UCSF
UC San Francisco Electronic Theses and Dissertations

Title
Confidence in Memories: Behavioral and Neural Approaches

Permalink
https://escholarship.org/uc/item/8116x843

Author
Joo, Hannah Reade

Publication Date
2020

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Confidence in Memories: Behavioral and Neural Approaches

by
Hannah Joo

DISSERTATION
Submitted in partial satisfaction of the requirements for degree of
DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

Approved:

DocuSigned by:
Vikaas Sohal
70A483C8DEB04E6...

Vikaas Sohal
Chair

DocuSigned by:
AB Nelson
[illegible]EB...

Alexandra Nelson

DocuSigned by:
Michael Brainard
[illegible]412...

Michael Brainard

DocuSigned by:
Loren Frank
EF21CD5A6EE643A...

Loren Frank

Committee Members

© Copyright 2020

by

Hannah Joo

For my father, Tae Hong Joo

Acknowledgments

This research was conducted with government funding under and awarded by the National Institutes of Mental Health Grant F30MH115582 to HRJ, National Institute of General Medical Sciences Medical Scientist Training Program grant T32GM007618 to the UCSF Medical Scientist Training Program, and support from the Howard Hughes Medical Institute for the Frank Lab.

This work was made possible by my advisor, Professor Loren Frank. Loren's standard of quantitative description with a convincing narrative grounded in psychology have shaped my thinking about this work and neuroscience in general. I am grateful to Loren for his encouragement, patience, and critical feedback. These five years have been challenging and rewarding because Loren gave me freedom and support to a degree I know I am lucky to have experienced.

One of many specific pieces of advice I've received from Loren is to avoid trying to be exhaustive when giving credit, because doing so properly is impossible. At the same time, however, Loren has never advised against trying an impossible thing.

There are a few whose names will not be here to whom I know I am indebted.

I am grateful to Adam Kepecs, who co-advised the memory confidence project's experimental design and analysis. Adam gave key feedback and encouragement, provided

the starting inspiration for the project, and gave analytical advice throughout.

I am grateful to the MSTP leadership, particularly our administrator Geri Ehle, Director Mark Anderson, and Associate Director Cathy Lomen-Hoerth. Thank you to the Neuroscience Graduate Program leadership, including the former Directors Roger Nicoll and Anatol Kreitzer and administrators Pat Veitch and Lucita Nacionales. My thesis committee was Vikaas Sohal, Michael Brainard, and Alexandra Nelson, and their technical and strategic advice has been essential in the progress of this work, as was advice from Karunesh Ganguly, who sat on my qualifying examination committee. I am grateful to Dirk Kleinhesselink for technical support particularly in management of the large neural datasets. I learned histology and related techniques from Viktor Kharazia, who contributed the histology for this project. Maxim Borius at SpikeGadgets was particularly helpful in designing and re-designing custom hardware, and often re-soldered and delivered these parts by bike on short notice, just in time for a scheduled implant surgery. Supin Chen, Allison Yorita, Razi Haque, and the rest of the LLNL team, especially Jeanie Pebbles, were great collaborators with whom it was fun to work. Anya Kiseleva, the Frank Lab manager, and Ellie Karlsson, our veterinarian, kept the rats healthy over the course of the experiments. A debt of gratitude is owed to the rats, who were as individual, interesting, affectionate and greedy as people. They were: Clooney, Yen, Pitt, Damon, Malala, Hillary, Tubman, Boudica, Beth, Beth's Friend, Empress Wu, Chien Shiung, Sojo, Dolo, Rosa, all the other HJ's, KF's, CGB's, and HL's.

I am also grateful to the former Frank lab members who created the Matlab pipelines and lab culture on which the present is built. In particular, Kenny Kay taught me the difference between good form and sloppy form in lifting and writing. Gideon Rothschild advised on the task design, project timeline, and other critical issues. Tom Davidson taught me how to set up my first Linux machine and with his curiosity and openness contributed to a healthy and productive environment. Jai Yu similarly gave advice on experimental design.

Mari Sosa introduced me to Frank lab culture, experimental design, analysis, and electrophysiology techniques, and was an unfailing example of careful and thorough science. Dan Liu taught me about hardware (we built my first RAID array together) and software (my realtime code was mysteriously glitchy until Dan suggested I add a line to flush my buffers), was always willing to debate what constitutes a plausible source of electrical noise, and was a supportive friend. Jason Chung developed a rat memory task, itself based on work from Demetris Roumis, that was the starting point for the memory confidence task I designed. Jason trained me in the surgeries for implantation of hybrid tetrode-polymer probe recording devices and has been a friend and source of uniquely reliable advice.

During my time in the Frank lab I was lucky to work with three talented technicians, starting with Jiang Lan (Kevin) Fan. Kevin was the person with whom I worked through many of the original ideas for the behavioral task, and we did many of the experiments for the sharpened stiffeners for polymer probe delivery together. I next worked with Charlotte Geaghan-Breiner, whose unusually keen observations of human and rat behavior were foundational to the success of the project and whose organizational, surgical, soldering, and design skills kept us closer to the original experimental timeline than we would have been otherwise. Charlotte's ongoing friendship since leaving the lab has been just as essential to the completion of this work as was her time in lab.

Hexin Liang worked on the project full-time for two years, contributing her sharp intelligence, her curiosity, and her clear, ego-free scientific outlook. Once, following a lab-wide discussion of whether a speaker had demonstrated enough passion in his talk, Hexin said, *I think that true passion can only be shown, through the work.* I agree, and I am lucky to have spent many long days, including a late Valentine's day evening and a very fun SfN sleepover, together with someone so passionate about the brain and behavior. Critically, Hexin acted as the main surgeon while I was recovering from a wrist injury and kept our data collection going. She pushed for a modeling approach to our task and was

instrumental in its genesis.

I have also learned from every other member I overlapped with in the Frank Lab: Michael Coulter, Eric Denovellis, Jen Guidera, David Kastner, Dani Astudillo Maya, Eric Miller, Jojo Yang, and those whose talent and drive I know from our work together, Alison Comrie and Clay Smyth. I learned in particular from Anna Gillespie and Demetris Roumis from their contributions to the lab code base and spike sorting software. In the Kepecs lab, I am grateful to Paul Masset for discussions of confidence reporting methods to use in the task, and to Torben Ott for analysis advice and discussions.

For their advice on science and career, I am grateful to Felice Dunn, Anna Molofsky, and Aimee Kao. When I wanted to stop, Kira Poskanzer encouraged me to keep going. During the long middle of my graduate training, Arun Prakash mentored my elective days in anesthesia, where control and rapid results were a relief from my scientific work. Alexandra Nelson, in addition to serving on my committee, is a role model who provided encouragement and advice when they were most critical. My MD/PhD advisor Dan Lowenstein has generously advised me on a range of topics since 2012. I am grateful to my MSTP big sibs Max Liu and Tess Veuthey.

I am grateful to the four cohorts that have helped me keep pace over these five years. First, the 2013 UCSF MSTP matriculating class, especially Tina Zheng, Jen Liu, Ruiji Jiang, and Ravi Desai. Even after we were no longer roommates, Ravi was my reliable partner for last-minute adventures to Crissy Field, Ocean Beach, and Ghirardelli Square. He listened to hours of venting and amateur psychoanalysis, always had interesting questions to discuss regarding the natural selection of various behaviors, and has literally lifted me when I have fallen (*e.g.*, on 16th street under the 280 onramp, where the sidewalk rises unexpectedly). Second, the 2012-13 Churchill scholars, especially: Ariana Peck, who treats everything with analytical care, knows the best recipes, and with whom I feel calm and clear; Judy Savitskaya, who is interested in absolutely everything; and Jacob Andreas, who helped

during the worst of the medical-to-graduate-school transition, when my window was smashed and my wallet stolen, and has been willing to discuss the important issues of consciousness and Anna Karenina. Third, the matriculating UCSF medical school class of 2013, who are now almost finished with their residency training. Finally, my graduate school class, especially Frances Cho. Other cohorts include my officemates in 515, Abhilasha ('burst no crackers') Joshi and relentlessly inquisitive Rhino Nevers, who were beside me during weeks of data munging, and Thiago Gouvea in the Kepecs lab, who started a parallel project at the same time.

I am grateful to the physicians at UCSF Student Health who treated me for various overuses and crashes.

I am also grateful to every teacher and mentor I have had so far who gave me books, freedom, encouragement, and hard edits, especially Gayle Holeton, Maryanne Stewart, Joanne Emmons, and my college professors William Evans, Samer Hattar, and Peter Holland. My view of what an ideal graduate school experience would look like was shaped a decade ago, by watching Jo Crook, Beth Peterson, Dennis Dacey, and Kate Mulligan during my summers at the University of Washington. Beth taught me how to write a paper and use Illustrator. She was my first true scientific collaborator, and I am not yet at peace with the reality that she is not here to read this.

I am thankful to Jennifer Chin, who has been a reliable and fun friend for over twenty years. Kelly Barry re-taught me how to write and has given me a powerful example of how to build a happy life. She has helped me work through the very hardest of problems. I am grateful to Christopher Polk, Sibella Kraus, and the rest of the Josephine compound for creating a quiet garden cottage where it is easy to focus and pleasant to quarantine.

In 2010, an important teacher who gave me my foundation in neuroscience and my start in research asked me if I like being alone. *Good*, he said after I told him that I do, *this is*

prerequisite for a life in science. It is a certain kind of aloneness, isn't it? More than the absence of others in the room, it can be a mental isolation. Even when I am in this state, there are a few people who have so significantly affected me that I feel them with me always*. First among them are my family.

I am grateful to my extended family on both sides for their encouragement and support, especially my aunt 케현 and her family, and my grandmothers. Grammie has always encouraged me to practice writing. 할머니 (Grandma Joo) understands me without language and seems to be the genetic source of a wild determination I have tried to make use of here.

My mother gave me three essential skills for this work. She taught me to read and write, which are still the ways I think and learn. She was my first editor. She also taught me that it is important to take breaks, and showed me how to do this by rearranging the furniture, taking a trip, reading a novel, or baking bread. She and her husband Mark helped me move in the middle of graduate school and when my wrist was braced, bought me a slow cooker and a breadmaker, which together saved me enough time and wrist flexibility to keep experiments going.

My little brother, Nathan, was my first lab partner and the only one who knows the full extent of my animal behavior studies carried out in childhood. Because of Nathan, I have known for almost my whole life what unconditional love for someone feels like. And I have never known what it is like to take on a challenging project without having at least one person who believes on faith you will succeed.

I am grateful to Benjamin Philip Nachman for his love. To be there for me, he has followed me to unexpected places – into the mountains, the night, and tiny closets. For the duration of my data collection, he had a daily trans-bay commute so that I could more easily return

*This might be a functional definition of love.

to lab after dinner and sometimes throughout the night. I am grateful to him for feeding me dried mango and nuts on some of those occasions so that my hands could be free to solder, to turn nutpieces, to drill or dissect; I am grateful to him for delivering midnight pizza and early-morning bagels to Charlotte, Jason, Hexin, Clay, Dan, and me (in various combinations) to eat during hard surgeries. I am grateful to him for giving of his time and his computing cluster's time, and for bringing me along to CERN, the Aspen Physics Center, Nebraska, and the national labs of the bay area when I needed a break and the jolting perspective of a different science. My thinking has been shaped by our hundreds of hours discussing the goals, the practice, and the sociological derangements of research. His unique balance of focus (patience) and momentum (impatience) in his own work have been instructive for me, and his boldness and strange intelligence inspiring. He has sacrificed time and attention on his own work to support, understand, and be so interested in mine that he asks whether I have made the next plots yet.

This work is dedicated to my father, who believed in me, taught me, gave me the freedom and protection to discover and be myself, never tolerated nonsense, made me dinner every night, signed me up for then drove me to philosophy classes and soccer practices, read the books I insisted he read so he would *finally understand me* (Madeliene L'Engle, Sharon Creech, Sarah Dessen, Megan McCafferty, Muriel Barbery, Jia Tolentino, Elena Ferrante), oriented me to science, oriented me to neuroscience, taught me to run, extended his love to the people I love, never believed confidence should be higher on hard than easy error trials, instilled in me old-fashioned scientific sensibilities, taught me how to use Matlab cell arrays, advised me on how to lead a scientific team, and modeled intellectual courage, lifelong curiosity, discipline, humor, and joy. He has always said that every PhD is different, but that the defining feature is that you shouldn't know at the start whether you can do it, or whether it can be done. *If you can see the end at the start, it's not worth doing.* It would not have been possible to attempt such a thing without him.

Contributions

All chapters were written by Hannah R. Joo (HRJ) and subsequently edited by Loren M. Frank (LMF) and others as specified below.

Contributions to chapter 2 include: *Conceptualization* - HRJ, Jason E. Chung (JEC), Adam Kepecs (AK), LMF; *Methodology* - HRJ, Jiang Lan Fan (JF), JEC, Hexin Liang (HL), Charlotte Geaghan-Breiner (CGB); *Software* - HRJ, JF, JEC, HL, Benjamin P. Nachman (BPN), CGB; *Formal analysis* - HRJ, HL; *Investigation* - HRJ, HL, CGB; *Resources* - LMF, AK; *Data curation* - HRJ, HL, KF, JEC; *Writing, original draft* - HRJ; *Writing, reviewing and editing* - HRJ, JEC, BPN, HL, JF, LMF, AK; *Visualization* - HRJ, HL, BPN; *Supervision* - AK, LMF; *Project administration* - HRJ; *Funding acquisition* - LMF, AK, HRJ, JEC.

The content of chapter 3 is reproduced largely as it appears in: H. R. Joo, J. L. Fan, S. Chen, J. A. Pebbles, H. Liang, J. E. Chung, A. M. Yorita, A. C. Tooker, V. M. Tolosa, C. Geaghan-Breiner, D. K. Roumis, D. F. Liu, R. Haque, L. M. Frank, *A microfabricated, 3D-sharpened silicon shuttle for insertion of flexible electrode arrays through dura mater into brain*, J. Neural Eng. 16 (2019) 066021.

Contributions to chapter 3 include: *Conceptualization* - HRJ, JLF, SC, JEC; *Methodology* -

HRJ, JLF, SC, JEC; *Formal analysis* - HRJ, JLF; *Investigation* - HRJ, JLF, JP; *Resources* - LMF, VT; *Data curation* - JLF; *Writing, original draft* - HRJ; *Writing, reviewing and editing* - HRJ, JLF, JEC, SC, LMF, DR, DL; *Visualization* - HRJ, JLF; *Supervision* - LMF; *Funding acquisition* - LMF, HRJ, SC, VT.

The content of chapter 4 is reproduced largely as it appears in: H. R. Joo and L. M. Frank, *The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation*, Nat. Rev. Neurosci. 19 (2018) 744.

Contributions to chapter 4 include: *Conceptualization* - HRJ, LMF; *Writing, original draft* - HRJ; *Writing, reviewing and editing* - HRJ, LMF; *Visualization* - HRJ; *Supervision* - LMF; *Funding acquisition* - LMF, HRJ.

Confidence in Memories: Behavioral and Neural Approaches

Hannah Reade Joo

Abstract: The brain computes and uses uncertainty to guide decision-making. While this is well established for information sensed externally in the form of perceptions, it is less established whether information retrieved from internal storage, in the form of episodic memory, is also treated probabilistically. To test this question, we developed a spatial episodic memory task in which rats gamble their time on a memory choice in each trial, indicating their confidence in its accuracy. We found that rats express higher confidence on correct trials than errors, indicating a degree of self-reflective consciousness thought previously to exist only in humans. We introduce a generative model for episodic memory confidence that predicts the observed patterns of memory confidence.

To investigate the neural correlates of memory confidence, we implanted four rats with triple-site, local field potential (LFP) and single-unit recording devices targeting the orbitofrontal cortex (OFC), nucleus accumbens (NAc), and dorsal hippocampus. To perform these surgeries, we developed a novel method for the implantation of thin-film polymer electrode arrays through the dura mater. We demonstrate that this technology can yield long-term, high quality single unit and LFP recordings.

To investigate the neural activity in these three regions as it may relate to memory confidence, we took as a starting point the decades-old observation that the hippocampus is required for memory, and the more recent finding that hippocampal neurons store and send information about past experience to the rest of the brain. In particular, a hippocampal neural activity pattern known as the sharp wave-ripple (SWR) is an LFP event associated with highly synchronous neural firing in the hippocampus and modulation of neural activity in distributed brain regions. A growing body of evidence indicates that SWRs support both memory consolidation and memory retrieval. This work is summarized in a synthetic review that introduces the perspective that the SWR may mediate the retrieval of stored representations that can be utilized immediately by downstream circuits in decision-making, planning, recollection, or a confidence evaluation, while simultaneously initiating memory consolidation processes. Finally, a proof-of-concept study of SWR function in the episodic memory task is presented.

Contents

List of Figures

List of Tables

Chapter 1

Introduction & Overview

Memory is essential to cognition[1] and is impaired in a wide range of neuropsychiatric disorders[2–7], but we still have only a preliminary understanding of its implementation in the brain. This may be because memory is one of many elements required for behaviorally useful cognition, and therefore cannot be fully understood except in relation to other cognitive processes. In particular, there are many behavioral scenarios in which decisions must be made using incomplete information from memory. Memory corruption is possible at every stage of memory processing, including encoding, consolidation, and retrieval[8–11]. Memory confidence is therefore necessary in mitigating losses from potentially incorrect decisions[12]. This can be achieved by reducing the resources invested in a decision according to an introspective assessment of memory quality (just as we evaluate sensory evidence – *do those look like rainclouds?* – we assess information stored in memory – *did yesterday's weather report predict rain?*). In addition to its indirect role in disease by its influence on memory-based decisions, a failure of memory confidence itself can account directly for psychiatric symptoms. For example, a patient with obsessive compulsive disorder (OCD) may wash her hands repeatedly because she feels unclean despite full knowledge of washing previously, but it is also possible that she may have abnormally low confidence in her

memory of having washed[13,14]. A similar theory proposes a memory confidence deficit in delusions and psychosis in schizophrenia[15–17]. Despite its clinical[18] and fundamental importance, it is not known how memory confidence is computed at the level of neural firing activity. The unifying aim of this work is to establish a foundation for further investigation of a memory confidence computation in the brain and its influence on behavior.

