades'. *Journal of Neurophysiology*, 49 (5), pp.1230–1253.
93. Holz, J., Piosczyk, H., Landmann, N., Feige, B., Spiegelhalder, K., Riemann, D., Nissen, C. and Voderholzer, U. (**2012**). 'The Timing of Learning before Night-Time Sleep Differentially Affects Declarative and Procedural Long-Term Memory Consolidation in Adolescents'. *PLOS ONE*, 7 (7), p.e40963.
94. Hopkins, R.O., Waldram, K. and Kesner, R.P. (**2004**). 'Sequences assessed by declarative and procedural tests of memory in amnesic patients with hippocampal damage'. *Neuropsychologia*, 42 (14), pp.1877–1886.
95. Huerta-Ocampo, I., Mena-Segovia, J. and Bolam, J.P. (**2014**). 'Convergence of cortical and thalamic input to direct and indirect pathway medium spiny neurons in the striatum'. *Brain Structure and Function*, 219 (5), pp.1787–1800.
96. Hunnicutt, B.J., Jongbloets, B.C., Birdsong, W.T., Gertz, K.J., Zhong, H. and Mao, T. (**2016**). 'A comprehensive excitatory input map of the striatum reveals novel functional organization' D. C. Van Essen (Ed.),. *eLife*, 5, p.e19103.
97. Hwang, E.J., Dahlen, J.E., Hu, Y.Y., Aguilar, K., Yu, B., Mukundan, M., Mitani, A. and Komiyama, T. (**2019**). 'Disengagement of motor cortex from movement control during long-term learning'. Science Advances, 5 (10), p.eaay0001.
98. Hwang, E.J., Dahlen, J.E., Mukundan, M. and Komiyama, T. (**2021**). 'Disengagement of Motor Cortex during Long-Term Learning Tracks the Performance Level of Learned Movements'. Journal of Neuroscience, 41 (33), pp.7029–7047.
99. Inostroza, M. and Born, J. (**2013**). 'Sleep for Preserving and Transforming Episodic Memory'. Annual Review of Neuroscience, 36 (1), pp.79–102.
100. Jackson, G.M., Jackson, S.R., Harrison, J., Henderson, L. and Kennard, C. (**1995**). 'Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit'. Neuropsychologia, 33 (5), pp.577–593.
101. Jackson, W.J. and Strong Jr., P.N. (**1969**). 'Differential effects of hippocampal lesions upon sequential tasks and maze learning by the rat'. Journal of Comparative and Physiological Psychology, 68 (3), pp.442–450.
102. Jacobacci, F., Armony, J.L., Yeffal, A., Lerner, G., Amaro, E., Jovicich, J., Doyon, J. and Della-Maggiore, V. (**2020**). 'Rapid hippocampal plasticity supports motor sequence learning'. Proceedings of the National Academy of Sciences, 117 (38), pp.23898–23903.

103. Jensen, K.T., Kadmon Harpaz, N., Dhawale, A.K., Wolff, S.B.E. and Ölveczky, B.P. (**2022**). 'Long-term stability of single neuron activity in the motor system'. Nature Neuroscience, 25 (12), pp.1664–1674.
104. Jin, X. and Costa, R.M. (**2015**). 'Shaping action sequences in basal ganglia circuits'. Current Opinion in Neurobiology, 33, pp.188–196.
105. Jin, X., Tecuapetla, F. and Costa, R.M. (**2014**). 'Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences'. Nature Neuroscience, 17 (3), pp.423–430.
106. Johansson, Y. and Silberberg, G. (**2020**). 'The Functional Organization of Cortical and Thalamic Inputs onto Five Types of Striatal Neurons Is Determined by Source and Target Cell Identities'. Cell Reports, 30 (4), pp.1178-1194.e3.
107. Josselyn, S.A. and Tonegawa, S. (**2020**). 'Memory engrams: Recalling the past and imagining the future'. Science (New York, N.Y.), 367 (6473), p.eaaw4325.
108. Jurado-Parras, M.-T., Safaie, M., Sarno, S., Louis, J., Karoutchi, C., Berret, B. and Robbe, D. (**2020**). 'The Dorsal Striatum Energizes Motor Routines'. Current Biology, 30 (22), pp.4362-4372.e6.
109. Kami, A., Meyer, G., Jezzard, P., Adams, M.M., Turner, R. and Ungerleider, L.G. (**1995**). 'Functional MRI evidence for adult motor cortex plasticity during motor skill learning'. Nature, 377 (6545), pp.155–158.
110. Kato, S., Fukabori, R., Nishizawa, K., Okada, K., Yoshioka, N., Sugawara, M., Maejima, Y., Shimomura, K., Okamoto, M., Eifuku, S. and Kobayashi, K. (**2018**). 'Action Selection and Flexible Switching Controlled by the Intralaminar Thalamic Neurons'. Cell Reports, 22 (9), pp.2370–2382.
111. Kato, S., Nishizawa, K. and Kobayashi, K. (**2021**). 'Thalamostriatal System Controls the Acquisition, Performance, and Flexibility of Learning Behavior'. Frontiers in Systems Neuroscience, 15, p.729389.
112. Kawai, R., Markman, T., Poddar, R., Ko, R., Fantana, A.L., Dhawale, A.K., Kampff, A.R. and Ölveczky, B.P. (**2015**). 'Motor cortex is required for learning but not for executing a motor skill'. Neuron, 86 (3), pp.800–812.
113. Kim, J., Gulati, T. and Ganguly, K. (**2019**). 'Competing Roles of Slow Oscillations and Delta Waves in Memory Consolidation versus Forgetting'. Cell, 179 (2), pp.514-526.e13.
114. Kim, J., Joshi, A., Frank, L. and Ganguly, K. (**2023**). 'Cortical–hippocampal coupling during manifold exploration in motor cortex'. Nature, 613 (7942), pp.103–110.
115. Klinzing, J.G., Niethard, N. and Born, J. (**2019**). 'Mechanisms of systems memory consolidation during sleep'. Nature Neuroscience, 22 (10), pp.1598–1610.
116. Korman, M., Doyon, J., Doljansky, J., Carrier, J., Dagan, Y. and Karni, A. (**2007**). 'Daytime sleep condenses the time course of motor memory consolidation'. Nature Neuroscience, 10 (9), pp.1206–1213.
117. Kornfeld, J., Januszewski, M., Schubert, P., Jain, V., Denk, W. and Fee, M.S. (**2020**). 'An anatomical substrate of credit assignment in reinforcement learning'. , p.2020.02.18.954354.
118. Kornysheva, K., Bush, D., Meyer, S.S., Sadnicka, A., Barnes, G. and Burgess, N. (**2019**). 'Neural Competitive Queuing of Ordinal Structure Underlies Skilled Sequential Action'. Neuron, 101 (6), pp.1166-1180.e3.
119. Krause, A.J., Simon, E.B., Mander, B.A., Greer, S.M., Saletin, J.M., Goldstein-Piekarski, A.N. and Walker, M.P. (**2017**). 'The sleep-deprived human brain'. Nature Reviews Neuroscience, 18 (7), pp.404–418.
120. Kreitzer, A.C. and Malenka, R.C. (**2008**). 'Striatal plasticity and basal ganglia circuit function'. Neuron, 60 (4), pp.543–554.
121. Krout, K.E., Belzer, R.E. and Loewy, A.D. (**2002**). 'Brainstem projections to midline and intralaminar thalamic nuclei of the rat'. Journal of Comparative Neurology, 448 (1), pp.53–101.
122. Kudrimoti, H.S., Barnes, C.A. and McNaughton, B.L. (**1999**). 'Reactivation of Hippocampal Cell Assemblies: Effects of Behavioral State, Experience, and EEG Dynamics'. Journal of Neuroscience, 19 (10), pp.4090–4101.
123. Kupferschmidt, D.A., Juczewski, K., Cui, G., Johnson, K.A. and Lovinger, D.M. (**2017**). 'Parallel, but Dissociable, Processing in Discrete Corticostriatal Inputs Encodes Skill Learning'. Neuron, 96 (2), pp.476-489.e5.
124. Kuriyama, K., Stickgold, R. and Walker, M. (**2004**). 'Sleep-dependent learning and motor-skill complexity'. Learning & memory (Cold Spring Harbor, N.Y.), 11, pp.705–13.