This work is presented in four sections: a novel behavioral task to study memory confidence; technology development for recording single unit neural activity in multiple brain regions stably over months; a review and perspective on current evidence for the hippocampal sharp wave-ripple (SWR) as neural correlate of memory; and preliminary investigation of the relationship between the SWR and memory confidence.

Behavior: In a deviation from the traditional hippocampal, brain-first (that is, 'inside out'[19]) approach applied very successfully by my mentor, Dr. Loren Frank, the starting point for this work has been behavior. We developed a task in which rats could use knowledge of memory accuracy to maximize reward. A rich history of psychological studies of confidence is focused almost exclusively on perceptual confidence, and frames it as a complex "metacognitive" computation indicative of self-awareness and available to only a subset of species, including rats[20,21], monkeys[22–24], dolphins[25,26], and human subjects[27–30]. Confidence in memories is significantly less studied than confidence in perceptions, even in human subjects, and it is unknown whether non-primate species can compute it at all. Thus, the first goal of designing a behavioral task was to test whether rats could compute memory confidence; the second goal was to investigate its neural correlates using the rat as a model system. Presented in chapter 2 is the final version in a series of three behavioral tasks I designed to test whether rats can access memory confidence. For posterity: my approach has been to personally observe every trial of every epoch or to have a technician very familiar with the task (Hexin Liang or Charlotte Geaghan-Breiner) do so. My current

perspective is that this type of observation is invaluable at our current state of understanding of brain and behavior, particularly for complex tasks. Our direct observation of behavior led to the analysis and modeling of putative memory confidence in rats presented in this dissertation.

Technology Development: After establishing the behavioral task, we were prepared to record from the multiple brain regions that we hypothesized could contribute to memory confidence. To investigate memory confidence at the level of neural activity, we focused on the hippocampus (HPC), which is required for memory[31,32] and is known to store and send information about past experience to the rest of the brain[33]. We focused on two of those other brain areas, the orbitofrontal cortex (OFC) and nucleus accumbens (NAc). Selection of the OFC was motivated by previous findings that the OFC can represent confidence in perceptual discrimination tasks[34], and that inactivation of OFC has resulted in a confidence deficit with no effect on choice accuracy[35]. Consistent with single-neuron representations of confidence seen in rat OFC, lesion and imaging studies in human subjects indicate the anterior frontal cortex in appraisal of uncertainty[18]. These studies investigate confidence in the context of perceptual discrimination; electrophysiological studies of confidence in memories are limited, restricted to human and non-human primates[36,37]. The foundation of our choice to record in OFC is the hypothesis that it may represent confidence not only across sensory modalities but in information retrieved from memory. Our working hypothesis has been that neural firing in OFC corresponds to memory confidence, and that behavioral confidence is influenced by hippocampal activity representing past experience. We included in our investigation the NAc for its responsiveness to reward and its known coordination with hippocampal activity[38].

The technology development work in chapter 3 was necessitated by the aim to record single units simultaneously in three brain regions over a data collection period spanning months. For this, we used a novel recording technology that combines flexible polymer probes and a

tetrode hyperdrive. To increase single unit yield in our electrophysiological recordings and reduce operative time, we developed a novel delivery strategy for polymer probes through the dura mater directly to their final recording sites. The force measurements were carried out with technicians Kevin Fan and Jeanie Pebbles, and the *in vivo* demonstration surgery and data collection with technician Hexin Liang. Together we validated that this approach would result in single unit recordings at 90 days post-implant. As part of this work, I used and gave feedback for ongoing development of Mountainsort 4.0, a spike sorting algorithm based on the original Mountainsort 3.0[39]. Together with Hexin Liang and Charlotte Geaghan-Breiner, and with guidance from Jason Chung, I implanted high-density polymer arrays into OFC (128 channels) and NAc (64 channels) and 24-tetrode hyperdrives (96 channels) to hippocampal CA1 to record population and spiking activity. We collected full combined neural and behavioral datasets from the four animals included in chapter 2 over an approximately three-month recording period each.

Synthetic review of the hippocampal sharp-wave ripple: In perceptual discrimination, confidence depends on the quality of sensory evidence (*e.g.*, the more obvious an odor cue, the more confident a subject is)[34,40]. We hypothesize a similar principle for hippocampal memory-dependent decisions, such that the computation of memory confidence is dependent on the accuracy and robustness of a memory trace. Such a link between the neural representations of memory and of memory confidence has not yet been found. Simultaneously with the technology development, I considered which neural phenomena to focus on as a potential memory substrate. A promising candidate mechanism for memory is the hippocampal sharp wave-ripple (SWR), an event corresponding to the synchronous activity of neurons representing recent past experience[33] that is a longstanding focus of the Frank Lab.

SWRs influence activity brain-wide (including OFC and NAc)[38,41–48], and have been proposed as a candidate mechanism for the rapid storage of memories in hippocampal

targets such as the prefrontal cortex[8,49–54]. Consistent with a role in memory consolidation, the incidence of SWRs is correlated with experiences thought to require learning, such as reward[55,56] and novelty[57]. SWRs have also been reported to structure place cell activity representing paths ahead of an animal's current location, potentially as a mechanism for memory retrieval in decision-making and planning[48,58–62]. Moreover, it has been shown that SWRs are necessary in the awake state for accurate choices[63,64]. Based on this evidence, we hypothesized that SWRs could be responsible for transmission of mnemonic information to a downstream memory confidence computation. Developing specific hypotheses as to when SWRs would be important and specifically what function they would serve required a review of the literature, which led to clarification and synthesis of what has been discovered so far regarding SWRs in consolidation and retrieval. This work constitutes chapter 3. Finally, we investigated the possible function of **SWRs in relation to memory confidence behavior**, presented in appendix A.5.

May 2020, shelter-in-place

Berkeley, California

Chapter 2

Episodic memory confidence in the rat: a novel behavioral task and memory model

Abstract: The brain computes and uses uncertainty to guide decision-making. While this is well established for information sensed externally in the form of perceptions, it is not known how information retrieved from internal storage, in the form of episodic memory, may also be treated probabilistically. To test this question, we developed an episodic memory task in which rats must visit the one of two cued locations on a six-arm maze that was visited longer ago in an ongoing sequence. On each trial, rats gamble a variable, self-determined length of time on the choice. For correct choices only, a reward amount dependent on the gambled time is delivered. This allows rats to maximize overall reward by gambling more on choices that are based on more accurate memories of where they went at what time. Taking gambled times as a graded, trial-by-trial behavioral readout of confidence, we found that rats express higher confidence on correct trials than errors. Using a modern machine learning approach, we defined a synthetic memory difficulty axis

and found that rats show a graded pattern of confidence as a function of memory difficulty. We introduce a generative model for episodic memory confidence that represents memories as distributions with asymmetric noise that increases elapsed time since the episodes. Due to its unique, time-dependent noise profile, the model differs significantly from previous signal detection theoretic models of perceptual confidence. The model predicts the observed patterns of memory confidence. Evidence of an ability to assess confidence in episodic memories implies animal *autonoetic*, or self-reflective, consciousness, previously thought to exist only in humans.

2.1 Introduction

An organism can access information from two sources: the external world as it exists in the present, and memory, the internal store of past experience[65]. Neither source is veridical: green can be perceived as blue[66]; an image never seen before can be mistaken for one that was. Some of this inaccuracy is due to adaptive, systematic shifts in perception or memory based on prior experience that bias them toward more likely scenarios (for example, by Bayesian inference in predictive coding[67]). In addition to this systematic error, statistical error results from limitations of perceptual or mnemonic processes. Metacognitive monitoring of this error, in the form of uncertainty in the perception itself or confidence in subsequent decisions[68–70], can valuably inform future action[71,72]. For instance, uncertainty can drive information seeking[73–75] or decreased resource investment in decisions[76]. For information perceived directly, in real time, estimates of uncertainty in sensory percepts are known to influence behavior in adult humans[77,78], dolphins[25], macaques[22–24,36,79–83], capuchin monkeys[84], honey bees[85], and rats[20,34,35,81,82,86]. Theoretical models of perceptual uncertainty based on signal detection theory (SDT) can successfully describe choice uncertainty[40] and, in primates, its dynamics[78,87,88]. Whether information sourced

internally, from memory, is also treated probabilistically, is less established[36,89]. The focus on confidence in perception is so exclusive that the term *decision confidence* is often used interchangeably in the literature with *perceptual decision confidence*[81,86,90–92].

Like perception, memory is unreliable. Its degradation is time dependent. While there is no such thing as an 'old' percept, an old memory, formed long ago, will persist and evolve over time. As it fades and loses fidelity, its inaccuracies can lead to suboptimal decisions. Memory inaccuracies likely result from propagation of perceptual errors in addition to errors in memory processes proper, including encoding and retrieval in consolidation or use[93,94]. Monitoring these processes directly, or monitoring memory quality by some other mechanism, would confer the same adaptive advantages that exist for perceptual uncertainty[70]. There is evidence that humans[95] and primates[36] access confidence in certain types of memory, and that an inability to do so can result in neurological or psychiatric symptoms[18]. Whether non-primate species estimate uncertainty in their memories or compute confidence in memory-based decisions, and whether these exert control over behavior, are completely unknown. Indeed, it was long believed that confidence computations in general may be restricted to primates or humans[96–99].

Given the relative complexity of memory, the studies that establish non-primate perceptual confidence do not necessarily imply an ability in these same species to also compute memory confidence. This idea, that an organism may be able to compute confidence in perceptual discrimination but not memory, originated alongside the concept of episodic memory itself[100,101]. In Tulving's original formulation, three different forms of phenomenal subjective experience emerge from metacognitive monitoring of different memory systems: from making judgments of external stimuli in the present moment, based on procedural memory, anoetic consciousness emerges; from retrieving semantic memories, noetic consciousness emerges; and from retrieving episodic memories, or mentally simulating the future, autonoetic consciousness emerges. Only the third level, autonoetic or 'self-reflective'

consciousness, is proposed to require — and to imply — conscious self-representation, or self-awareness[65,102,103]. This form of consciousness has been argued to exist only in humans[99,104].

Studies that have identified neuronal correlates of memory confidence in human and non-human primates have focused exclusively on visual[36,37,80] or object[105] recognition memory[106]. In a recent review of human neuroimaging studies of metacognition, Vaccaro and Fleming identified thirty published metamemory studies[95]. The majority of these were also recognition tasks, for words[107], face-name associations[108], facts[109], or visual scenes[110]. Even those studies that described the task as requiring episodic memory did not require retrieval of personally experienced events in their temporal and spatial contexts[111,112], the closest being a test of memory for events in a recently watched film[113]. Although there is ongoing debate regarding their relative contributions, familiarity (or *knowing*) and recollection (or *remembering*, equivalent to episodic recall)[101] are both known to support recognition memory. An isolated test of animal autonoetic consciousness, in contrast, requires a behavioral task dependent on episodic memory retrieval[114]. If such an animal model could be established, it could be used to interrogate the neural mechanisms underlying this important function.

Here, we developed a task to study confidence in memories for personally experienced events in their temporal and spatial contexts. In other words, high choice accuracy on the task requires knowledge of 'where', 'when', and 'what', defining features of episodic-like memories in animals[99,115,116]. The task design prevents it from being solved using a decaying memory trace strength/familiarity. Given the evidence that such memories are supported by the hippocampus[117], we designed our task for the rat, in which the hippocampus and its contribution to temporal and spatial sequences in memory are well established[33,117,118]. Other advantages of the rat as a model system for such a study are the ability to perform relatively many trials of a complex task that takes place in real space,

and the recording technology for high density recordings in freely moving animals[119,120].

Of the various methods used to elicit confidence ratings in non-verbal subjects (including human infants), post-decision wagering has advantages[76,121]. For instance, it demonstrates both metacognitive monitoring and control of behavior[122]. Post-decision wagers are typically binary, providing high or low-bet options, but verbal subjects can be instructed to report confidence on a continuous scale (for example, by using a button press of self-determined length[123]). To measure confidence on a continuous scale in non-verbal subjects, researchers have used willingness to wait for a delayed reward of fixed size that is delivered only following correct choices[34,35]. This yields a graded confidence report on all error trials and a subset of correct trials on which reward is withheld, allowing experimenters to determine how long the rat would be willing to wait before leaving. Here we use a new form of confidence report in which rats make a memory-based choice, then gamble their time on the outcome. Time gambling provides a graded confidence report on every trial, regardless of outcome. In combination with an episodic memory logic, time gambling as a confidence report allowed us to study whether a non-primate species, the rat, uses metacognitive memory retrieval to guide decision-making and resource investment.

2.2 Results

2.2.1 The episodic memory confidence reporting task requires memory for the order of past spatial trajectories

In our trial-based task, the use of memory confidence is required to maximize reward. Each trial requires a memory-based decision, then presents an opportunity for the rat to gamble its time on the outcome of that decision. A larger reward amount is delivered for longer gambled times, but only for correct choices. If rats are able to compute confidence in their

memories, they should gamble more time on choices that are based on memories in which they have higher confidence.

The task takes place in fixed, hour-long epochs, with self-paced trials. On every trial of the normal task logic, two of six total ports are cued (Figure 2.1). The correct choice, or *target*, is the one that was last visited longer ago in the cumulative, ongoing sequence. The *distractor* is the one that was last visited more recently. To indicate a choice, the rat must visit and nose-poke at the port. This visit is then appended to sequence of past visits. If the rat makes an error, it is this choice, not what would have been the correct choice on that trial, that is appended to the sequence of past visits. The task differs from a recognition memory task or classic radial arm maze sequence tasks in that the rat does not visit a series of spatial locations, experience a delay, then make judgments about the order of those visits in a brief sequence. Rather, the sequence is ongoing for the entire hour-long epoch. This requires rats to judge which of two episodes occurred earlier in the ongoing sequence.

This logic is implemented on a large, branched maze with a central home port, a back port, and three main branches, which bifurcate into six stems total, each of which has a choice port at its end (Figure 2.1). Every port can be cued with a light and deliver liquid milk reward. During an initialization period, the rat is cued to visit each of the six choice ports in a randomly generated order. After every port has been visited at least once, the normal task logic begins.

Target and distractor are selected randomly on every trial, with two restrictions. First, there is a spatial restriction on target types (Figure 2.1). The target and distractor are always located on topologically adjacent maze arms; A/B and B/C are allowable pairs, but A/C is not. This results in six total spatial trial types. Three of these, A/B, C/D, and E/F, are *stem trials*, pairs with trajectories that share a branch but differ by the stem portion of the trajectory; the three others, B/C, D/ E, and A/F, are *branch trials* that cue

pairs with trajectories that differ in both branch and stem. The branch trial A/F is special among these in that the angle between the trajectories is larger, and they are farther apart in space. With the exception of the branch type A/F, stem and branch trials do not differ in the visual angle or Euclidian distance between the cues.

The second restriction on trial types is temporal, and restricts distractor and target ages. The distractor is selected such that its last visit (*age*) is always 1, 2, or 3 trials ago. By definition, the target's age is always higher than that of the distractor. Note, however, that this does not mean that targets are always age 4 or greater: distractor-target pairs 1-2, 1-3, and 2-3 are allowable. Because of this, the rat cannot simply remember and avoid the ports aged 1, 2, and 3 on every trial, as targets aged 2 and 3 may be presented as targets. By presenting distractor-target pairs spanning a range of ages, we could test choice accuracy as a function of how long ago the to-be-remembered episodes occurred. Because the presented target and distractor can be influenced by the rat's particular sequence of previous visits, we checked for each rat that a balance of trial types was presented (Figure 2.7). The proportion of trials with distractor ages 1, 2, and 3 is approximately one third each. The distribution of presented distractor-target age pairs was the same for each port as the overall distribution. Figure 2.1D shows the full trial sequence for a correct (top row, blue) and an error (bottom row, red) choice.

2.2.2 Rats have the option to gamble their time following each trial, indicating confidence in the memory-based choice

Once at the port of its choice, the rat has an option to gamble its time by maintaining the nosepoke position for a variable, self-determined length of time (Figure 2.1). Withdrawal from the reward port terminates the gambling period. For correct choices only, a liquid reward amount dependent on the length of gambled time is delivered. This reward amount

is delivered at the end of the investment period at the choice port, and at the back port at the end of the trial. The reward amount dependence on gambled time allows the rat to maximize overall reward by gambling more time for choices that are based on memories it assesses to be more accurate. Longer gambled times have a higher possible reward payout in the case of a correct outcome, but a higher penalty in the case of an incorrect outcome. By assessing whether gambled times are higher for choices that turn out to be correct, we test whether they use memory confidence to guide their gambling behavior. Importantly, because distractor and target ages are not directly observable 'first order' sensory stimuli, rats cannot simply use an associative strategy, gambling more time on specific stimuli that also have higher accuracy rates [124].

This form of confidence report tests two aspects of metacognition at once, both monitoring and control[125]. It also overcomes limitations of previously used post-decision wagering paradigms including those that use invested time[35]: it is graded rather than binary, available on every trial, and relatively more ethological than selection of a token on a screen. Furthermore, no new information regarding the probability of choice outcome is gained by the subject over the investment period, in contrast to post-decision wagering tasks in which a delayed reward is eventually delivered on correct trials. Preventing acquisition of new information in the post-decision period is important for testing whether confidence evolves following the decision, even in the absence of new external evidence (*i.e.*, as in post-decisional locus models of confidence)[28,126,127].