125. Kurth-Nelson, Z., Behrens, T., Wayne, G., Miller, K., Luettgau, L., Dolan, R., Liu, Y. and Schwartenbeck, P. (**2023**). 'Replay and compositional computation'. Neuron, 111 (4), pp.454–469.
126. Kurth-Nelson, Z., Economides, M., Dolan, R.J. and Dayan, P. (**2016**). 'Fast Sequences of Non-spatial State Representations in Humans'. Neuron, 91 (1), pp.194–204.
127. Lacey, C.J., Bolam, J.P. and Magill, P.J. (**2007**). 'Novel and Distinct Operational Principles of Intralaminar Thalamic Neurons and Their Striatal Projections'. Journal of Neuroscience, 27 (16), pp.4374–4384.
128. Laforce, R. and Doyon, J. (**2001**). 'Distinct Contribution of the Striatum and Cerebellum to Motor Learning'. Brain and Cognition, 45 (2), pp.189–211.
129. Lai, L., Huang, A.Z. and Gershman, S.J. (**2022**). 'Action chunking as policy compression'. Available [Online] at: https://psyarxiv.com/z8yrv/.
130. Lanciego, J.L., Luquin, N. and Obeso, J.A. (**2012**). 'Functional Neuroanatomy of the Basal Ganglia'. Cold Spring Harbor Perspectives in Medicine, 2 (12), p.a009621.
131. Lashley, K.S. (**1951**). 'The Problem of Serial Order in Behavior'. , p.36.
132. Latchoumane, C.-F.V., Ngo, H.-V.V., Born, J. and Shin, H.-S. (**2017**). 'Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms'. Neuron, 95 (2), pp.424-435.e6.
133. Lauwereyns, J., Watanabe, K., Coe, B. and Hikosaka, O. (**2002**). 'A neural correlate of response bias in monkey caudate nucleus'. Nature, 418 (6896), pp.413–417.
134. Laventure, S., Fogel, S., Lungu, O., Albouy, G., Sévigny-Dupont, P., Vien, C., Sayour, C., Carrier, J., Benali, H. and Doyon, J. (**2016**). 'NREM2 and Sleep Spindles Are Instrumental to the Consolidation of Motor Sequence Memories'. PLOS Biology, 14 (3), p.e1002429.
135. Lee, A.K. and Wilson, M.A. (**2002**). 'Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep'. Neuron, 36 (6), pp.1183–1194.
136. Lemke, S.M., Ramanathan, D.S., Darevksy, D., Egert, D., Berke, J.D. and Ganguly, K. (**2021**). 'Coupling between motor cortex and striatum increases during sleep over long-term skill learning' M. J. Frank, A. H. Gittis, E. Yttri, and D. Robbe (Eds.),. eLife, 10, p.e64303.
137. Li, W., Ma, L., Yang, G. and Gan, W.-B. (**2017**). 'REM sleep selectively prunes and maintains new synapses in development and learning'. *Nature Neuroscience*, 20 (3), pp.427–437.
138. Lim, D., LeDue, J., Mohajerani, M., Vanni, M. and Murphy, T. (**2013**). 'Optogenetic approaches for functional mouse brain mapping'. *Frontiers in Neuroscience*, 7.
139. Liu, Y., Dolan, R.J., Kurth-Nelson, Z. and Behrens, T.E.J. (**2019**). 'Human Replay Spontaneously Reorganizes Experience'. *Cell*, 178 (3), pp.640-652.e14.
140. Logiaco, L., Abbott, L.F. and Escola, S. (**2021**). 'Thalamic control of cortical dynamics in a model of flexible motor sequencing'. *Cell Reports*, 35 (9), p.109090.
141. Long, J., Feng, Y., Liao, H., Zhou, Q. and Urbin, M.A. (**2018**). 'Motor Sequence Learning Is Associated With Hippocampal Subfield Volume in Humans With Medial Temporal Lobe Epilepsy'. *Frontiers in Human Neuroscience*, 12, p.367.
142. Lustenberger, C., Boyle, M.R., Alagapan, S., Mellin, J.M., Vaughn, B.V. and Fröhlich, F. (**2016**). 'Feedback-Controlled Transcranial Alternating Current Stimulation Reveals a Functional Role of Sleep Spindles in Motor Memory Consolidation'. *Current Biology*, 26 (16), pp.2127–2136.
143. Ly, J.Q.M., Gaggioni, G., Chellappa, S.L., Papachilleos, S., Brzozowski, A., Borsu, C., Rosanova, M., Sarasso, S., Middleton, B., Luxen, A., Archer, S.N., Phillips, C., Dijk, D.-J., Maquet, P., Massimini, M. and Vandewalle, G. (**2016**). 'Circadian regulation of human cortical excitability'. *Nature Communications*, 7 (1), p.11828.
144. Maboudi, K., Ackermann, E., de Jong, L.W., Pfeiffer, B.E., Foster, D., Diba, K. and Kemere, C. (**2018**). 'Uncovering temporal structure in hippocampal output patterns' L. Colgin (Ed.),. *eLife*, 7, p.e34467.
145. Maes, A., Barahona, M. and Clopath, C. (**2021**). 'Learning compositional sequences with multiple time scales through a hierarchical network of spiking neurons'. *PLOS Computational Biology*, 17 (3), p.e1008866.

146. Magee, J.C. and Johnston, D. (**1997**). 'A Synaptically Controlled, Associative Signal for Hebbian Plasticity in Hippocampal Neurons'. *Science*, 275 (5297), pp.209–213.
147. Maingret, N., Girardeau, G., Todorova, R., Goutierre, M. and Zugaro, M. (**2016**). 'Hippocampo-cortical coupling mediates memory consolidation during sleep'. *Nature Neuroscience*, 19 (7), pp.959–964.
148. Maltry, L., Ohlendorf, D., Bendels, M., Groneberg, D.A. and Mickel, C. (**2021**). 'Does sleep contribute to the consolidation of motor memory?'. *Manuelle Medizin*, 59 (1), pp.26–32.
149. Mandelbaum, G., Taranda, J., Haynes, T.M., Hochbaum, D.R., Huang, K.W., Hyun, M., Umadevi Venkataraju, K., Straub, C., Wang, W., Robertson, K., Osten, P. and Sabatini, B.L. (**2019**). 'Distinct Cortical-Thalamic-Striatal Circuits through the Parafascicular Nucleus'. *Neuron*, 102 (3), pp.636-652.e7.
150. Manohar, S.G., Finzi, R.D., Drew, D. and Husain, M. (**2017**). 'Distinct Motivational Effects of Contingent and Noncontingent Rewards'. *Psychological Science*, 28 (7), pp.1016–1026.
151. Markowitz, J.E., Gillis, W.F., Beron, C.C., Neufeld, S.Q., Robertson, K., Bhagat, N.D., Peterson, R.E., Peterson, E., Hyun, M., Linderman, S.W., Sabatini, B.L. and Datta, S.R. (**2018**). 'The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection'. *Cell*, 174 (1), pp.44-58.e17.
152. Markowitz, J.E., Gillis, W.F., Jay, M., Wood, J., Harris, R.W., Cieszkowski, R., Scott, R., Brann, D., Koveal, D., Kula, T., Weinreb, C., Osman, M.A.M., Pinto, S.R., Uchida, N., Linderman, S.W., Sabatini, B.L. and Datta, S.R. (**2023**). 'Spontaneous behaviour is structured by reinforcement without explicit reward'. *Nature*, 614 (7946), pp.108–117.
153. Marshall, J.D., Aldarondo, D.E., Dunn, T.W., Wang, W.L., Berman, G.J. and Ölveczky, B.P. (**2021**). 'Continuous Whole-Body 3D Kinematic Recordings across the Rodent Behavioral Repertoire'. *Neuron*, 109 (3), pp.420-437.e8.
154. Márton, C.D., Schultz, S.R. and Averbeck, B.B. (**2020**). 'Learning to select actions shapes recurrent dynamics in the corticostriatal system'. *Neural Networks*, 132, pp.375–393.
155. Masullo, L., Mariotti, L., Alexandre, N., Freire-Pritchett, P., Boulanger, J. and Tripodi, M. (**2019**). 'Genetically Defined Functional Modules for Spatial Orienting in the Mouse Superior Colliculus'. *Current Biology*, 29 (17), pp.2892-2904.e8.
156. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W. and Bethge, M. (**2018**). 'DeepLabCut: markerless pose estimation of user-defined body parts with deep learning'. *Nature Neuroscience*, 21 (9), pp.1281–1289.
157. Matsumoto, N., Minamimoto, T., Graybiel, A.M. and Kimura, M. (**2001**). 'Neurons in the thalamic CM-Pf complex supply striatal neurons with information about behaviorally significant sensory events'. *Journal of Neurophysiology*, 85 (2), pp.960–976.
158. Mayerich, D., Abbott, L. and McCORMICK, B. (**2008**). 'Knife-edge scanning microscopy for imaging and reconstruction of three-dimensional anatomical structures of the mouse brain'. *Journal of Microscopy*, 231 (1), pp.134–143.
159. Mednick, S.C., Cai, D.J., Shuman, T., Anagnostaras, S. and Wixted, J.T. (**2011**). 'An opportunistic theory of cellular and systems consolidation'. *Trends in Neurosciences*, 34 (10), pp.504–514.
160. Mehaffey, W.H. and Doupe, A.J. (**2015**). 'Naturalistic stimulation drives opposing heterosynaptic plasticity at two inputs to songbird cortex'. *Nature neuroscience*, 18 (9), pp.1272–1280.
161. Meij, R. van der and Voytek, B. (**2018**). 'Uncovering Neuronal Networks Defined by Consistent Between-Neuron Spike Timing from Neuronal Spike Recordings'. *eNeuro*, 5 (3).
162. Melief, E.J., McKinley, J.W., Lam, J.Y., Whiteley, N.M., Gibson, A.W., Neumaier, J.F., Henschen, C.W., Palmiter, R.D., Bamford, N.S. and Darvas, M. (**2018**). 'Loss of glutamate signaling from the thalamus to dorsal striatum impairs motor function and slows the execution of learned behaviors'. *npj Parkinson's Disease*, 4 (1), p.23.
163. Minamimoto, T., Hori, Y. and Kimura, M. (**2005**). 'Complementary Process to Response Bias in the Centromedian Nucleus of the Thalamus'. *Science*, 308 (5729), pp.1798–1801.

164. Miyachi, S., Hikosaka, O., Miyashita, K., Kárádi, Z. and Rand, M.K. (**1997**). 'Differential roles of monkey striatum in learning of sequential hand movement'. *Experimental Brain Research*, 115 (1), pp.1–5.
165. Mizes, K.G.C., Lindsey, J., Escola, G.S. and Ölveczky, B.P. (**2022**). 'Similar striatal activity exerts different control over automatic and flexible motor sequences'. p.2022.06.13.495989.
166. Moll, F.W., Kranz, D., Corredera Asensio, A., Elmaleh, M., Ackert-Smith, L.A. and Long, M.A. (**2023**). 'Thalamus drives vocal onsets in the zebra finch courtship song'. *Nature*, 616 (7955), pp.132–136.
167. Morris, R.G.M., Garrud, P., Rawlins, J.N.P. and O'Keefe, J. (**1982**). 'Place navigation impaired in rats with hippocampal lesions'. *Nature*, 297 (5868), pp.681–683.
168. Murray, J.M. and Escola, G.S. (**2020**). 'Remembrance of things practiced with fast and slow learning in cortical and subcortical pathways'. *Nature Communications*, 11 (1), p.6441.
169. Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J. and Buzsáki, G. (**1999**). 'Replay and time compression of recurring spike sequences in the hippocampus'. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 19 (21), pp.9497–9507.
170. Nagai, H., de Vivo, L., Bellesi, M., Ghilardi, M.F., Tononi, G. and Cirelli, C. (**2017**). 'Sleep Consolidates Motor Learning of Complex Movement Sequences in Mice'. *Sleep*, 40 (2), p.zsw059.
171. Nagai, H., de Vivo, L., Bellesi, M., Ghilardi, M.F., Tononi, G. and Cirelli, C. (**2016**). 'Sleep Consolidates Motor Learning of Complex Movement Sequences in Mice'. *Sleep*, 40 (2), p.zsw059.
172. Nanda, B., Galvan, A., Smith, Y. and Wichmann, T. (**2009**). 'Effects of stimulation of the centromedian nucleus of the thalamus on the activity of striatal cells in awake rhesus monkeys'. *European Journal of Neuroscience*, 29 (3), pp.588–598.
173. Nishida, M. and Walker, M.P. (**2007**). 'Daytime Naps, Motor Memory Consolidation and Regionally Specific Sleep Spindles'. *PLoS ONE*, 2 (4), p.e341.
174. Ólafsdóttir, H.F., Bush, D. and Barry, C. (**2018**). 'The Role of Hippocampal Replay in Memory and Planning'. *Current Biology*, 28 (1), pp.R37–R50.
175. O'Neill, J., Pleydell-Bouverie, B., Dupret, D. and Csicsvari, J. (**2010**). 'Play it again: reactivation of waking experience and memory'. *Trends in Neurosciences*, 33 (5), pp.220–229.
176. Parent, M. and Parent, A. (**2006**). 'Single-axon tracing study of corticostriatal projections arising from primary motor cortex in primates'. *The Journal of Comparative Neurology*, 496 (2), pp.202–213.
177. Park, J., Coddington, L.T. and Dudman, J.T. (**2020**). 'Basal Ganglia Circuits for Action Specification'. *Annual Review of Neuroscience*, 43 (1), pp.485–507.
178. Pavlides, C., Miyashita, E. and Asanuma, H. (**1993**). 'Projection from the sensory to the motor cortex is important in learning motor skills in the monkey'. *Journal of Neurophysiology*, 70 (2), pp.733–741.
179. Perrin, E. and Venance, L. (**2019**). 'Bridging the gap between striatal plasticity and learning'. *Current Opinion in Neurobiology*, 54, pp.104–112.
180. Peters, A.J., Chen, S.X. and Komiyama, T. (**2014**). 'Emergence of reproducible spatiotemporal activity during motor learning'. *Nature*, 510 (7504), pp.263–267.
181. Ragan, T., Kadiri, L.R., Venkataraju, K.U., Bahlmann, K., Sutin, J., Taranda, J., Arganda-Carreras, I., Kim, Y., Seung, H.S. and Osten, P. (**2012**). 'Serial two-photon tomography for automated ex vivo mouse brain imaging'. *Nature Methods*, 9 (3), pp.255–258.
182. Ramanathan, D.S., Gulati, T. and Ganguly, K. (**2015**). 'Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation'. *PLOS Biology*, 13 (9), p.e1002263.
183. Rand, M.K., Hikosaka, O., Miyachi, S., Lu, X., Nakamura, K., Kitaguchi, K. and Shimo, Y. (**2000**). 'Characteristics of sequential movements during early learning period in monkeys'. *Experimental Brain Research*, 131 (3), pp.293–304.
184. Rasch, B. and Born, J. (**2013**). 'About Sleep's Role in Memory'. *Physiological Reviews*, 93 (2), pp.681–766.
185. Rasch, B., Pommer, J., Diekelmann, S. and Born, J. (**2009**). 'Pharmacological REM sleep suppression paradoxically improves rather than impairs skill memory'. *Nature Neuroscience*, 12 (4), pp.396–397.