The reward function (Figure 2.1) was designed to counter the possible effects of temporal discounting on gambled times. Like humans, rats prefer smaller rewards sooner to larger rewards later[128,129]. The expected effect of such temporal discounting is that rats would reduce their gambled times to receive a smaller reward sooner rather than waiting for a larger one. This effect may be greater on trials where they are highly confident in their memories and choice, as the option of a smaller reward sooner is more certain. This effect

could obscure the difference between gambled times on correct and error trials by inducing a left shift of gambled times on correct trials. To counter this possible effect, the reward amount delivered was a piecewise function of gambled time with a relatively low derivative for the first 2.2 seconds and a relatively high derivative after 2.2 seconds (see Methods). To ensure a high enough number of trials per epoch to sample trial types evenly, we discouraged extremely long gambled times greater than 9.5 seconds by choosing a reward function with a derivative that fell by 9.5 seconds to the level it was prior to 2.2 seconds. A non-zero intercept ensured that the rat received an appreciable reward amount (350 ms, or 5 mL) even for very short waits on correct trials, preventing the development of uncertainty in the memory rule itself following a series of short gambled times.

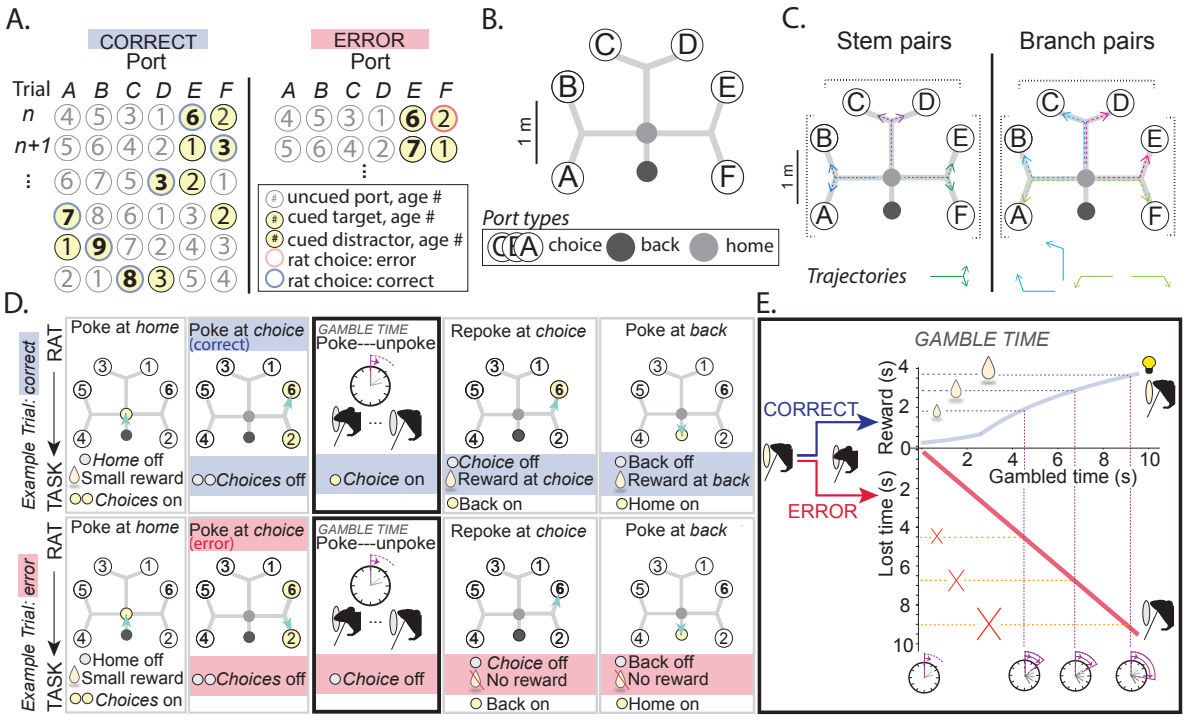

Figure 2.1: **Episodic memory confidence task. A. Memory-based logic.** On each trial (top to bottom), two of ports A-F (top, left to right) are cued with a light (yellow fill, each row). On each trial, the correct choice is to visit the port that was visited more trials ago (*i.e.*, has a higher age, indicated by number in circle, called the *target*). Bold number indicates the target, unbold number indicates distractor. The rat's choice on each trial is indicated by the colored circle (correct choice, purple; error choice, orange). Distractor age is restricted to 1, 2, or 3. On each trial, ages of unvisited ports increment and the last-visited port is set to age 1. Left panel: a sequence of correct choices. Right panel: if the rat made an error to distractor port F (orange outline), then on trial *n+1*, port E would be correct. That is, the correct sequence depends on where the rat visited previously, not what was presented. **B. Task environment.** The task is implemented on an elevated track with back port, home port, and six choice ports (*A*-*F*), located at the ends of three branches that bifurcate into stems. **C. Spatial restriction of trial types.** Cued ports are always adjacent to one another, producing: stem pairs *AB*, *CD*, *EF* (left panel), for which the trajectories to the choices share a branch and differ by their stems (clockwise from left: blue, purple, green dotted lines); and branch pairs *BC*, *DE*, *FA*, which differ in both branch and stem (right panel; clockwise from upper left, cyan, magenta, lime). **D. Example trial, correct (top) and error (bottom).** Trial sequence (left to right) with rat action at top and task response at bottom of each panel, for each of correct (upper row) and error (lower row) choices. Highlighted text indicates events that occur only on correct trials (blue) or error trials (red); events that occur regardless of outcome are not. Panel 1, top and bottom: trials begin with rat (turquoise triangle) at home port (center, gray), where the rat receives a small, fixed reward. Target and distractor choices are cued with lights. Panel 2: The rat runs to and nosepokes at his port of choice, target (port E, age **6** in bold) on a correct trial (top) or distractor (port F, age 2) on an error trial (bottom). For both, the choice lights turn off. Panel 3: at the choice port, the rat maintains the nosepoke position for a variable, self-determined length of time to indicate his confidence that the choice outcome will be correct. When he withdraws, the choice light will turn on again if he was correct (top) but stay unlit if he was in error (bottom). Panel 4: The rat re-pokes at the same choice port, and its light turns off and delivers reward on a correct trial (top) or no reward (bottom). The back port light turns on. Panel 5: The rat then runs to and pokes at the back port, and receives a reward amount equal to that received at the choice port. For an error trial (bottom), this is no reward. **E. Graded, trial-by-trial, confidence report.** When the rat nosepokes at a port on a correct (blue) or error (red) trial, the gambling period begins. The reward amount he receives on a correct trial (top) is a function of how much time he gambles on the choice by maintaining the nosepoke position. On an error trial (bottom) he receives no reward; the time is lost. Withdrawal from the port indicates the end of the gambled time. The outcome of the trial is revealed: on a correct trial (top), the choice port re-lights; on an error trial, the port does not re-light. On a correct trial, the rat can re-poke, still at the choice port, to receive a reward amount dependent on the length of time he gambled.

Rats achieve high choice accuracy overall. We collected approximately 3000 trials from each rat over a period of stable performance accuracy. Each epoch was approximately 50 - 100 trials long, with trial types and errors distributed throughout the epoch (Figure 2.2). Rats achieved a stable performance average of $80.2 \pm .04$ percent accuracy on the task (mean $\pm$ SEM, $N = 192$ epochs pooled from 4 rats), exceeding what could be achieved by selecting uniformly at random between the six ports or between the two cued ports (Figure 2.2). Unlike in a perceptual discrimination task, where trial difficulty can be strictly controlled by manipulating the stimulus features (*e.g.*, the proportion of dots moving left or right; the relative proportions of odorants in a mixture), the memorability of an episode is confounded by various internal factors. For instance, an episode to which the subject paid more attention may be remembered better, but this is not controlled or even observable by an experimenter. A starting point is that memory becomes less accurate over time[130]. We therefore hypothesized that rats would show higher choice accuracy on lower distractor ages, which happened more recently. As expected, the average choice accuracies for distractor ages 1, 2, and 3 respectively were 89.5 ± 0.5, 77.7 ± 0.7, and 72.7 ± 0.7 percent ($N = 192$ epochs pooled from 4 rats; Figure 2.2). The hierarchy of accuracies is statistically significant at $p \ll 10^{-4}$. We hypothesized an additional effect of the difference between target and distractor ages, as it has previously been shown that episodes that occur closer together in time are more likely to be confused with one another[131]. Controlling for distractor age, we found that choice accuracy was higher for older targets, which had an overall range of 2 to about 20. The choice accuracy for each distractor/target pair is shown in Figure 2.2.

2.2.3 Rats show higher choice accuracy on stem than branch trials

Performance accuracy on stem trials was 88.6 ± 0.4 percent, and on branch trials was 71.1 ± 0.8 percent (means $\pm$ SEM across $N = 4$ rats; Figure 2.2). We considered whether the difference in performance could be the result of different strategies applied to the two trial types. On stem trials, cued trajectories are distinguishable only by the time the rat last traversed the stem where the port is located (the stem age). On branch trials, in contrast, cued trajectories are distinguishable by both their branch and stem ages (Figure 2.1). For the majority of branch trials, the trajectory to the correct port is defined by the oldest branch and the oldest stem. However, on approximately twenty percent of branch trials (18.7, 18.4, 17.7, 18.2 percent for rats *T, S, D, R*), the port on the oldest stem is not also on the oldest branch. These 'conflict trials' gave us an opportunity to examine the relative effects of branch versus stem age on memorability.

We examined conflict trials and found that rat choices were more consistent with use of a branch than stem rule: an average of 64.5 ± 1.6 percent of choices were consistent with a branch rule versus 32.8 ± 1.5 percent consistent with a stem rule (average across rats $\pm$ SEM, $N = 4$ rats). That is, on conflict trials, rats more often selected the cue with the older branch, not the older stem. This suggests greater reliance on the memory of branch segments of previous trajectories, of which there are fewer than there are stems, and which are physically longer trajectories. However, rats do not rely solely on a branch age rule. For non-conflict branch trials, rats' performance accuracy was no better predicted by branch age than stem age (logistic regression accuracies and AUCs consistent; $p \approx 0.3$). Because branch and stem trials were presented in random order throughout the epoch, and stem trials require knowledge of the stem history, the rat was required to remember at all times the stem history, not just the branch history. Together these results suggest that recency judgments of episodes may be influenced by memory of each of the different

segments of previously traversed paths.

Excluding conflict trials, performance accuracy was still higher on stem than branch trials (88.6 ± 0.4 vs. 78.0 ± 1.1 percent across $N = 4$ rats). We hypothesize that this could be because stem but not branch trials allow the rat to partially reenact the trajectories leading to the cued ports before making a decision. Intriguingly, there was no difference in accuracy between the different-arm trial types BC and DE versus different-arm trial type AF (Figure 2.2; t-test, 2-tailed $p > 0.5$ for all rats except R, with $p = .0003$), whose ports A and F are separated by a greater distance and wider visual angle (Figure 2.1). This suggests that difficulty depends on whether the distractor and target ports share a partially overlapping trajectory; if they do not, there is no additional effect of how far apart the non-overlapping trajectories are in space.

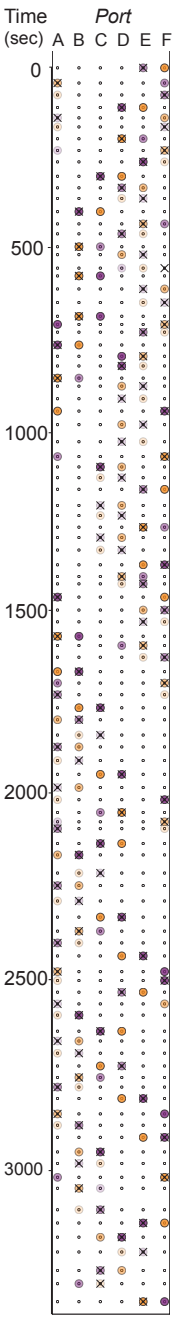

(a)

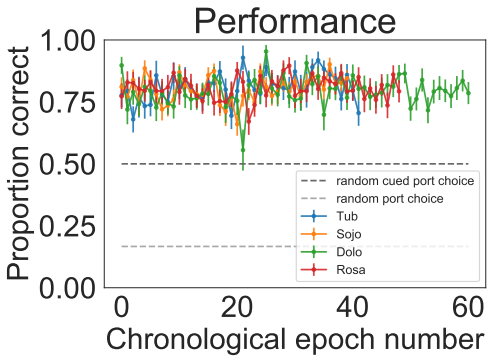

(b)

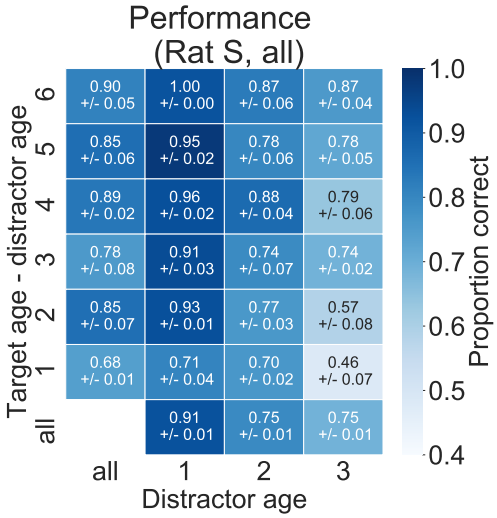

(c)

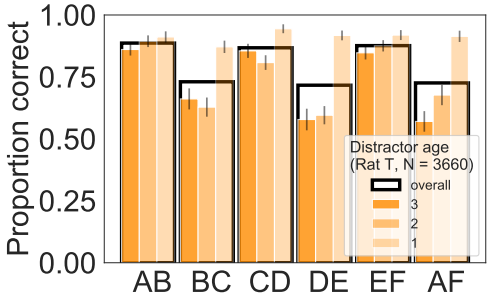

(d)

Figure 2.2: **A. Full, hour-long epoch for one rat.** Y axis: trial start time in seconds. X axis: State for each of the six choice ports. Target (correct) port is purple and distractor (error) port is orange, with lighter shade indicating age (*i.e.*, which occurred longer ago in the sequence). Three orange shades correspond to distractor ages 1, 2, 3; three purple shades correspond to target ages 2-3, 4-6, or 7 or older. The *X* over the port indicates the port selected by the rat (which can be an uncued port, no color). **B. Performance is stable across epochs.** Choice accuracy, defined as the proportion correct, is stable per rat for each epoch included in the dataset. Error bars are standard error for a binomial process. **C. Choice accuracy** as a function of target and distractor ages, representative rat S. **D. Choice accuracy for stem versus branch trials**, representative rat S ($N = 4111$ *trials*). **E. Trial types are balanced over target and distractor age for each port**, representative rat S, ($N = 4111$ *trials*).

2.2.4 Performance accuracy is more consistent with the memory rule than alternative strategies

We tested whether the true memory rule, to visit the least recently visited of the two cued ports, could be degenerate with one or more other strategies that did not require memory and would be easier to apply. Each of ports A-F was equally likely to be presented as a target or a distractor, preventing success of a selection rule that favored particular ports (see Methods). We also verified that choice accuracy was higher than could be achieved by learning a fixed sequence of port visits (Figure 2.8). The probability of any of the other five ports being presented subsequently is approximately equal, at 0.2 (Figure 2.8).

Next, we considered an alternative strategy in which each of the six spatial trial types is treated as a separate alternation task. Under this hypothetical alternative strategy, rats would remember for each of the six port pairs which of the two ports was last visited and choose the other. To test this alternative strategy, we simulated choices that the strategy would dictate for the actual sequence of trials experienced by each rat in each epoch. We calculated the proportion of these trials that yielded a correct choice under the true rule. Such an alternation strategy was not consistent with the true rule, producing a choice accuracy that was below the accuracy achieved by rats performing the task (Figure 2.8). Moreover, this strategy was not consistent with the rat's actual sequence of choices (Figure 2.8). Performing this analysis per epoch allowed us to rule out the possibility that the rat employed an alternation strategy selectively, such as earlier or later in learning.

Using the same approach, we considered and ruled out additional alternative strategies (Figure 2.8) to select from the two cued ports the port that was: farthest from the last-visited, the leftmost, the least visited, the least cumulatively rewarded over the epoch in total, the most cumulatively rewarded over the epoch in total, or the least recently rewarded. The three of these strategies that had the highest agreement with the true rule,

exceeding 50 percent, were to visit the least cumulatively rewarded, the farthest from the last visited, or the least recently rewarded. This is unsurprising, as each of these alternatives is correlated with the true rule. However, none of these strategies was more consistent with rat choices than was the true rule (Figure 2.8).

2.2.5 Spatial type and distractor age are the major determinants of performance accuracy

We next investigated whether features other than spatial trial type, target age, and distractor age influenced choice accuracy less overtly than by an alternative strategy. A secondary reason to identify contributors to choice accuracy is that these can be inferred to also be the determinants of trial *difficulty* and therefore define a memory difficulty axis, along which confidence for each of correct and error trials is expected to vary. To investigate additional effects of features other than spatial and temporal trial type on choice accuracy, we trained a deep neural network (DNN) to predict rat choice based on an exhaustive and a reduced feature set. The exhaustive feature set included the age of the target and distractor ports, the spatial trial type, as well as the last, cumulative and maximum single reward amount received and the last and cumulative dwell time spent at each of the target and distractor (Table 2.3). The reduced feature set included only the target and distractor ages, spatial trial type, and last reward and dwell time at each of the target and distractor.