186. Reber, P.J. and Squire, L.R. (**1998**). 'Encapsulation of Implicit and Explicit Memory in Sequence Learning'. *Journal of Cognitive Neuroscience*, 10 (2), pp.248–263.
187. Reynolds, J.N., Hyland, B.I. and Wickens, J.R. (**2001**). 'A cellular mechanism of reward-related learning'. *Nature*, 413 (6851), pp.67–70.
188. Ribeiro, S., Gervasoni, D., Soares, E.S., Zhou, Y., Lin, S.-C., Pantoja, J., Lavine, M. and Nicolelis, M.A.L. (**2004**). 'Long-Lasting Novelty-Induced Neuronal Reverberation during Slow-Wave Sleep in Multiple Forebrain Areas'. *PLoS Biology*, 2 (1), p.e24.
189. Rigoux, L. and Guigon, E. (**2012**). 'A Model of Reward- and Effort-Based Optimal Decision Making and Motor Control'. *PLOS Computational Biology*, 8 (10), p.e1002716.
190. Rioult-Pedotti, M.S., Friedman, D., Hess, G. and Donoghue, J.P. (**1998**). 'Strengthening of horizontal cortical connections following skill learning'. *Nature Neuroscience*, 1 (3), pp.230–234.
191. Robertson, E.M., Pascual-Leone, A. and Miall, R.C. (**2004**). 'Current concepts in procedural consolidation'. *Nature Reviews Neuroscience*, 5 (7), pp.576–582.
192. Rosenbaum, D.A., Cohen, R.G., Jax, S.A., Weiss, D.J. and van der Wel, R. (**2007**). 'The problem of serial order in behavior: Lashley's legacy'. *Human Movement Science*, 26 (4), pp.525–554.
193. Rosenbaum, D.A., Kenny, S.B. and Derr, M.A. (**1983**). 'Hierarchical control of rapid movement sequences'. *Journal of Experimental Psychology: Human Perception and Performance*, 9, pp.86–102.
194. Roumis, D.K. and Frank, L.M. (**2015**). 'Hippocampal sharp-wave ripples in waking and sleeping states'. *Current Opinion in Neurobiology*, 35, pp.6–12.
195. Roux, L., Hu, B., Eichler, R., Stark, E. and Buzsáki, G. (**2017**). 'Sharp wave ripples during learning stabilize the hippocampal spatial map'. *Nature Neuroscience*, 20 (6), pp.845–853.
196. Royce, G.J. (**1983**). 'Cells of origin of corticothalamic projections upon the centromedian and parafascicular nuclei in the cat'. *Brain Research*, 258 (1), pp.11–21.
197. Rubin, D.B., Hosman, T., Kelemen, J.N., Kapitonava, A., Willett, F.R., Coughlin, B.F., Halgren, E., Kimchi, E.Y., Williams, Z.M., Simeral, J.D., Hochberg, L.R. and Cash, S.S. (**2022**). 'Learned Motor Patterns Are Replayed in Human Motor Cortex during Sleep'. *The Journal of Neuroscience*, 42 (25), pp.5007–5020.
198. Rueda-Orozco, P.E. and Robbe, D. (**2015**). 'The striatum multiplexes contextual and kinematic information to constrain motor habits execution'. *Nature neuroscience*, 18 (3), pp.453–460.
199. Saalmann, Y.B. (**2014**). 'Intralaminar and medial thalamic influence on cortical synchrony, information transmission and cognition'. *Frontiers in Systems Neuroscience*, 8, p.83.
200. Sadikot, A.F., Parent, A. and François, C. (**1992**). 'Efferent connections of the centromedian and parafascicular thalamic nuclei in the squirrel monkey: A PHA-L study of subcortical projections'. *Journal of Comparative Neurology*, 315 (2), pp.137–159.
201. Sakai, K., Kitaguchi, K. and Hikosaka, O. (**2003**). 'Chunking during human visuomotor sequence learning'. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 152, pp.229–42.
202. Samborska, V., Butler, J.L., Walton, M.E., Behrens, T.E.J. and Akam, T. (**2022**). 'Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems'. *Nature Neuroscience*, 25 (10), pp.1314–1326.
203. Sánchez-Fuentes, A., Ramírez-Armenta, K.I., Verma-Rodríguez, A.K., Díaz-Hernández, E., Aguilar-Palomares, A., Ramírez-Jarquín, J.O. and Tecuapetla, F. (**2021**). 'The Contribution of Premotor Cortico-Striatal Projections to the Execution of Serial Order Sequences'. *eNeuro*, 8 (5), p.ENEURO.0173-21.2021.
204. Sánchez-Valpuesta, M., Suzuki, Y., Shibata, Y., Toji, N., Ji, Y., Afrin, N., Asogwa, C.N., Kojima, I., Mizuguchi, D., Kojima, S., Okanoya, K., Okado, H., Kobayashi, K. and Wada, K. (**2019**). 'Corticobasal ganglia projecting neurons are required for juvenile vocal learning but not for adult vocal plasticity in songbirds'. *Proceedings of the National Academy of Sciences of the United States of America*, 116 (45), pp.22833–22843.
205. Santos, F.J., Oliveira, R.F., Jin, X. and Costa, R.M. (**2015**). 'Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning'. *eLife*, 4, p.e09423.

206. Schapiro, A.C., Reid, A.G., Morgan, A., Manoach, D.S., Verfaellie, M. and Stickgold, R. (**2019**). 'The hippocampus is necessary for the consolidation of a task that does not require the hippocampus for initial learning'. *Hippocampus*, 29 (11), pp.1091–1100.
207. Scharff, C., Kirn, J.R., Grossman, M., Macklis, J.D. and Nottebohm, F. (**2000**). 'Targeted Neuronal Death Affects Neuronal Replacement and Vocal Behavior in Adult Songbirds'. *Neuron*, 25 (2), pp.481–492.
208. Schendan, H.E., Searl, M.M., Melrose, R.J. and Stern, C.E. (**2003**). 'An fMRI Study of the Role of the Medial Temporal Lobe in Implicit and Explicit Sequence Learning'. *Neuron*, 37 (6), pp.1013–1025.
209. Schmid, D., Erlacher, D., Klostermann, A., Kredel, R. and Hossner, E.-J. (**2020**). 'Sleep-dependent motor memory consolidation in healthy adults: A meta-analysis'. *Neuroscience & Biobehavioral Reviews*, 118, pp.270–281.
210. Schmidt, M.F. (**2003**). 'Pattern of Interhemispheric Synchronization in HVc During Singing Correlates With Key Transitions in the Song Pattern'. *Journal of Neurophysiology*, 90 (6), pp.3931–3949.
211. Schönauer, M., Geisler, T. and Gais, S. (**2014**). 'Strengthening procedural memories by reactivation in sleep'. *Journal of Cognitive Neuroscience*, 26 (1), pp.143–153.
212. Schreiner, T., Petzka, M., Staudigl, T. and Staresina, B.P. (**2021**). 'Endogenous memory reactivation during sleep in humans is clocked by slow oscillation-spindle complexes'. *Nature Communications*, 12 (1), p.3112.
213. Schwarting, R.K.W. and Busse, S. (**2017**). 'Behavioral facilitation after hippocampal lesion: A review'. *Behavioural Brain Research*, 317, pp.401–414.
214. Scoville, W.B. and Milner, B. (**1957**). 'Loss of Recent Memory After Bilateral Hippocampal Lesions'. *Journal of Neurology, Neurosurgery & Psychiatry*, 20 (1), pp.11–21.
215. Shadmehr, R., Huang, H.J. and Ahmed, A.A. (**2016**). 'A Representation of Effort in Decision-Making and Motor Control'. *Current Biology*, 26 (14), pp.1929–1934.
216. Shima, K. and Tanji, J. (**2000**). 'Neuronal Activity in the Supplementary and Presupplementary Motor Areas for Temporal Organization of Multiple Movements'. *Journal of Neurophysiology*, 84 (4), pp.2148–2160.
217. Shin, J.D., Tang, W. and Jadhav, S.P. (**2019**). 'Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making'. *Neuron*, 104 (6), pp.1110-1125.e7.
218. Siapas, A.G. and Wilson, M.A. (**1998**). 'Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep'. *Neuron*, 21 (5), pp.1123–1128.
219. Singer, A.C. and Frank, L.M. (**2009**). 'Rewarded outcomes enhance reactivation of experience in the hippocampus'. *Neuron*, 64 (6), pp.910–921.
220. Skaggs, W.E. and McNaughton, B.L. (**1996**). 'Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience'. *Science (New York, N.Y.)*, 271 (5257), pp.1870–1873.
221. Smeal, R.M., Keefe, K.A. and Wilcox, K.S. (**2008**). 'Differences in excitatory transmission between thalamic and cortical afferents to single spiny efferent neurons of rat dorsal striatum'. *European Journal of Neuroscience*, 28 (10), pp.2041–2052.
222. Smith, null and MacNeill, null (**1994**). 'Impaired motor memory for a pursuit rotor task following Stage 2 sleep loss in college students'. *Journal of Sleep Research*, 3 (4), pp.206–213.
223. Smith, C. (**2001**). 'Sleep states and memory processes in humans: procedural versus declarative memory systems'. *Sleep Medicine Reviews*, 5 (6), pp.491–506.
224. Smith, C., Aubrey, J. and Peters, K. (**2004**). 'Different Roles for REM and Stage 2 Sleep In Motor Learning: A Proposed Model'. *Psychologica Belgica. Special Issue: Cognition in Slumberland. Mechanisms of Information Processing in the Sleep-Wake Cycle*, 44, pp.81–104.
225. Smith, Y., Galvan, A., Ellender, T.J., Doig, N., Villalba, R.M., Huerta-Ocampo, I., Wichmann, T. and Bolam, J.P. (**2014**). 'The thalamostriatal system in normal and diseased states'. *Frontiers in Systems Neuroscience*, 8, p.5.
226. Sommer, M.A. and Wurtz, R.H. (**2002**). 'A pathway in primate brain for internal monitoring of movements'. *Science (New York, N.Y.)*, 296 (5572), pp.1480–1482.

227. Sommer, M.A. and Wurtz, R.H. (**2004**). 'What the Brain Stem Tells the Frontal Cortex. II. Role of the SC-MD-FEF Pathway in Corollary Discharge'. *Journal of Neurophysiology*, 91 (3), pp.1403–1423.
228. Spens, E. and Burgess, N. (**2023**). 'A Generative Model of Memory Construction and Consolidation'. , p.2023.01.19.524711.
229. Squire, L.R. (**2009**). 'The Legacy of Patient H.M. for Neuroscience'. *Neuron*, 61 (1), pp.6–9.
230. Staresina, B.P., Niediek, J., Borger, V., Surges, R. and Mormann, F. (**2023**). 'How coupled slow oscillations, spindles and ripples control neuronal processing and communication during human sleep'. , p.2023.01.08.523138.
231. Stella, F., Baracskay, P., O'Neill, J. and Csicsvari, J. (**2019**). 'Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion'. *Neuron*, 102 (2), pp.450-461.e7.
232. Stephenson-Jones, M., Samuelsson, E., Ericsson, J., Robertson, B. and Grillner, S. (**2011**). 'Evolutionary Conservation of the Basal Ganglia as a Common Vertebrate Mechanism for Action Selection'. *Current Biology*, 21 (13), pp.1081–1091.
233. Summerside, E.M., Shadmehr, R. and Ahmed, A.A. (**2018**). 'Vigor of reaching movements: reward discounts the cost of effort'. *Journal of Neurophysiology*, 119 (6), pp.2347–2357.
234. Takakusaki, K., Chiba, R., Nozu, T. and Okumura, T. (**2016**). 'Brainstem control of locomotion and muscle tone with special reference to the role of the mesopontine tegmentum and medullary reticulospinal systems'. *Journal of Neural Transmission*, 123 (7), pp.695–729.
235. Takikawa, Y., Kawagoe, R., Itoh, H., Nakahara, H. and Hikosaka, O. (**2002**). 'Modulation of saccadic eye movements by predicted reward outcome'. *Experimental Brain Research*, 142 (2), pp.284–291.
236. Tanaka, M. (**2006**). 'Inactivation of the central thalamus delays self-timed saccades'. *Nature Neuroscience*, 9 (1), pp.20–22.
237. Tang, W., Shin, J.D., Frank, L.M. and Jadhav, S.P. (**2017**). 'Hippocampal-Prefrontal Reactivation during Learning Is Stronger in Awake Compared with Sleep States'. *Journal of Neuroscience*, 37 (49), pp.11789–11805.
238. Tervo, D.G.R., Tenenbaum, J.B. and Gershman, S.J. (**2016**). 'Toward the neural implementation of structure learning'. *Current Opinion in Neurobiology*, 37, pp.99–105.
239. Tingley, D. and Buzsáki, G. (**2020**). 'Routing of Hippocampal Ripples to Subcortical Structures via the Lateral Septum'. *Neuron*, 105 (1), pp.138-149.e5.
240. Tingley, D. and Peyrache, A. (**2020**). 'On the methods for reactivation and replay analysis'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375 (1799), p.20190231.
241. Tolman, E.C., Ritchie, B.F. and Kalish, D. (**1946**). 'Studies in spatial learning. II. Place learning versus response learning'. *Journal of Experimental Psychology*, 36 (3), pp.221–229.
242. Tresch, M.C., Saltiel, P. and Bizzi, E. (**1999**). 'The construction of movement by the spinal cord'. *Nature Neuroscience*, 2 (2), pp.162–167.
243. Tsodyks, M. (**2005**). 'Attractor Neural Networks and Spatial Maps in Hippocampus'. *Neuron*, 48 (2), pp.168–169.
244. Tyson, A.L. and Margrie, T.W. (**2022**). 'Mesoscale microscopy and image analysis tools for understanding the brain'. *Progress in Biophysics and Molecular Biology*, 168, pp.81–93.
245. Tyson, A.L., Vélez-Fort, M., Rousseau, C.V., Cossell, L., Tsitoura, C., Lenzi, S.C., Obenhaus, H.A., Claudi, F., Branco, T. and Margrie, T.W. (**2022**). 'Accurate determination of marker location within whole-brain microscopy images'. *Scientific Reports*, 12 (1), p.867.
246. Vaidya, A.R., Jones, H.M., Castillo, J. and Badre, D. (**2021**). 'Neural representation of abstract task structure during generalization' M. Liljeholm, R. B. Ivry, C. Ranganath, and S. Michelmann (Eds.),. *eLife*, 10, p.e63226.
247. Vandermaelen, C.P. and Kitai, S.T. (**1980**). 'Intracellular analysis of synaptic potentials in rat neostriatum following stimulation of the cerebral cortex, thalamus, and substantia nigra'. *Brain Research Bulletin*, 5 (6), pp.725–733.