We trained a separate model for each rat, using a five-fold cross-validation procedure (Figure 2.9). By comparing the performance of the model on the exhaustive and reduced feature sets, we determined whether the additional features significantly improved model performance. We compared model performance using two metrics. First, the average model accuracy in predicting rat choice outcome (*correct, error*) was approximately 70 percent

Figure 2.3: **Table:** $\mathcal{T}$, $\mathcal{D}$, t, and r stand for target, distractor, time, and reward amount, respectively. The abbreviations 'max.' and 'cum.' are short for 'maximum' and 'cumulative', respectively. An 'x' marks the combinations that are part of the full model and the red 'x' denote the combinations of features used in the reduced set. **B. Comparison of DNN and logistic regression models** for full and reduced feature sets by ROC AUC. Error bars = standard deviation across the five folds used in training. **C. Comparison of a logistic regression model trained on different feature sets** including only spatial and temporal features or spatial and temporal features plus last reward amount or spatial and temporal features plus last dwell time **D. Comparison of a logistic regression model trained on the base feature set including spatial and temporal features of time in terms of number of trials versus clock time.**

	1	arm	$\mathcal{T}$ age	$\mathcal{D}$ age	$\mathcal{T}$ age - $\mathcal{D}$ age	last dwell t at $\mathcal{T}$	last dwell t at $\mathcal{D}$	last r at $\mathcal{T}$	last r at $\mathcal{D}$
1	x								
arm	x								
$\mathcal{T}$ age	x	x		x					
$\mathcal{D}$ age	x	x							
$\mathcal{T}$ age - $\mathcal{D}$ age	x								
last dwell t at $\mathcal{T}$	x	x	x				x		
last dwell t at $\mathcal{D}$	x	x		x				x	
last r at $\mathcal{T}$	x	x	x						x
last r at $\mathcal{D}$	x	x		x					
max. single r at $\mathcal{T}$	x								
max. single r at $\mathcal{D}$	x								
cum. r at $\mathcal{T}$	x								
cum. r at $\mathcal{D}$	x								
cum. dwell t at $\mathcal{T}$	x								
cum. dwell t at $\mathcal{D}$	x								
t since last r at $\mathcal{T}$	x								
t since last r at $\mathcal{D}$	x								
t since last visit to $\mathcal{T}$	x								
t since last visit to $\mathcal{D}$	x								

(a)

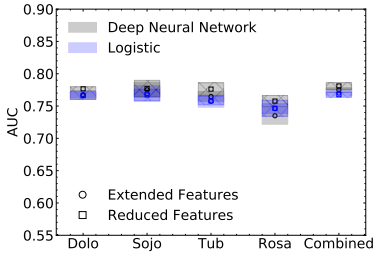

(b)

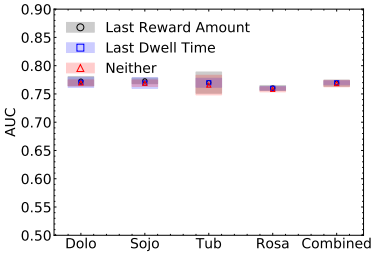

(c)

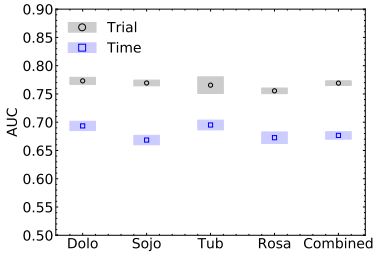

(d)

(Figure 2.9), and the ROC AUC was approximately .75 (Figure 2.9). We also trained and tested a model on a combined dataset across rats, which yielded similar results. The primary advantage of a DNN was its capacity to learn interactions we did not explicitly include as interaction terms. For greater interpretability, we also trained a logistic regression model, which had comparable performance accuracy.

We next trained a logistic regression model on subsets of the reduced feature set to test their relative contributions to choice accuracy. We compared a base feature set including five predictors, DISTRACTOR AGE, TARGET AGE, BRANCH/STEM TYPE and the interaction terms BRANCH/STEM × DISTRACTOR AGE, BRANCH/STEM × TARGET AGE to one that included either the last reward amount or the last dwell time at the target and distractor ports. We found that including each of these improved model performance only slightly, suggesting that the most important features in predicting choice accuracy were spatial and temporal (Figure 2.9).

To investigate whether the dependence on distractor age was the result of elapsed trials or elapsed time itself, we compared the performance of two logistic regression models, one with target and distractor ages in units of trials and another with target and distractor ages in units of absolute time (Figure 2.9). For each rat, substituting port age in elapsed time rather than trial count yielded significantly less accurate (AUC) predictions of rat performance accuracy, indicating that what matters is the number of elapsed experiences rather than elapsed time. A similar result has been reported in episodic sequence memory tasks in monkeys[132]. Having identified distractor age in units of trials as the most important determinant of choice accuracy, we could ask whether gambled times predicted choice outcome over a range of memory difficulty levels, which would provide evidence of a memory confidence computation.

2.2.6 Rats use memory confidence to guide time investments

To maximize overall reward in this task, rats should invest more time in choices that are more likely to be correct. Our previous analysis demonstrates that choice accuracy is determined by memory variables, primarily distractor age. If choice outcome depends solely on memory, the only way for a rat to selectively gamble more on choices that turn out to be correct is to introspectively evaluate memory accuracy. Longer gambled times on trials that turn out to be correct would therefore indicate a memory confidence computation.

We found that rats gambled more time preceding correct outcomes relative to errors overall (Figure 2.4; average AUC 0.74 $\pm$ .03 SEM, N=4 rats; for each rat, two-tailed KS test $p < 1 \times 10^{-5}$). On average, gambled times were 1.45 ± 0.33 seconds higher for correct than error trials (average $\pm$ SEM, N=4 rats). Gambled times were higher for correct than error trials even when controlling for trial stimulus features, which in this task are the target and distractor ports themselves. Figure 2.4 shows that for every port pair, gambled time on correct trials was longer than gambled time on error trials (two-sided KS test $p << 1 \times 10^{-4}$; all rats shown in Figure 2.10).

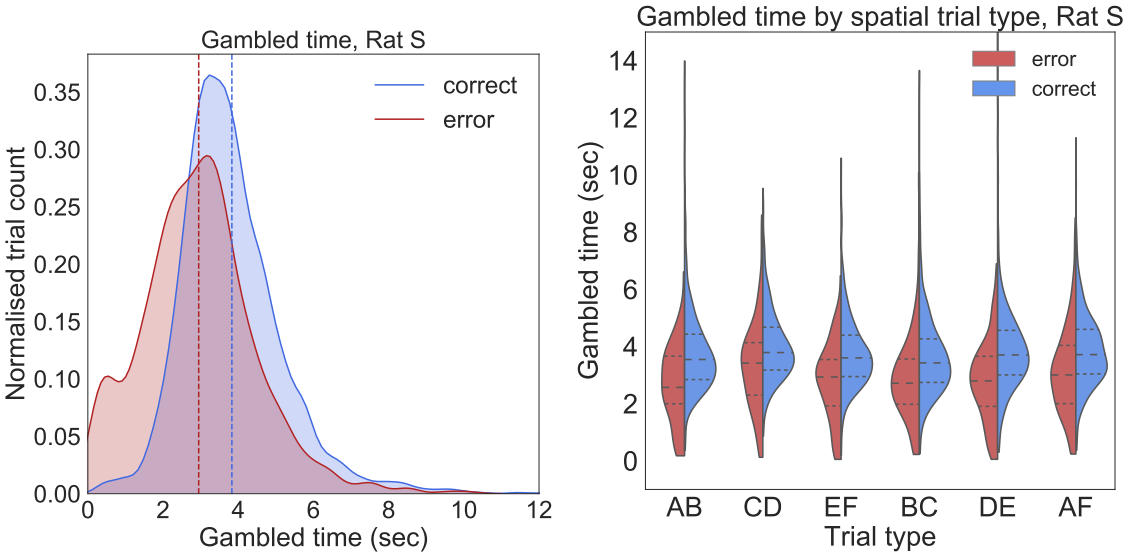

(a)

(b)

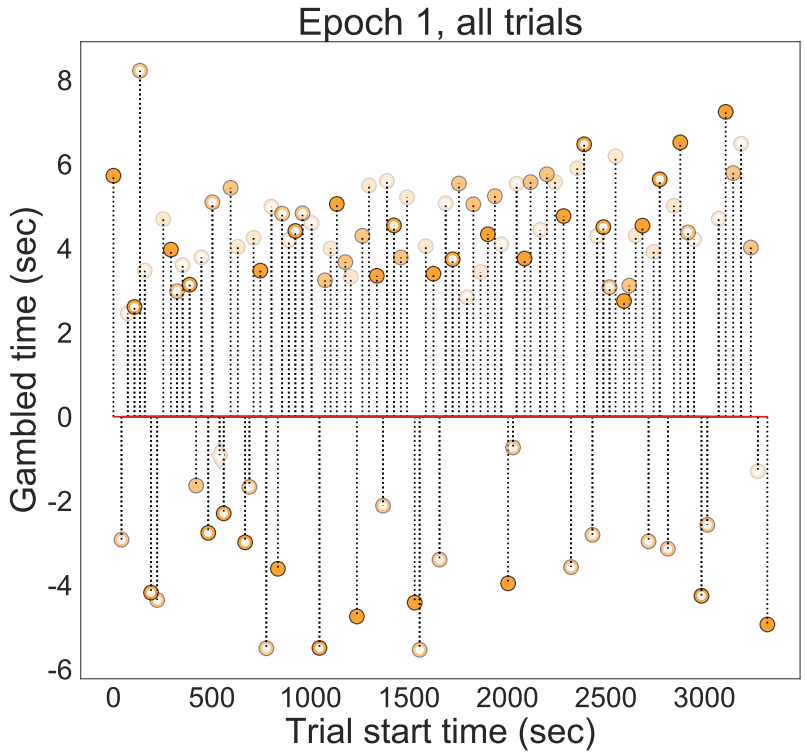

(c)

Figure 2.4: **Gambled times predict choice outcome. A.** Gambled times were significantly higher for correct (blue) than error choices (red), as shown for representative rat S, inclusive over all trials in all epochs (N error trials = 738; M correct trials = 3306). Dashed vertical lines = average gambled times. **B.** Gambled times were significantly higher for correct than error choices within each port pair, representative rat S. **C. Gambled times are stable within epochs.** Each trial is depicted as a vertical line, located on the x-axis at the time it started. Correct trials show lines extending above the x-axis; error trials show lines below the x-axis. The length of each line corresponds to the amount of time gambled on the trial. Marker color indicates distractor age, with darker colors corresponding to harder trials: age 1 is light orange, age 2 is orange, age 3 is dark orange. Open circles = branch trials; filled circles = stem trials. Diamond marker indicates an error to an uncued port.

Previous confidence studies in perception and memory have reported lower latency to choice for trials on which subjects are more confident. In our task, too, we found that the latency to choice was significantly lower on correct than error trials for each rat (average AUC of 0.55 ± SEM 0.26, N=4 rats, and two-sided KS test, $p < 0.05$ for each rat; excludes data with latency to choice above 30 seconds), and 0.69 seconds lower on correct than error trials on average (N = 4 rats). Latency to choice was also lower on correct than error trials for branch and stem trials evaluated separately.

A longer latency to choice can be understood as an information seeking behavior secondary to uncertainty. From this perspective, longer latency to choice on error trials corroborates relatively short gambled times on these trials. To rule out the possibility that gambled times were determined based primarily on trial-by-trial observation of choice latency rather than metacognitive assessment of memory, we found that latency to choice did not significantly predict choice outcome when controlling for memory trial type. Pooling all data, we found only a slight negative correlation between latency to choice and gambled time with $R^2 < 0.05$, explaining little of the spread in gambled times.This indicates that rats do not rely on latency to choice as an external decision variable to determine their gambled times.

2.2.7 Rats gamble less time on error visits to uncued ports

While watching rats perform the task, we observed that they occasionally made errors to a port that was not one of the two cued ports (*uncued errors*). Excluding conflict trials, uncued errors represented an average of 20.9 ± 4.3 percent of total errors (mean and SEM across N=4 rats), evenly divided between branch and stem trials. If rats have learned the component of the rule that the correct choice is always one of the two lit ports, we expect that they should have relatively low confidence on this subset of error trials and thus

gamble short times. Critically, these gambled times would represent confidence in the decision itself but not confidence in the age of the target or distractor. Consistent with this prediction, the time gambled on uncued errors was significantly lower than for other errors (Figure 2.5 and Figure 2.11; average AUC 0.74 ± .01 SEM, N=4 rats; each rat, two-tailed KS test $p < 1 \times 10^{-5}$).

Interestingly, uncued errors had a significantly longer latency to choice than did errors to the distractor port (8.9 ± .7 seconds versus 6.3 ± 0.5, average ± SEM across N=4 rats; t-test for independent populations $p < 0.0005$ for each rat). Given that faster responding can be correlated with higher motivation, attention, and anticipation, it is possible that the longer latency to choice can be explained by low task engagement or foreknowledge that no reward will be received. Alternatively, the longer latency to choice could represent a longer period of deliberation to select between the two cued ports, before rats give up and select an uncued port. The majority (69.1 ± 3.2 percent, N=4 rats) of uncued errors occurred on trials with distractor age 1, with no significant difference in occurrence as a function of target age (Figure 2.5). Given that this is the distractor age with the highest choice accuracy, suggesting it is the easiest, it is possible that the correct answer is known and that uncued errors correspond to an alternative strategy such as intentionally 'throwing' a trial for alternative gain. In particular, these errors may correspond to a type of exploration to ascertain that the optimal strategy is still to follow the learned rule. Gambled times were still significantly different for correct versus error trials after excluding uncued errors (Figure 2.5 and Figure 2.10; average AUC 0.71 ± .02 SEM, N=4 rats; each rat, two-tailed KS test $p < 1 \times 10^{-5}$).

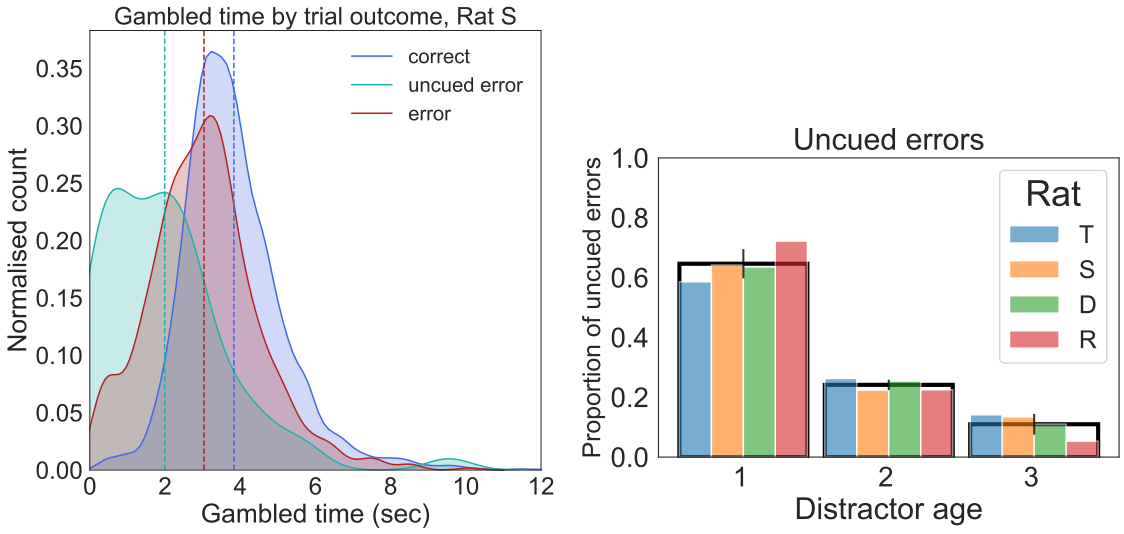

(a) (b)

Figure 2.5: **Less time is gambled on uncued errors. A.** The majority of uncued errors occurred on trials with distractor age = 1, consistent across rats. **B.** Gambled times were significantly shorter for uncued errors relative to other errors. Excluding uncued errors, gambled times on error trials were still significantly lower than gambled times on correct trials.

2.2.8 A generative episodic memory model to predict confidence

Gambled times that predict choice outcome are a strong indication of metacognitive memory recall, but in their strictest interpretation they indicate decision confidence in a memory task, not uncertainty in a memory *per se*. To test whether gambled times were explained based on graded memory confidence, we built a generative model of episodic memory.

In our model, each previous port visit is an episode stored in memory, represented as a distribution centered at the actual visit time with mnemonic noise as variance on the temporal axis. Each trial cues a target and distractor port, with ages M_0 and M_1, respectively. By definition, $M_0 > M_1$. The rat has an internal representation for each of these episodes, M_0' and M_1'. The spread in the distribution of $M_i'|M_i$ is due to mnemonic noise as a result of errors in encoding, consolidation, and/or retrieval. Except for encoding, each of these processes is assumed to occur repeatedly, with cumulative error, resulting in an increasing spread over time. We also observed that it should never be possible to mistakenly retrieve a memory that occurred in the past as having occurred in the future. To satisfy these two requirements, we selected as our noise distribution the lognormal. Specifically, we modeled each M_i' as a lognormal distribution with mean $a_0 M_i$ and standard deviation $\sigma_0(1 + a_1 M_i + a_2 M_i^2)$. The *separation* parameter a_0 sets the degree of separation in memory between two episodes, the *spread* parameter σ_0 sets the baseline spread of each memory distribution, and the coefficients a_1 and a_2 set the rate of change in spread for each memory distribution as a function of its age M_i. By fitting the standard deviation with a second-order polynomial we give it freedom to increase or decrease as a function of time, though our hypothesis is that it should increase. Each memory retrieval can be thought of as sampling from the memory distribution. Figure 2.5 depicts the memory distributions corresponding to a sequence of previously visited ports. When two ports are cued (Figure 2.5), the rat retrieves the time of last visit to each by sampling from the

memory distributions. The rat is correct if the retrieved time of last visit for the target is older than the retrieved time of last visit for the distractor (*i.e.*, $m'_0 > m'_1$). We defined accuracy and confidence:

$$\text{Accuracy} = \Pr(\text{correct}|M_0 = m_0, M_1 = m_1, \sigma_0, \vec{a}) \tag{2.1}$$

$$\text{Confidence} = \Pr(\text{correct}|M'_0 = m'_0, M'_1 = m'_1, \sigma_0, \vec{a}), \tag{2.2}$$

for $\vec{a} = (a_0, a_1, a_2)$. Under this model, confidence is monotonically related to the distance from the decision boundary, $|m' - b|$, and for ease of computation we re-define confidence to be $|m'_1 - m'_0|$.

Our overall approach was to *(i)* fit the model parameters a_0, a_1, a_2 and σ_0 based on choice accuracy for each rat, then (*ii*), from the memory distributions defined by these parameters, generate predictions for memory confidence. Based on our finding that trial count was a better predictor of choice accuracy than elapsed time, we parameterized the temporal axis of the model by episode age in trials.

For each rat, the model parameters were fit to the observed error rates across the 12 trial types with distractor ages $M_0 \in \{1, 2, 3\}$ and the difference between target and distractor ages $M_1 \in \{1, 2, 3, 4\}$, excluding errors to uncued ports. We define accuracy in the data as simply the proportion of correct trials, and the error rate (ϵ_{data}) as 1−Accuracy. The model-predicted Accuracy is given by Eq. 2.1 as the difference of the lognormal random variables M'_0 and M'_1. The probability density of the difference between two lognormal distributions does not have a closed form analytic solution, so we simulated 10^4 trials for each trial type within the fit. Each simulated trial generated an m'_1 and m'_0, from which we computed an outcome (correct or error). Across many simulated trials, this returned a predicted error rate pattern across trial types for the current set of parameters.

A χ^2 metric was used to evaluate model performance and find the best fit parameters:

$$\sum_{\text{trial type } i} \left(\frac{\epsilon_{i,\text{data}} - \epsilon_{i,\text{model}}}{\sigma_{\epsilon_{i,\text{data}}}} \right)^2, \tag{2.3}$$

where ϵ is the error rate and σ_ϵ is the uncertainty in the error rate (see Methods).

For each rat, the model fit to choice accuracy results in memory distributions with increasing spread as a function of time (a_1 and a_2 are positive; Figure 2.5 and Figure 2.12). The model did not require this, as the standard deviation is fit by a polynomial and so could increase or decrease over time. This result is consistent with our expectation of decreasing memory fidelity (or *precision*) with time. The spread parameter, $\sigma_0 \ll 1$, which sets the overlap between neighboring memory distributions, is relatively low as expected given the high choice accuracy of 80 percent correct. For each rat, the fitted value of the separation parameter $a_0 > 1$ (Figure 2.5 and Figure 2.12). A fitted value $a_0 = 1$ would correspond to single trial unit separation in memory between episodes that in fact did occur one trial apart in time. Under our model, time is dilated: the recalled time at which an episode occurred is on average longer ago than the actual time at which it occurred.

The quality of the model fit to the data is shown in Figure 2.5. Model biases (average of (fit-data)/uncert.) for fits to stem, branch, and all combined data are -0.24 ± 0.19, 0.03 ± 0.19, and 0.24 ± 0.21, respectively. As such, model predictions of rat error rates are consistent with the data at $p \approx 0.3$. The standard deviations of (fit-data)/uncert. are 1.28 ± 0.12, 1.29 ± 0.11, and 1.46 ± 0.15 for stem, branch, and stem+branch, respectively. These values are all slightly above one, indicating a slight under-estimation of the uncertainty used in the fit. This is consistent with a small systematic uncertainty, as the fit only includes statistical uncertainties.

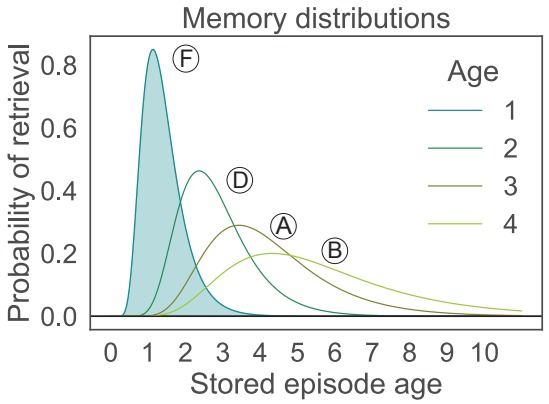
Memory distributions
Probability of retrieval
0.0
0.2
0.4
0.6
0.8
F
D
A
B
Age
1
2
3
4
0 1 2 3 4 5 6 7 8 9 10
Stored episode age

(a)

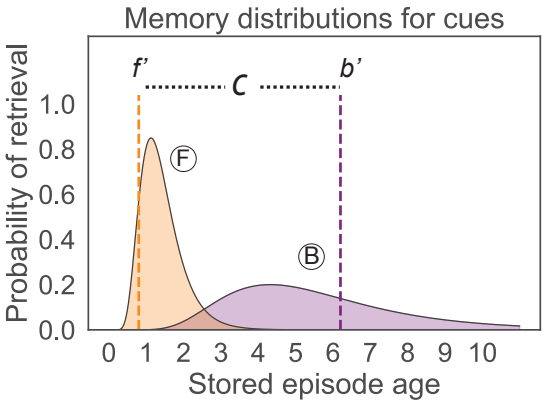
Memory distributions for cues
Probability of retrieval
0.0
0.2
0.4
0.6
0.8
1.0
f'
c
b'
F
B
0 1 2 3 4 5 6 7 8 9 10
Stored episode age

(b)

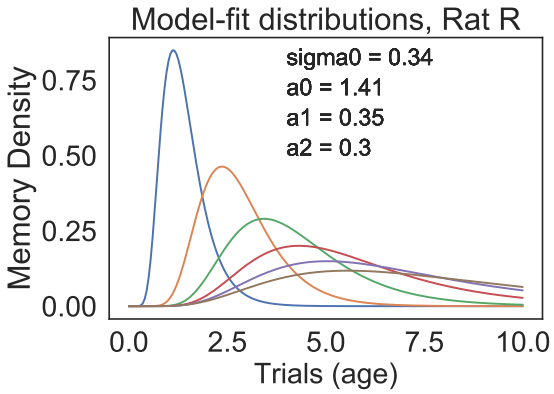
Model-fit distributions, Rat R
Memory Density
0.00
0.25
0.50
0.75
sigma0 = 0.34
a0 = 1.41
a1 = 0.35
a2 = 0.3
0.0 2.5 5.0 7.5 10.0
Trials (age)

(c)

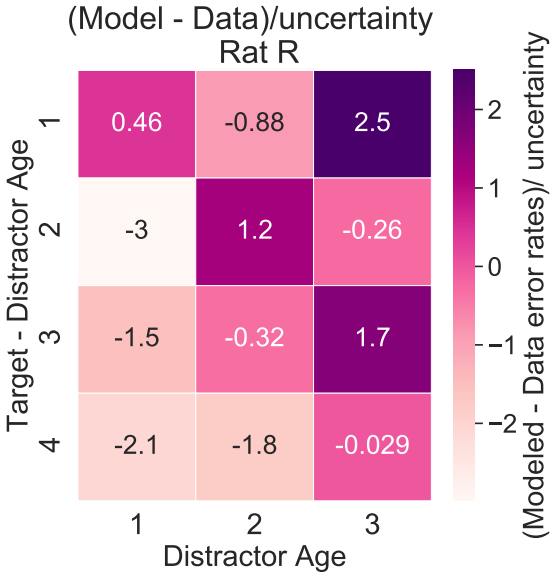
(Model - Data)/uncertainty
Rat R
Target - Distractor Age
1
2
3
4
0.46
-0.88
2.5
-3
1.2
-0.26
-1.5
-0.32
1.7
-2.1
-1.8
-0.029
1 2 3
Distractor Age
(Modeled - Data error rates)/ uncertainty
2
1
0
−1
−2

(d)

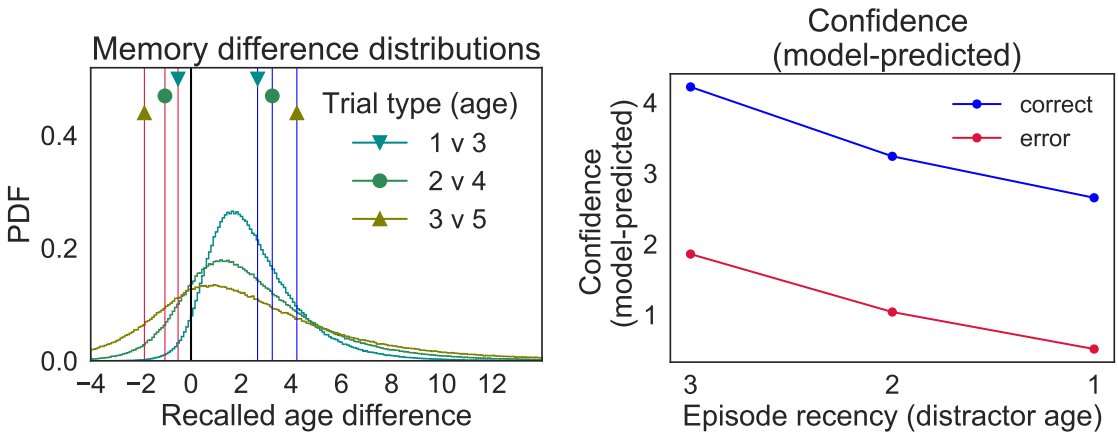

(e) (f)

Figure 2.5: **Episodic memory model.** Each episode, corresponding to a visit to a choice port, is stored in memory as having occurred at some time (its *age*), with a lognormal distribution about its center representing the probability of retrieval for a given age. Distributions are from representative rat R with parameters were fit to choice accuracy. **A. An example sequence of memory distributions.** Most to least recent port visits F, D, A, B are shown (teal filled, green, olive, lime). **B. Example trial.** Ports F and B are cued. The rat retrieves the time of last occurrence for each, f' and b', and compares them. If the sign of F - B is the same as f' - b', the choice outcome is correct; otherwise it is an error. Confidence is defined as the retrieved duration between f' and b', or $|f' - b'|$. **C.** Fitted model lognormal distributions corresponding to the memory distributions (M_i') from which the rat samples in memory retrieval. Each distribution is a different color and corresponds to the memory distribution, from left to right, for the last trial to six trials ago. Fit parameters for representative rat R. **D.** The relative difference between the model and the fit: $(\epsilon_{\text{model}} - \epsilon_{\text{data}})/\epsilon_{\sigma_{\text{data}}}$. **E.** Confidence distributions for DISTRACTOR AGE = 1. Considering trials that present distractor age = 1 (in **A**, teal filled distribution for labeled F) in combination with target age 2, 3, and 4 (green, olive, lime distributions to the right), each target and distractor combination yields a distribution that represents the distribution of target - distractor retrievals (colors matched to target in top panel). The density to the right of zero (black horizontal line) represents those retrievals for which a correct choice will be made, and the average confidence on these trials is the center of this density (blue vertical lines). The density to the left of zero represents those retrievals for which an error choice will be made, and the average confidence on these trials is the absolute value of the center of this density (red vertical lines). **F.** As the inter-episode interval increases (TARGET AGE - DISTRACTOR AGE) increases, the average model-predicted confidence (corresponding to the 'recalled age difference' in E) for correct trials increases, but the average confidence for error trials is virtually constant.

2.2.9 Confidence predictions of the model deviate from Gaussian noise SDT models

Next, we generated specific predictions for confidence as a function of memory difficulty. We first studied the strongest graded predictor of memory difficulty, the distractor age. We generated predictions for accuracy as a function of gambled time, for accuracy on high versus low gambled times as a function of memory difficulty, and, finally, the gambled time on correct versus error trials as a function of memory difficulty. These three predictions correspond to the three signatures of statistical confidence established in perceptual discrimination tasks. Under our model, the third prediction deviates significantly from that of the statistical framework for confidence developed for perceptual discrimination. In these studies, the Gaussian distributions of fixed variance give rise to a characteristic X shape for confidence on correct versus error trials as a function of difficulty (the *vevaiometric curve*). Due to its asymmetric noise profile and non-constant variance, the model predicts a differently shaped vevaiometric curve. Figure 2.5 illustrates three difference distributions corresponding to fixed inter-episode interval (distractor age - target age) 2 for incrementing distractor ages, or the three distractor-target trial types 1-3, 2-4, and 3-5. The model-predicted confidence for correct and error trials is marked on the distributions. Given that choice accuracy is inversely correlated with distractor age, plotting confidence for each of correct and error trials for each distractor age yields predicted confidence trends as a function of memory difficulty. On this axis, gambled times on both correct and error trials are expected to decrease as a function of distractor age (Figure 2.5). These trends are the result of the model form and hold over the range of the specific parameters fit based on experimental rat data.

To test whether these model-predicted trends in confidence were consistent with behavioral data, we converted the predicted confidence values to gambled times. This was done separately for each rat by calibrating the model confidence CDF to the empirical gambled

time CDF (see Methods). The mapping was applied to the entire distribution of gambled times and did not map separately correct versus error trials or any other subset of the data (Figure 2.12). By construction, there is a monotonic relationship between model-predicted confidence and model-predicted gambled time.

2.2.10 The statistical episodic memory model predicts confidence as a function of episode age

We tested whether the observed gambled times were consistent with model predictions as a function of memory difficulty defined by the distractor age axis. Figures 2.5A, C, and E show model predictions with data overlaid for representative rat R, excluding uncued errors and conflict trials. Figure 2.5A shows model-predicted accuracy as a function of gambled time, with data overlaid, for representative rat R. Intuitively, longer gambled times predict higher choice accuracy. Figure 2.5C shows that gambled time is longer for correct than error trials at every difficulty level, with trends for correct and error trials as a function of difficulty in agreement with those predicted by the model. Figure 2.5E shows that for correct trials, gambled times are higher than error trials with a slight negative slope as a function of distractor age, whereas gambled times decrease more dramatically with distractor age. The same qualitative patterns were observed for all rats (Figure 2.13), and for branch (Figure 2.15) and stem trials (Figure 2.17) separately. Figures 2.5B, D, and F show averages across the four rats, with trends corresponding to the model.

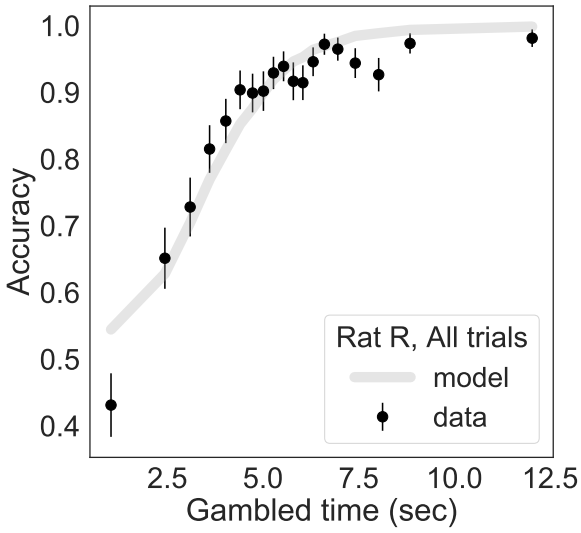
Accuracy
Gambled time (sec)
Rat R, All trials
model
data

(a)

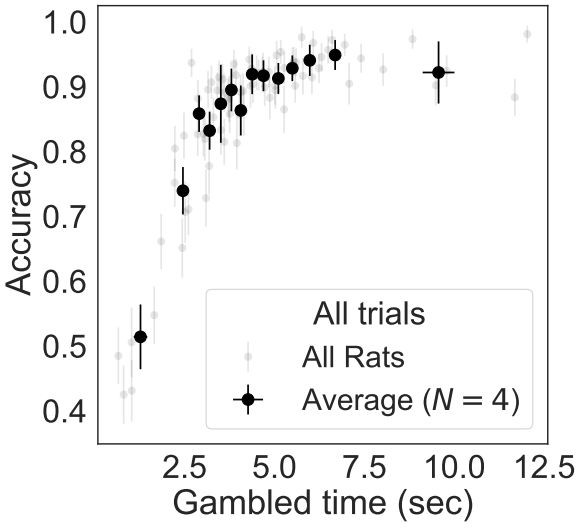
Accuracy
Gambled time (sec)
All trials
All Rats
Average (N = 4)

(b)

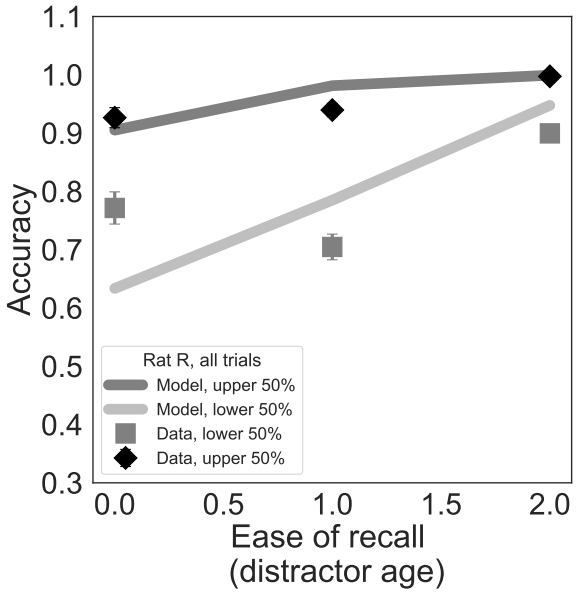
Accuracy
Ease of recall
(distractor age)
Rat R, all trials
Model, upper 50%
Model, lower 50%
Data, lower 50%
Data, upper 50%

(c)

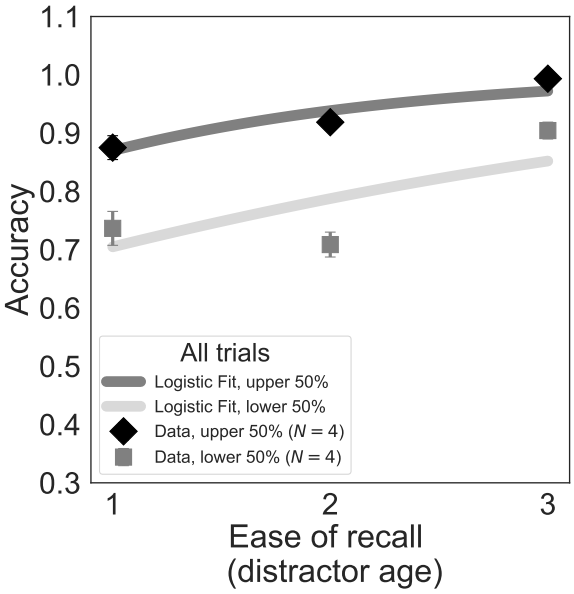
Accuracy
Ease of recall
(distractor age)
All trials
Logistic Fit, upper 50%
Logistic Fit, lower 50%
Data, upper 50% (N = 4)
Data, lower 50% (N = 4)