248. Velayos, J.L., Ojeda, P. and Picó, J.M. (**1989**). '[Organization of the projections of the cortex of the sulcus cruciatus to the centromedian and parafascicular nuclei in cats]'. *Revista Espanola De Fisiologia*, 45 Suppl, pp.151–160.
249. Vicente, A.M., Martins, G.J. and Costa, R.M. (**2020**). 'Cortico-basal ganglia circuits underlying dysfunctional control of motor behaviors in neuropsychiatric disorders'. *Current Opinion in Genetics & Development*, 65, pp.151–159.
250. Walker, M.P., Brakefield, T., Hobson, J.A. and Stickgold, R. (**2003**). 'Dissociable stages of human memory consolidation and reconsolidation'. *Nature*, 425 (6958), pp.616–620.
251. Walker, M.P., Brakefield, T., Morgan, A., Hobson, J.A. and Stickgold, R. (**2002**). 'Practice with Sleep Makes Perfect: Sleep-Dependent Motor Skill Learning'. *Neuron*, 35 (1), pp.205–211.
252. Walker, M.P., Brakefield, T., Seidman, J., Morgan, A., Hobson, J.A. and Stickgold, R. (**2003**). 'Sleep and the time course of motor skill learning'. *Learning & Memory (Cold Spring Harbor, N.Y.)*, 10 (4), pp.275–284.
253. Wheatcroft, T., Saleem, A.B. and Solomon, S.G. (**2022**). 'Functional Organisation of the Mouse Superior Colliculus'. *Frontiers in Neural Circuits*, 16.
254. White, N.M. and McDonald, R.J. (**2002**). 'Multiple Parallel Memory Systems in the Brain of the Rat'. *Neurobiology of Learning and Memory*, 77 (2), pp.125–184.
255. Wiestler, T. and Diedrichsen, J. (**2013**). 'Skill learning strengthens cortical representations of motor sequences'. *eLife*, 2, p.e00801.
256. Will, J.L., Eckart, M.T., Rosenow, F., Bauer, S., Oertel, W.H., Schwarting, R.K.W. and Norwood, B.A. (**2013**). 'Enhanced sequential reaction time task performance in a rat model of mesial temporal lobe epilepsy with classic hippocampal sclerosis'. *Behavioural Brain Research*, 247, pp.65–72.
257. Williams, A.H., Degleris, A., Wang, Y. and Linderman, S.W. (**2020**). 'Point process models for sequence detection in high-dimensional neural spike trains'. *Advances in neural information processing systems*, 33, pp.14350–14361.
258. Williams, H. and Vicario, D.S. (**1993**). 'Temporal patterning of song production: Participation of nucleus uvaeformis of the thalamus'. *Journal of Neurobiology*, 24 (7), pp.903–912.
259. Wilson, M.A. and McNaughton, B.L. (**1994**). 'Reactivation of Hippocampal Ensemble Memories During Sleep'. *Science*, 265 (5172), pp.676–679.
260. Wilson, R.C., Takahashi, Y.K., Schoenbaum, G. and Niv, Y. (**2014**). 'Orbitofrontal Cortex as a Cognitive Map of Task Space'. *Neuron*, 81 (2), pp.267–279.
261. Wiltschko, A.B., Johnson, M.J., Iurilli, G., Peterson, R.E., Katon, J.M., Pashkovski, S.L., Abraira, V.E., Adams, R.P. and Datta, S.R. (**2015**). 'Mapping Sub-Second Structure in Mouse Behavior'. *Neuron*, 88 (6), pp.1121–1135.
262. Wolff, S.B.E., Ko, R. and Ölveczky, B.P. (**2022**). 'Distinct roles for motor cortical and thalamic inputs to striatum during motor skill learning and execution'. *Science Advances*, 8 (8), p.eabk0231.
263. Xiao, L. and Roberts, T.F. (**2021**). 'What Is the Role of Thalamostriatal Circuits in Learning Vocal Sequences?'. *Frontiers in Neural Circuits*, 15, p.724858.
264. Xu, T., Yu, X., Perlik, A.J., Tobin, W.F., Zweig, J.A., Tennant, K., Jones, T. and Zuo, Y. (**2009**). 'Rapid formation and selective stabilization of synapses for enduring motor memories'. *Nature*, 462 (7275), pp.915–919.
265. Yamasaki, D.S.G., Krauthamer, G.M. and Rhoades, R.W. (**1986**). 'Superior collicular projection to intralaminar thalamus in rat'. *Brain Research*, 378 (2), pp.223–233.
266. Yang, G., Lai, C.S.W., Cichon, J., Ma, L., Li, W. and Gan, W.-B. (**2014**). 'Sleep promotes branch-specific formation of dendritic spines after learning'. *Science*, 344 (6188), pp.1173–1178.
267. Yin, H.H. (**2010**). 'The Sensorimotor Striatum Is Necessary for Serial Order Learning'. *The Journal of Neuroscience*, 30 (44), pp.14719–14723.
268. Yin, H.H. and Knowlton, B.J. (**2006**). 'The role of the basal ganglia in habit formation'. *Nature Reviews. Neuroscience*, 7 (6), pp.464–476.

269. Yokoi, A., Arbuckle, S.A. and Diedrichsen, J. (**2018**). 'The Role of Human Primary Motor Cortex in the Production of Skilled Finger Sequences'. *Journal of Neuroscience*, 38 (6), pp.1430–1442.
270. Zhu, H. and Roth, B.L. (**2014**). 'Silencing synapses with DREADDs'. *Neuron*, 82 (4), pp.723–725.
271. Zimnik, A.J. and Churchland, M.M. (**2021**). 'Independent generation of sequence elements by motor cortex'. *Nature neuroscience*, 24 (3), pp.412–424.

APPENDIX

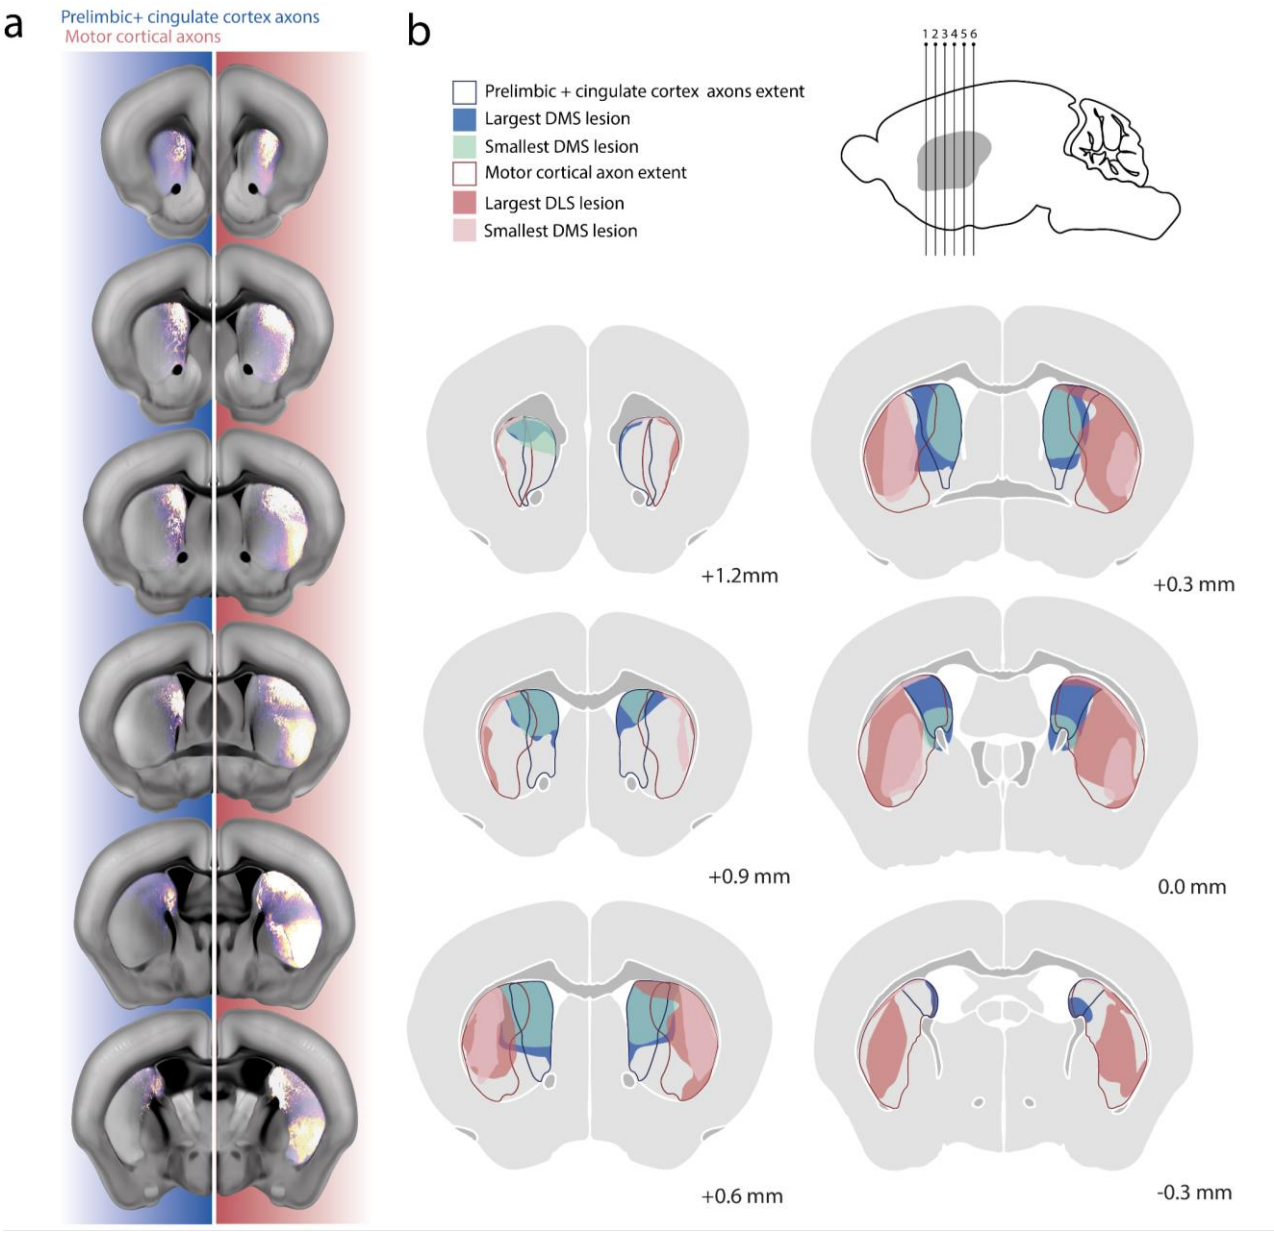

Appendix 1: DLS and DMS lesion extent

(**a**) Allen reference projections from prelimbic and cingulate cortex (left, Allen experiment: 157711748 & 112514202) and motor cortex (right, Allen experiment: 180720175 & 180709942). **(b)** Allen projection defined DMS (blue line) with largest DMS lesion (blue shaded) and smallest DMS lesion (turquoise shaded) for the extent of striatum. Allen projection defined DLS (red line) with largest DLS lesion (red shaded) and smallest DLS lesion (pink shaded) for the extent of striatum.