(d)

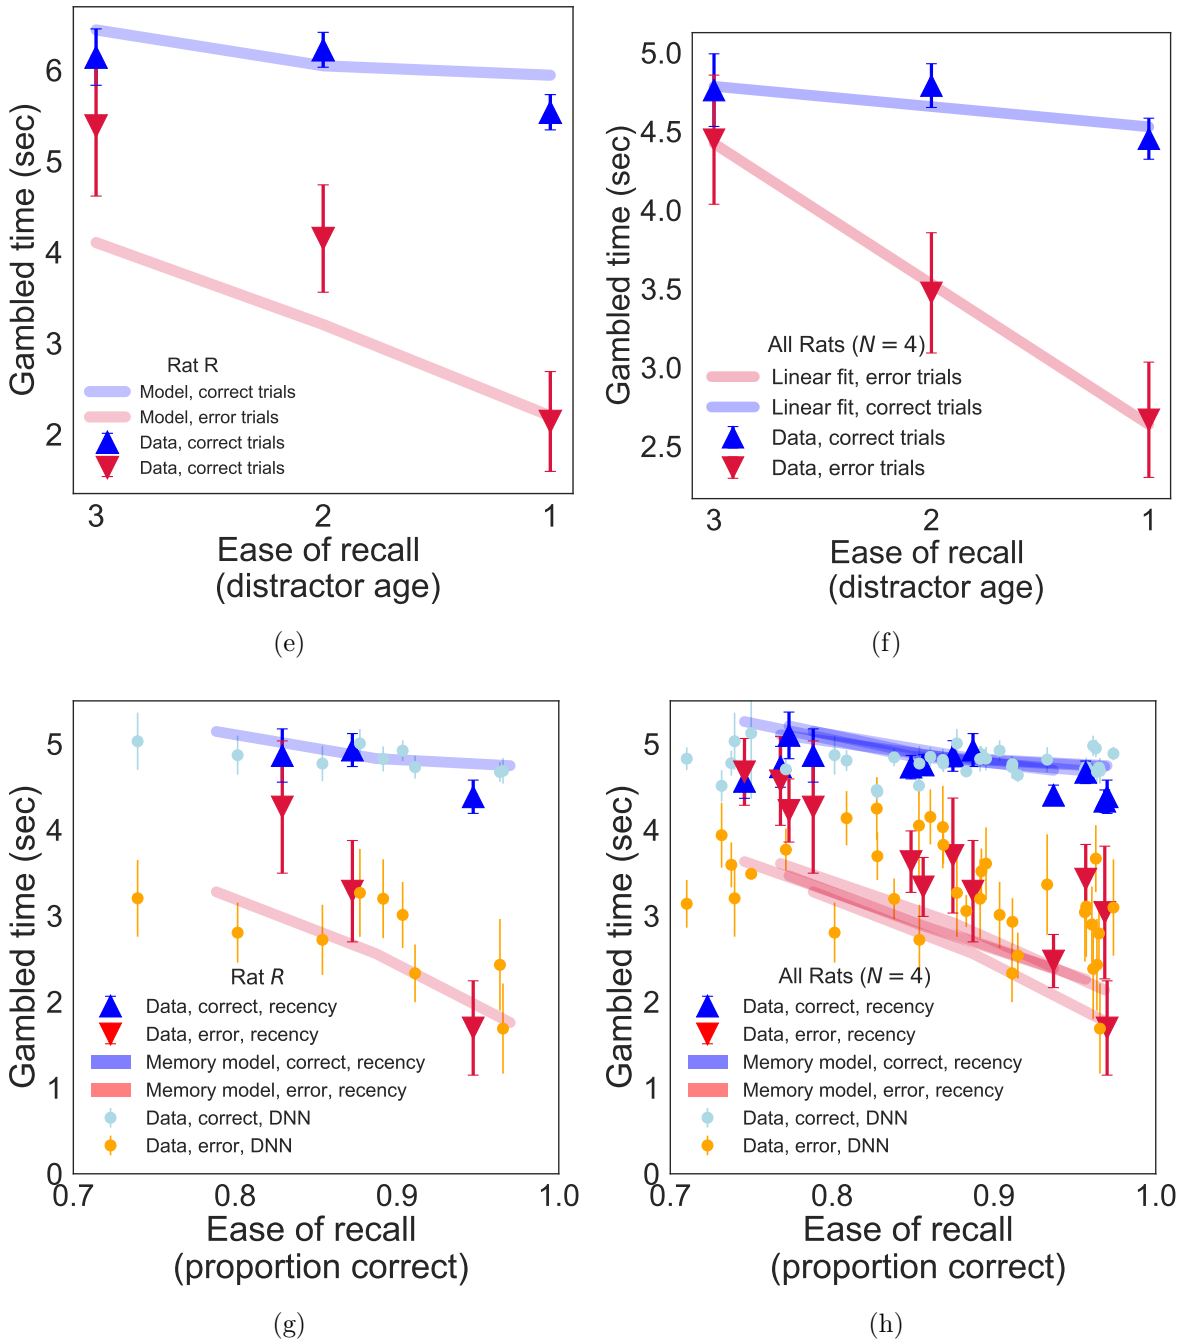

Figure 2.5: **Model-predicted gambled times match data.** Model-predicted gambled times (lines) with data (points) overlaid. **A.** Model predictions (gray line) for gambled time as a function of choice accuracy, for representative rat R, excluding uncued errors and conflict trials (black points are mean ± Y-error in 20 bins of X trials each). **E.** Model-predicted trends in gambled time as a function of distractor age match data. bf For representative rat R, long (upper 50 percent, dark gray points are mean ± SEM) and short (upper 50 percent, light gray points are mean ± SEM) gambled times are predicted by the model (dark and light gray lines, respectively). **C.** For representative rat R, correct (blue) and error (red) trials, trends in gambled times (points, mean ± SEM) are predicted by the model (lines). **D - F** Average trends across all rats are consistent with model. The average trends depicted in subfigures D-F match the model predictions shown in A-C, respectively. **G.** Comparison of gambled times on correct and error trials for SEM predictions (lines) versus data (triangles) versus DNN predictions (circles), representative rat R. **H.** As in G but with data from all four rats overlaid.

2.2.11 The probabilistic memory confidence model predicts confidence as a function of a synthetic difficulty axis

An advantage of the distractor age as a memory difficulty axis is its interpretability; it is intuitive that memory accuracy should be lower for episodes that occurred longer ago. However, memorability in general and on this task depend on many features. To take this into account, we generated a synthetic difficulty axis by again training a DNN to predict choice accuracy on the exhaustive feature set (Table 2.3) plus target and distractor branch ages. The network gave as output a probability that the rat would make a correct choice. Previously, we assessed the network accuracy in predicting trial outcome by interpreted probabilities above 0.5 as correct and below 0.5 as errors. Here, we interpret the probability of a correct choice as a difficulty score for each trial. We then defined trial types by binning based on these synthetic difficulty scores, and calculated rat performance accuracy in each. This yielded a synthetic difficulty axis, inferred from performance accuracy, based on the features in the exhaustive set. Calculation of performance accuracy for the trial types identified by this method allowed direct comparison to confidence predictions of our generative model and rat performance accuracy characterized by distractor age trial type. The DNN predicted the same trends observed for the distractor age memory difficulty axis for representative rat R (Figure 2.5) and for all four rats.

2.2.12 The statistical episodic memory confidence model predicts confidence as a function of episode temporal discriminability

Our model also generates predictions for confidence as a function of another determinant of choice accuracy, the target - distractor age. Treating this as another memory difficulty

axis, we can generate predictions for confidence on correct versus error trials as for the distractor age axis. Figure 2.6 shows three target - distractor differences for fixed distractor age 1, corresponding to trial types 1-2, 1-3, and 1-4. The model-predicted vevaiometric curve on this axis has a flat trend in gambled times for errors, and an increasing trend in gambled time for correct trials, as a function of target - distractor age. This is the result of the long-tailed memory distribution shapes, which result in a relatively greater shift for the positive density than the negative density of the difference distribution between a fixed distractor and increasingly older target distributions.

We found that predicted gambled times for correct and error trials on this memory difficulty axis also match trends in the confidence data for individual rats (representative rat R is shown in Figure 2.6) and for an average over all rats. Like for distractor age as a memory difficulty axis, gambled times predicted accuracy and were higher for correct than error trials (Figure 2.18). The same compression of gambled times seen for distractor age is also seen on this difficulty axis. The same qualitative patterns were observed for all rats (Figure 2.19).

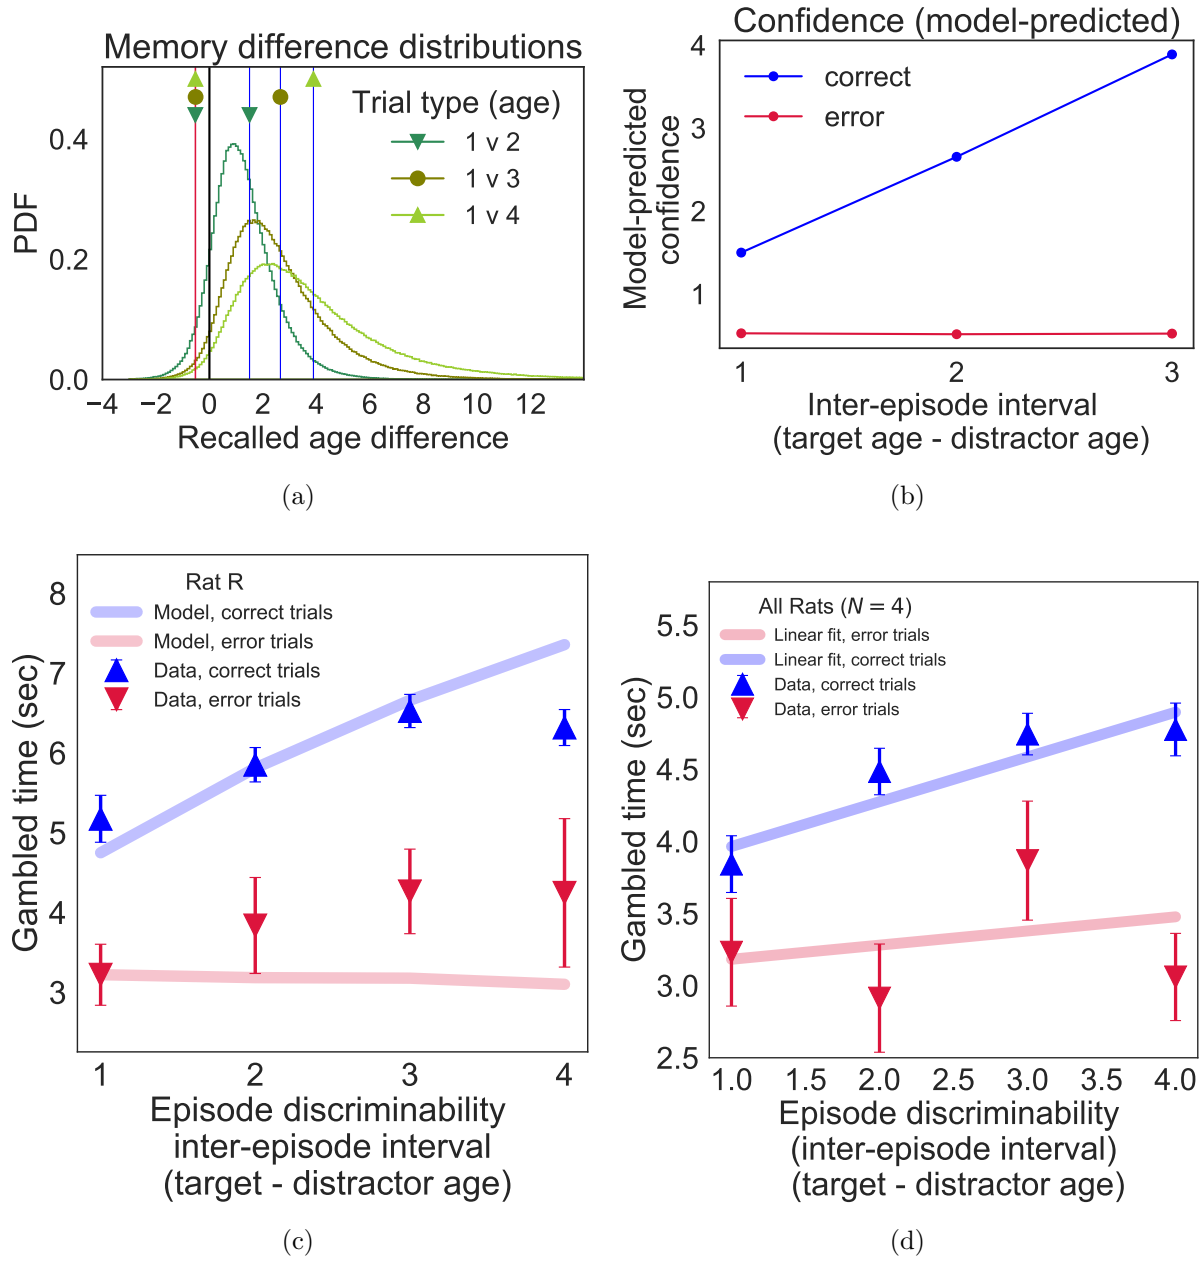

Figure 2.6: **A. Confidence distributions for distractor age = 1.** Considering trials that present distractor age 1 with target age 2, 3, and 4 (green, olive, lime distributions to the right), each target and distractor combination yields a distribution that represents the distribution of target - distractor retrievals (colors matched to target in top panel). The density to the right of zero (black horizontal line) represents those retrievals for which a correct choice will be made, and the average confidence on these trials is the center of this density (blue vertical lines). The density to the left of zero represents those retrievals for which an error choice will be made, and the average confidence on these trials is the absolute value of the center of this density (red vertical lines). **B.** Episodic memory model prediction. As the inter-episode interval (TARGET AGE - DISTRACTOR AGE) increases, the average confidence for correct trials increases, but the average confidence for error trials is virtually constant. **C.** Comparison of gambled times on correct and error trials for SEM predictions (lines) versus data (triangles) for correct (blue) and error (red) trials for representative rat R. **D.** As in G but with data from all four rats overlaid.

2.3 Discussion

Episodes occur at a temporal and a spatial coordinate, and when they are stored and accessible in memory by these coordinates as episodic memories, they are the most detailed, least abstract form of stored experience. For these memories, there is uncertainty on the temporal and spatial axes representing when and where the episode occurred, as well as the third, 'what' axis, which represents all other features. Uncertainty arises from perception itself and the ongoing memory processes underlying storage and use. We found evidence that rats have metacognitive knowledge of these uncertainties in memory and use this information to maximize reward. Specifically, rats can recall which of two locations they visited first in an ongoing sequence and use confidence in these memories to guide the amount of time they gamble on their choices. Time gambling as a confidence report requires both metacognitive monitoring of memory and demonstrates metacognitive control of future action[122]. The task design precludes associative learning of specific ports and their probability of being correct, and we found that latency to choice could not explain the pattern of metacognitive gambling[133].

In studies of confidence in a perceptual discrimination context, a single decision variable with Gaussian noise can often explain choice accuracy and confidence well. A longstanding challenge in studying metacognitive memory is in identifying the appropriate decision variable and noise axis[134–136]. Previous memory confidence tasks in human and non-human primates have used ROC analysis and an implicit Gaussian noise model[134]. Here, we applied to our task a modern machine learning approach to define a synthetic memory difficulty axis based on an exhaustive feature set. We used a DNN to predict accuracy as a function of task features, and treated the predicted accuracy as a synthetic memory difficulty score. This allowed us to treat choice accuracy, predicted by the DNN as a function of many features, as a proxy decision variable.

To better understand the determinants of memory confidence, we developed a novel generative model of episodic memory in which past episodes are modeled as memory distributions on a mental timeline, representing the animal's belief in when the episode occurred. We made two observations: first, that memory should decrease in fidelity with time; second, that an event that occurred in the past can never be thought to have occurred in the future. To satisfy these requirements, we modeled each memory distribution as a lognormal distribution with time-dependent variance. For each rat, we fit the model parameters based on choice accuracy, then generated predicted confidence values. Model-predicted gambled times as a function of the DNN-synthetic decision variable were in agreement with the data. We next generated model predictions as a function of the single feature that was most predictive of choice accuracy, the elapsed time since the distractor episode. This has been found in previous studies to determine difficulty[22]. Like the synthetic decision variable, model-predicted gambled times as a function of this feature were in agreement with the data. Finally, our model also predicted the observed patterns of choice accuracy and confidence on a third difficulty axis, the number of trials elapsed between the target and distractor episodes.

Critically, model predictions differed from predictions of the standard SDT model of confidence that employs Gaussian noise and a fixed variance[40,77]. Such a model could not explain our data. In particular, the trend in gambled time for correct versus error trials a function of memory difficulty (the vevaiometric curve) differs starkly from predictions of the Gaussian noise model. It has been observed previously that deviations from this signature can be seen under different model forms, for instance, a second-order computation where decision and confidence variables are not identical[137], varying stimulus strengths[138], or different confidence heuristics[137,139,140]. No model variant with Gaussian noise can produce the trends we observed in memory confidence, which require an asymmetric noise profile and unequal variance. Furthermore, these unique trends are not an effect of our specific choice of lognormal distributions, but arises as a direct consequence

of the two basic assumptions of the model. To predict gambled times from modeled confidence, we applied a transformation to the CDF of predicted confidence values such that its shape matched that of the empirical gambled times. This calibration step was applied over the empirical distribution of all gambled times for each rat, and therefore cannot account for the observed difference in gambled times between correct and error trials. Moreover, such a transformation in combination with a Gaussian noise model could not account for the observed trends in gambled time as a function of memory difficulty.

A trend in the data that is not captured well by the model is the saturating rise for gambled times as a function of target - distractor age (Figure 2.6), the second identified memory difficulty variable. One explanation for the leveling off of gambled times may be that rats have a temporal discounting bias for shorter/sooner rewards relative to longer/later, which will bias them toward shorter gambled times. Although we designed our reward function with this in mind, it may not be fully compensated by our reward function. This effect is likely to be strongest on trials for which animals are more confident and therefore have the highest reward expectation (these tasks can be understood as self-control tasks; they must exert the most on this trial type). This would result in shorter wait times on easy correct trials, flattening the curve of gambled times as a function of difficulty. This effect would be seen in real data but not the computational model predictions, which are calibrated to empirical gambled times overall but not selectively as a function of memory difficulty.

Previous studies of memory confidence have focused on visual or other recognition memory tasks[36,37], which present a series of stimuli then, after a delay, require the subject to identify them among never before seen distractors. In human subjects, recognition memory accuracy and confidence have been modeled by 'dual process' models of two parallel, distinct memory processes, recollection and familiarity[141,142] Recollection is remembering the full episodic details, based on an episodic memory system; familiarity is merely knowing that the event occurred, without its details, based on the semantic memory

system. A signal detection theory approach has been applied to each of familiarity and recollection[143] but differ fundamentally in from our approach: targets and lures share a distribution rather than each having one, with Gaussian noise, and on a memory strength axis[134,144]. In our task, such a familiarity computation alone cannot account for high accuracy in identifying the specific temporal ordering of port visits, both of which have been previously visited, often close in time. Rather, our task requires memory confidence in recent actions and their outcomes in spatial and temporal context. High choice accuracy indicates that rats can access such information. Gambled times consistent with the model imply animal metacognitive recall and autonoetic or self-reflective consciousness, previously reported only in humans. Our axis of time of occurrence is correlated with memory strength, but is not equivalent to it.

The model represents memory distributions at the present moment for episodes that happened n trials ago. Memory dynamics are implicit in the model; each memory distribution could alternatively be interpreted as corresponding to a different time. The model does not, however, address the evolution of confidence within a trial; it is a static SDT confidence model compatible with theories of confidence dynamics developed for perceptual discrimination, such as a drift diffusion model or other diffusion to bound variants[30].

Although we use a temporal axis, our model could be adapted to episodic memory on a spatial or other axis, or multiple axes. That the predictions of the model, which takes only timing information, are consistent with the data as a function of a synthetic difficulty axis that includes non-memory features, such as the last reward amount at each of the choice ports, suggests that a model on the temporal axis captures much of the variability. An advantage of parametrizing our model by a temporal axis is that this axis is correlated with memory strength and ongoing memory creation and use processes, allowing interpretation of our parameters as a form of hypothesis generation. Under the model, each

memory is represented at the time of encoding (t = 0) as a delta function. While this could be interpreted as representing perfectly accurate memory encoding, that is likely unrealistic; another interpretation is that there is merely uniform encoding noise across episodes, with no lapses in attention. The accumulating noise on older episode memories corresponds to memory processes subsequent to encoding. Every time a memory is retrieved for consolidation or use it is known to become labile again. That episodes are increasingly difficult to remember the longer ago they happened is explained by ongoing retrieval processes and is reflected in our model as the time-dependent increase in variance of each memory distribution.

The evidence we find that rats are able to compute confidence in episodic memories suggests they may have autonoetic consciousness. A theorized aspect of autonoetic consciousness is the ability also to imagine or plan the future . There is evidence of overlap between episodic memory and an ability to predict the future[145,146], raising the possibility that rats may also compute confidence in their future plans or projections.

We note that the data in this study were taken from rats implanted with multi-site recording devices, that the reported number of trials and the time gambling behavior can be performed by rats with neural implants, opening the door to a mechanistic understanding of the underlying neural processes[147]. In addition to its philosophical implications for the self and consciousness, memory confidence has significant legal, social, and political importance, for instance in judging the credibility of eyewitness testimony (*e.g.*, in the 2018 Kavanaugh hearings)[148].

Aberrations in confidence computations have been proposed to account for a variety of psychiatric symptoms[149]. These proposals have built conceptually on models of perception[150]. For instance, hallucinations and delusions may result from overly strong priors as described in a predictive coding framework[151]. Aberrations in memory confidence, particularly in episodic memories, have been proposed to account for symptoms of

OCD[13,152], schizophrenia[15,153], and others[18]. Relative to perceptual confidence, however, memory confidence has not had the requisite behavioral tasks, animal models, or theoretical framework for understanding its contribution to disease. A statistical treatment of memory confidence may enable study of the negative symptoms of schizophrenia and other conditions thought to be related to memory. As these negative symptoms are often more debilitating than positive symptoms like hallucinations, it is our hope that this line of investigation will eventually be of therapeutic value to patients and their families.

2.4 Methods

2.4.1 Behavioral training and task

All procedures followed the guidelines from the University of California San Francisco Institutional Animal Care and Use Committee and US National Institutes of Health. Male Long-Evans hooded rats were trained to perform an episodic memory task with time gambling for liquid reward. Behavioral testing was controlled by custom software written in Python using data acquisition hardware (Trodes ECU, SpikeGadgets LLC) to record rat pokes and unpokes at the ports and to control reward delivery.

Three cohorts of Long Evans male rats (3-4 months old, 450-600g, 6-8 rats per cohort) were habituated to daily handling for a week and to hand-delivered liquid food reward (evaporated milk plus 5 percent sucrose) from a syringe in the home cage for three days. Animals were then food deprived to 85-90 percent of their baseline weight and pre-trained on a raised linear track for 3-4 days, 2-3 epochs/day, 10 mins/epoch (Stage I). A port was located at each end of the track, each equipped with an LED light and an IR beam, which detects entry and exit from the port[154]. Each port can automatically deliver reward, which flows for a specified length of time through the port at a rate of .17mL/sec; it must be

consumed during this period, as the reward does not remain in the port after this period. A variable delay between nose-poke and reward delivery drawn from exponential distribution with $\tau = 0.2 - 0.5$ seconds that was gradually incremented to $\tau = 1 - 8$ seconds. Rats learned to run back and forth on the track to visit the lit well, and to wait for the delayed reward From each cohort, 2-3 rats with the best performance were selected for training on the episodic confidence task.

Following linear track pre-training in Stage I, rats were trained on the full track where the EC task took place. The track had eight ports in total: one home port at the center, one back port, six choice ports at each end of six branches. Each epoch was of a fixed length per animal, during which trials were self-paced. In Stage II, rats learned the basic task structure (Figure 2.1) with only one cue lit per trial. The lit cue corresponded to the target selected by the same code as in the final task logic; lighting of the distractor port was suppressed. The sequence of visits within a trial was: home port light on; rat pokes at home port for a small fixed reward (350 ms long); home port light off; after a variable `cue delay`, one choice port lights; rat pokes lit choice port; choice port light off and port delivers `initial reward` 350 ms long and, after a variable `reward delay`, a `wait-dependent reward`; back port light on; rat pokes back port; back port light off and port delivers `back reward`. Choice accuracy was measured as the percentage of trials for which the rat visited the lit choice well.

The `cue delay` was introduced to jitter the events of each trial relative to every other trial, to control for across-trial temporal correlations between behavioral and neural events. To train rats to wait for the cue lights to come on, the `cue delay` was gradually increased from range [0.2, 0.5] to [0.5, 2]. Initially, the `back reward` delivered the same reward amount as the `wait-dependent reward` regardless of trial outcome, which encouraged the animals to solidify knowledge of the port visit sequence (*i.e.*, to not skip the back port). After three epochs, `back reward` was only delivered on correct trials.

The `reward delay` was determined by sampling from an exponential distribution with rate parameter $\lambda = 1/2$, accepting only samples that were between 1-3s at the start of this training phase and 2-10s by the end, with a `wait-dependent reward` amount that increased accordingly, to allow rats to learn that a longer period spent nosepoked in the port would result in a larger reward. After rats were consistently performing at above 80 percent choice accuracy and waiting for the full `reward delay`, the `initial reward` was omitted. Once rats were able to wait for the majority of the `reward delay`s (6-10s), the switch was made to gambling logic. In the gambling logic of the final task, rats voluntarily reported the time they were willing to wait for the reward. The gambled time began at the time of nosepoke in the choice port and ended when rats withdrew from the port. Nosepoke withdrawal was detected with a 'grace period' (800 ms for rats T, S, D; 700 ms for rat R in final behavior, calibrated based on how quickly each rat moved) to allow for small head movements during the gambling period: rats were only declared to have ended the gambling period after a grace period had passed between un-poke and the next re-poke. After gambled times were observed to be stable across at least three epochs, the distractor cue was introduced alongside the target cue, starting with distractor age 1. Distractors age 2 and 3 were introduced when choice accuracy was approximately 80 percent and stable.

Approximately 3000 - 4000 trials were collected from each of four rats. Each rat had a typical length of time for which he would continuously play the game, after which he would occasionally perform trials but otherwise sleep or lean off the edge of the track and attempt to eat the milk tubes or CAT6 cables. Epochs shorter than 20 minutes (rat T, N = 5 excluded epochs), 40 minutes (rats S, N = 0 excluded epochs, and D, 2 excluded epochs) or 45 minutes (rat R, N = 2 excluded epochs) were excluded from final analyses. This resulted in the following epoch and trial counts: From rat T, 2978 epochs over 42 trials; from rat S, 4111 trials over 40 epochs; from rat R, 3660 trials over 49 epochs; from rat D, 4369 trials over 61 epochs. Typically rats ran an average of 350-400 meters per day (the human equivalent of approximately five miles) and consumed 50 mL of sweetened

evaporated milk.

The selection of distractor and target is random with temporal weighting. The logic used for selection of the two cued ports on each trial is: from the list of possible pairs [AB, BC, CD, DE, EF, FA], select candidate pairs for which at least one of the ports has an allowable distractor age (1, 2, or 3). If there is more than one candidate pair in this list, remove from it the candidate pairs with distractor ages equal to those presented on the last trial, the penultimate trial, and the trial before that, in that order, until candidate pairs with only one distractor age remain. If there is only one candidate pair in this set, select it as the presented pair. If there is more than one candidate pair in this set, randomly select between them with equal probability. For example, if the last three trials were distractor ages 1, 2, 3 (*N.B.: regardless of which ports these distractor ages corresponded to*), then on the upcoming trial, the candidate pair with distractor age 3 would be removed first, then the candidate pair with distractor age 2. The candidate pair with distractor age 1 would be selected; if there were more than one candidate pair with distractor age 1 remaining, the cued pair would be selected randomly from this set. On every trial, there will necessarily be a candidate pair with distractor age 1. There will not, however, be candidate pairs with distractor ages 2 and 3 on every trial; this can occur in the case of revisits, where the port with distractor age 3 is the same as the port with distractor age 1 (or the age 2 port = the age 1 port, or the age 3 port = the age 2 port = the age 1 port). For this reason, the sequence of presented distractor ages is similar to a repeating cycle 1-2-3, but not exactly. This selection algorithm has the effect of sampling evenly across distractor types and preventing an alternation sequence from developing.

The reward function was designed to counter the potential effects of temporal discounting. On correct trials, for investments less than 2.2 seconds, the length of time for which a sweetened evaporated milk reward was delivered at a constant rate of 0.17 mL/sec was given by $R = 0.27e^{0.34(t+0.8)}$; for investments greater than 2.2 seconds,

$R = 2.6 \times log(0.44 \times (t + 0.8))$. This function delivered a minimal reward (350 ms long, or about 5 mL) for any correct trial, even if the gambled time was zero. Four seconds of reward are delivered for a ten-second wait. The reward function has a relatively low, increasing derivative for gambled times less than or equal to 2.2 seconds and a relatively high, decreasing derivative for gambled times greater than 2.2 seconds. The desired effect was to bias the rat toward longer gambled times on trials for which he would already have waited at least 2.2 seconds, as he could double the reward amount by waiting just one second longer. If rats were able to access memory confidence, these longer waits should be enriched in correct trials and the reward function could help better resolve them from error trials. By selecting a function with a derivative that fell to approximately the value it had in the [0, 2.2] second range at 9.5 seconds, rats were still encouraged to run enough trials to sample evenly over the trial types within each epoch: rats take an average of 15 seconds to perform the rest of a trial, with a 9.5 second gambled time, 4-second reward delivered at both choice and back ports, this yields approximately 30-second trials, for our aim of 80 trials per 40-minute epoch.

Rats that performed many trials per epoch with a large spread in gambled times were implanted with hardware for recording neural data. Following a week or more of recovery, behavioral data in the final task were acquired from implanted rats.

2.4.2 Correlation of choice latency and gambled time

For analysis of correlation between gambled times and latency to choice, outliers with gambled times greater than 10 seconds or latency to choice greater than 20 seconds were excluded, leaving over ninety percent of the data per rat. Linear regression was implemented in `SciPy`.

2.4.3 Evaluation of alternative strategies

For each rat, the proportion of times that each port was presented as target versus distractor were compared. Per epoch, these values were rarely above or below 50 percent by greater than 3 percent, and the majority of differences were not statistically significant by a t-test for independent samples.

We tested whether there existed an alternative strategy that could better explain the rat's choices than the true rule, which is to select the least recently visited of the two cued ports. For every trial in every epoch, for each rat, we determined whether the alternative rule would have resulted in the same choice as the one the rat made, or the same choice dictated by the true rule. This resulted in two proportions per epoch for each rat.

2.4.4 Evaluation of logistic regression and neural network models of choice accuracy

We used logistic regression and a DNN model to predict choice outcome (*correct* or *error*) as a function of an exhaustive or a reduced feature set. The features were each standardized to have zero mean and unit variance.

The DNNs were implemented in KERAS[155] using the TENSORFLOW[156] backend and optimized using ADAM[157]. Each network has three hidden layers with 32 nodes each and the rectified linear unit activation. The output of the last layer is a sigmoid and the binary cross-entropy is the loss function. Networks were trained with 200 epochs with early stopping using a patience of 5 epochs. A $k = 5$-fold training procedure is used whereby $1/k^{\text{th}}$ of the data are withheld for testing, $1/k^{\text{th}}$ are withheld for validation and the rest are used for training. The trials that comprise each fold are uniformly selected at random. A total of 10 networks are trained for this configuration and the network with the best

validation loss is used to evaluate on the test set. The test set is then rotated k times until all data are used for testing. The loss is weighted during training so that the weighted number of instances from the two trial outcomes (*i.e.*, correct or error) are the same.

Logistic regression was tested in multiple frameworks. In Keras, logistic regression is simply a neural network without any hidden layers. It was found that the performance of such a regression model was comparable to the default logistic regression in SCI-KIT LEARN[158] as well as STATSMODELS[159]. Unless otherwise indicated, the SCI-KIT LEARN tool was used for logistic regression studies.

2.4.5 Fitting the statistical episodic memory model parameters

The model was fit on a subset of distractor-target trial types for which there was enough data, excluding uncued errors and conflict trials. The reduced datasets were 1716, 2356, 2493, and 2092 trials for rats T, S, D, R, respectively. Model parameters a_0, a_1, a_2 and σ_0 were fit for each rat based on its performance across trial types defined by distractor and target - distractor ages (excluding uncued error trials and target - distractor ages > 4). To find the best fit parameters using Eq. 2.3, we use the Nelder-Mead method with 200 maximum iterations as implemented in `scipy`, minimizing the χ^2 fit to error rates across trial types. Then, using these parameters, we generated the distributions corresponding to each episode memory and sampled from each 100,000 times to generate target memories, distractor memories, the outcome of the trial (*correct/error*) and a confidence (absolute value of the difference between target and distractor).

To convert the simulated confidence values to invested times, we mapped the confidence (C) probability density onto the probability density of the rat's invested times (T). Let $F(x) = \Pr(C \leq x)$ be the cumulative distribution function (CDF) for C and $G(x) = \Pr(T < x)$ be the CDF of the invested times. Then, the mapping procedure

proceeds as follows:

1. Compute the empirical CDF of the confidence values from the model $\hat{F}$ using `ECDF` from `statsmodels`. Trials are generated from the model such that the number of trials from each trial type follows the relative rates in data which are not uniform. The minimum number of trials generated is 10^4.

2. Compute the empirical CDF of the wait times from data $\hat{G}$ using `ECDF` from `statsmodels`. This is inclusive over trial types.

3. For each confidence value c, evaluate $\hat{G}^{-1}(\hat{F}(c))$. The inverse $\hat{G}^{-1}$ is computed via linear interpolation (using `numpy`'s `interp` function) inverting the x and y coordinates.