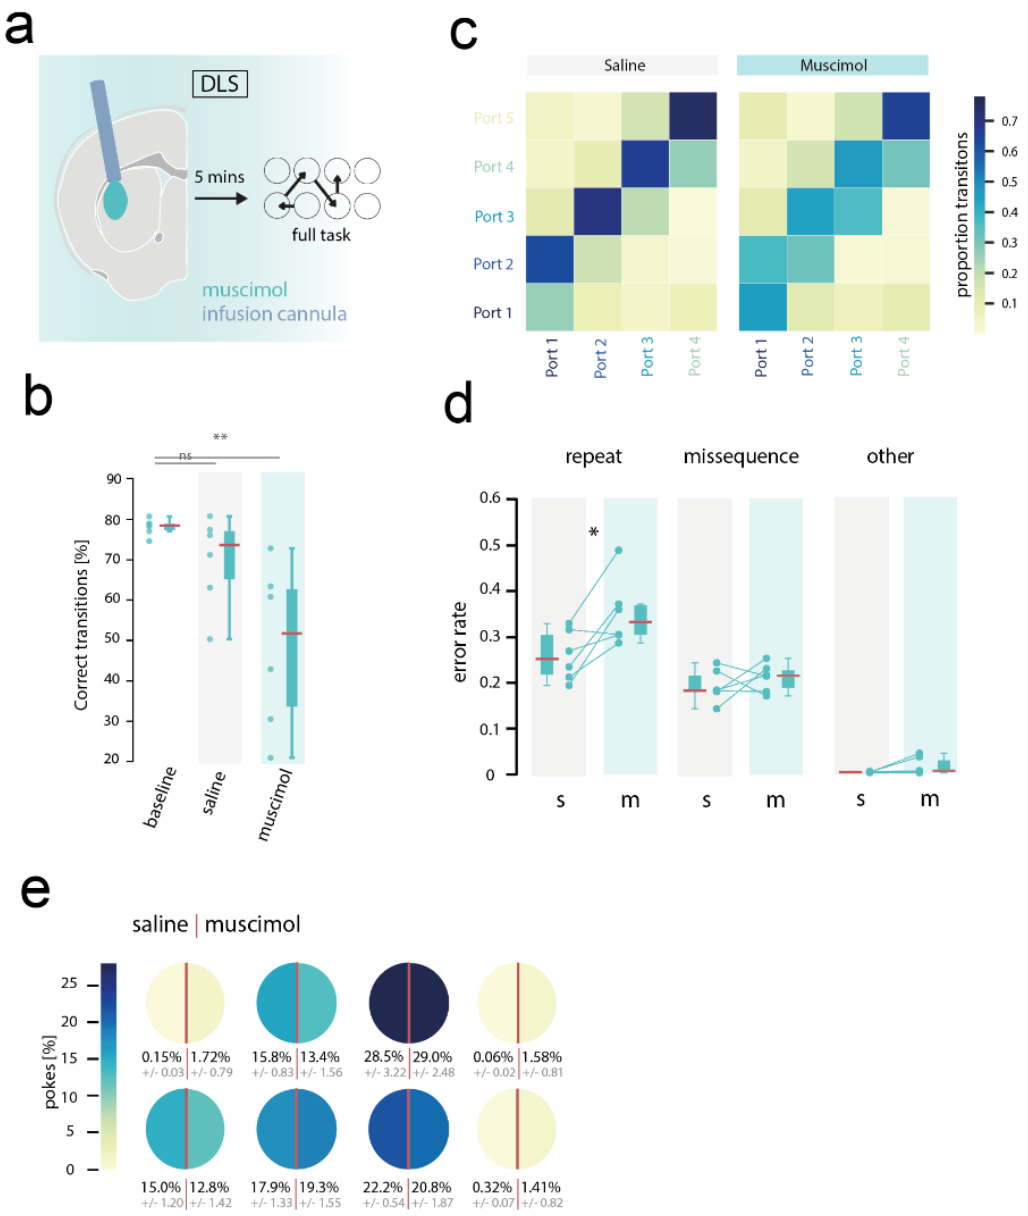

Appendix 2: Muscimol silencing of DLS impairs task execution

(a) Schematic showing experimental design. Trained mice were infused with muscimol via bilaterally implanted cannulas and tested on the task. **(b)** Percentage correct port-to-port transitions for baseline sessions, saline infusions and muscimol infusions. **(c)** Transition histograms for saline and muscimol infusion sessions (**d**) Error type occurrences as a proportion of all transitions (**e**) Port poke proportions after saline (left semicircle) and muscimol (right semicircle)

a

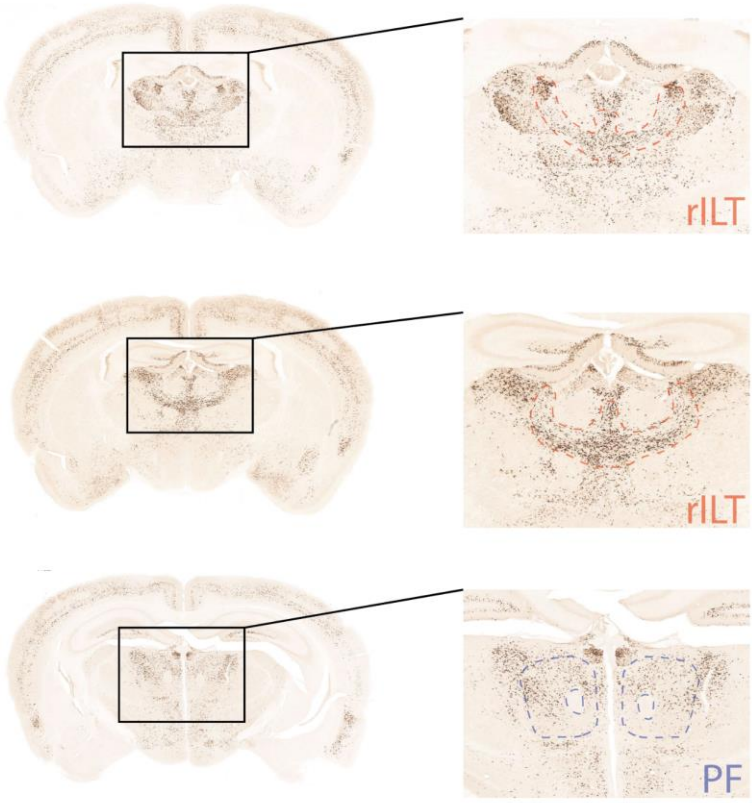

http://www.gensat.org/imagenavigator.jsp?imageID=103561

Appendix 3: GRM2 positive cells and preliminary rILT input tracing

(a) GRM2 expression shown for 3 example slices of intralamainar thalamus. Rostal ILT labelled in red, caudal ILT (parafasicular nuclei) labelled in blue. Data taken from Genstat: http://genstat.ord/imgenavigator.jsp?imageID=103561

b

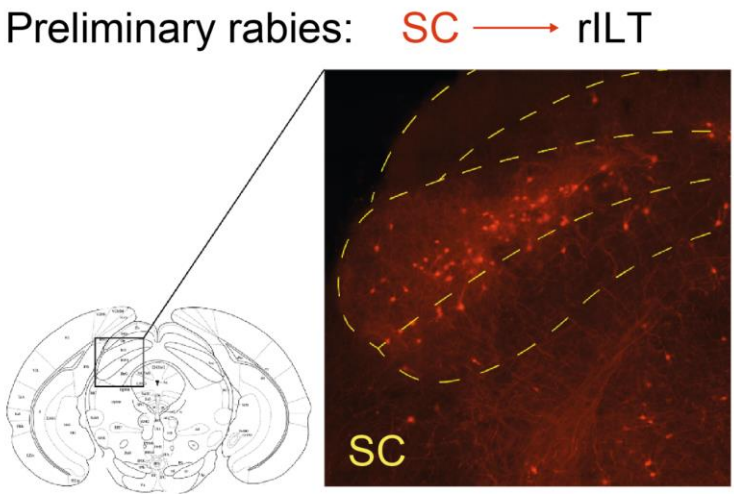

(b) Example fluorescence image showing rabies traced cells in superior colliculus that project to striatum projecting cells in rILT.