2.5 Supplement

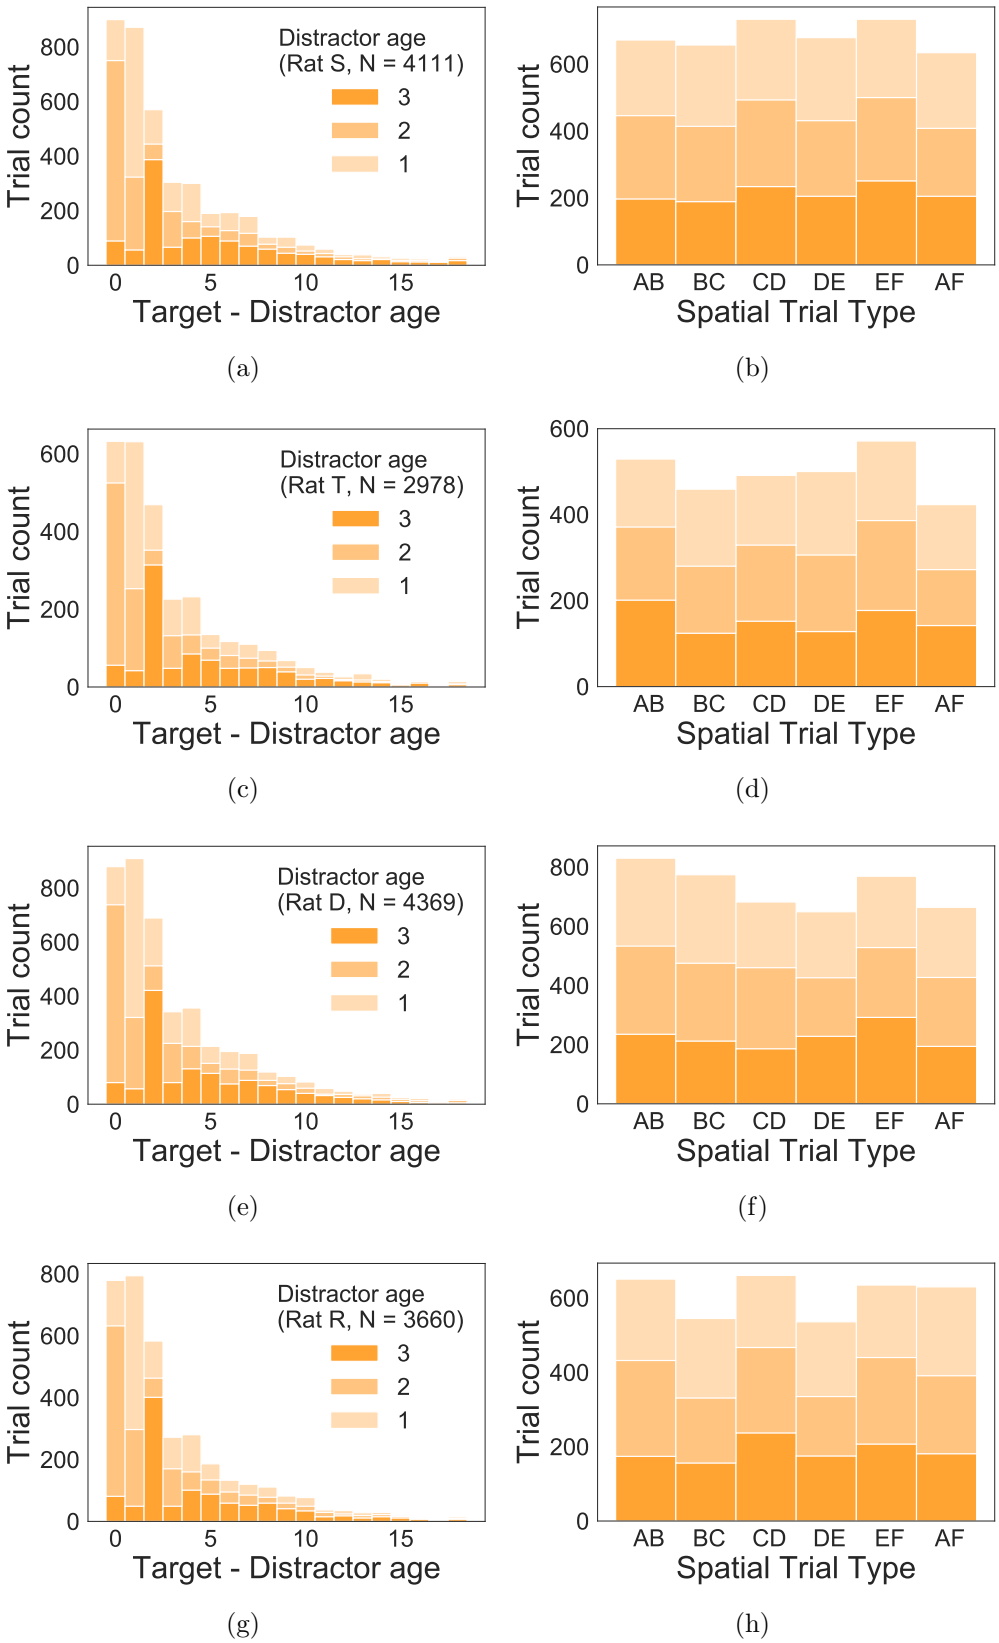

Figure 2.7: Even sampling across memory age and spatial trial type.

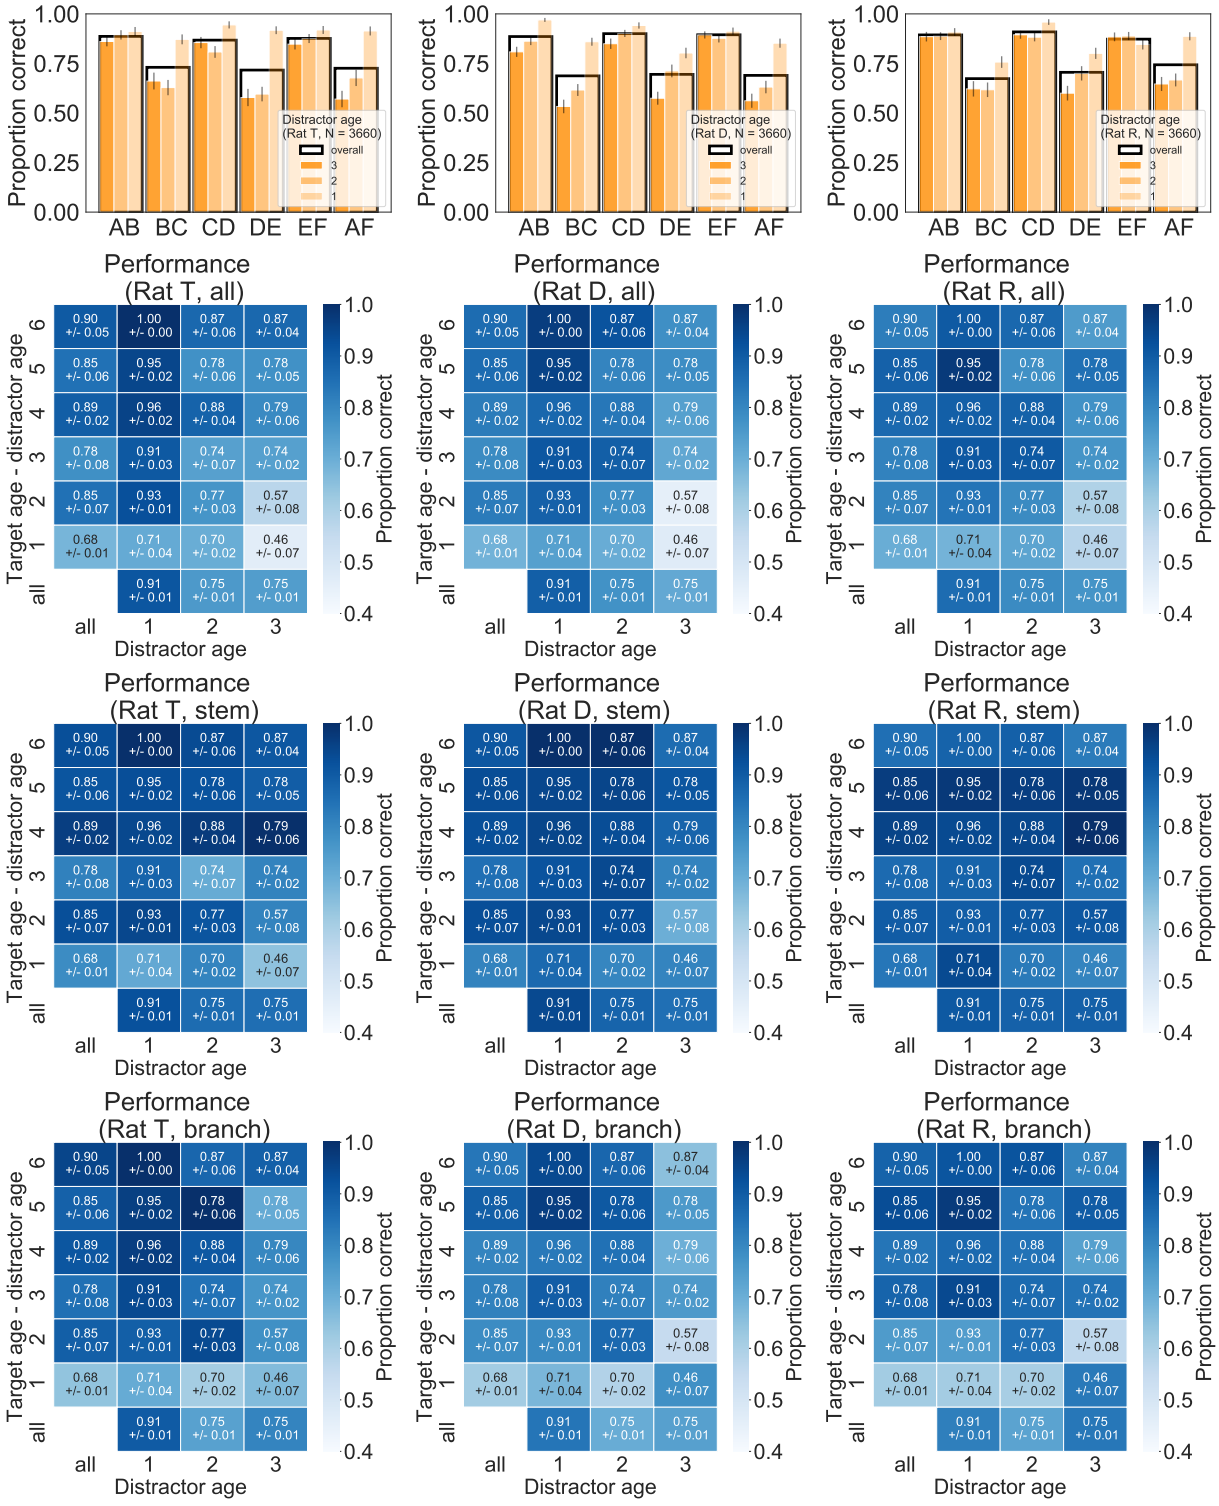

Figure 2.8: **A. Performance is higher for branch than stem trials.** Includes conflict trials.

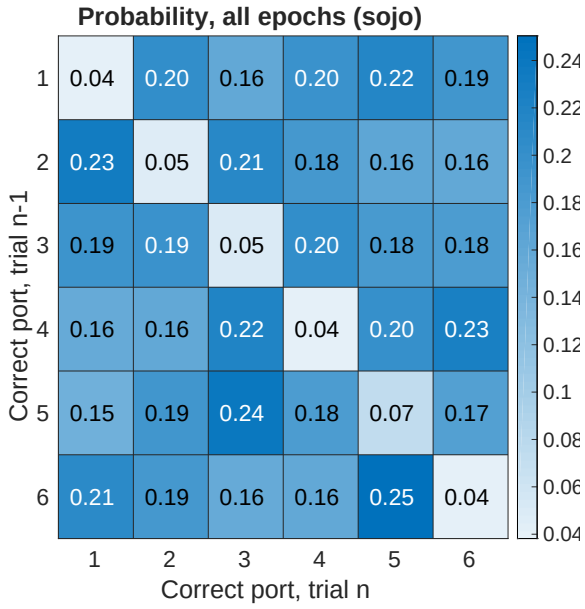
Probability, all epochs (sojo)
Correct port, trial n-1
Correct port, trial n
1 2 3 4 5 6
0.04 0.20 0.16 0.20 0.22 0.19
0.23 0.05 0.21 0.18 0.16 0.16
0.19 0.19 0.05 0.20 0.18 0.18
0.16 0.16 0.22 0.04 0.20 0.23
0.15 0.19 0.24 0.18 0.07 0.17
0.21 0.19 0.16 0.16 0.25 0.04

(a)

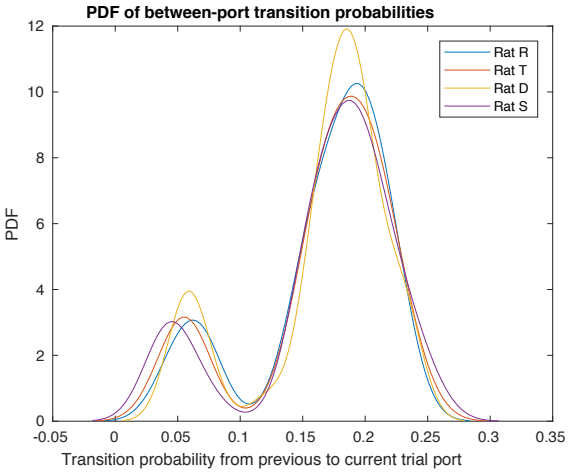
PDF of between-port transition probabilities
PDF
Transition probability from previous to current trial port
Rat R
Rat T
Rat D
Rat S

(b)

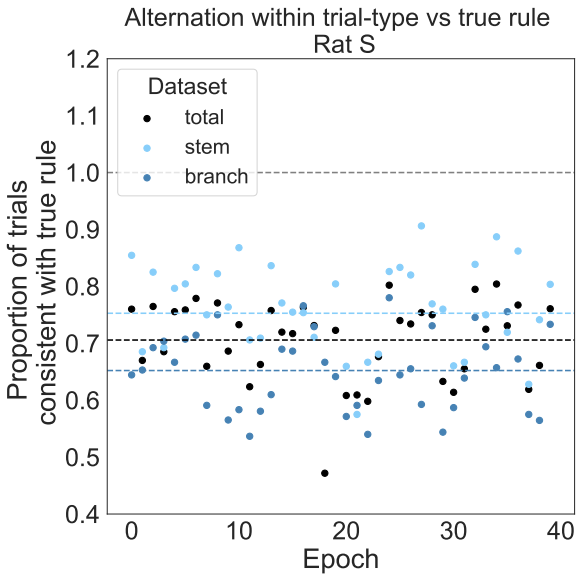
Alternation within trial-type vs true rule
Rat S
Proportion of trials consistent with true rule
Epoch
Dataset
total
stem
branch

(c)

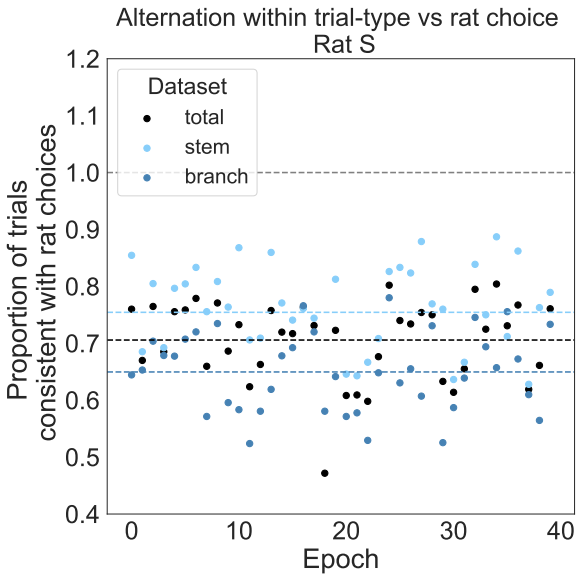
Alternation within trial-type vs rat choice
Rat S
Proportion of trials consistent with rat choices
Epoch
Dataset
total
stem
branch

(d)

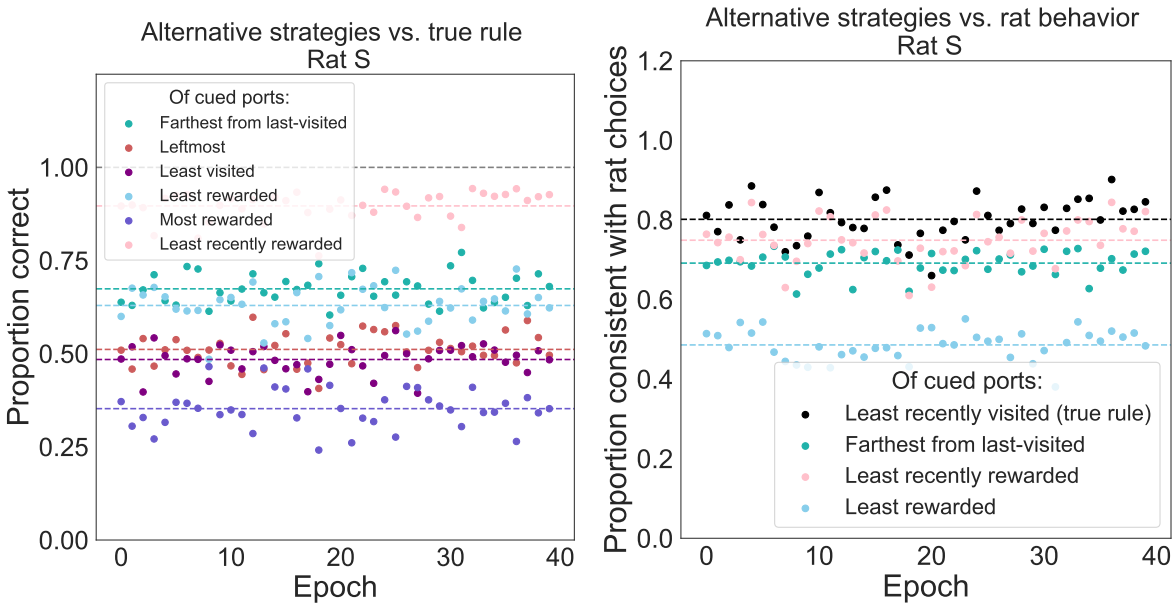

(e) (f)

Figure 2.8: **Alternative strategies cannot explain performance accuracy. A.** For all trials from representative rat S, the correct port on trial n-1 (y-axis) does not strongly predict the correct port on trial n (x-axis), as illustrated by the nearly equivalent transition probabilities across rows. **B.** PDF of transition probabilities in (A) for each of four rats shows a smooth distribution with two peaks: a low-probability peak corresponding to the same port being presented as target twice in a row, and a peak centered at 0.2, corresponding to the probability of transition from port N to any of the other five ports. **C.** There is no epoch (x-axis) for which an alternation rule matches the true rule on every trial (black points), stem trials only (light blue) or branch trials (dark blue). Dashed lines = averages across epochs. **D.** As in C, for agreement with rat choices rather than true rule. There is no epoch for which an alternation strategy can explain greater than 90 percent of the rat's choices. **E.** Alternative strategies are not consistent with the true rule. The proportion of trials that would be correct under application of each of six alternative strategies (y-axis) is shown for each epoch (x-axis) for representative rat S. Strategies that result in a correct outcome greater than .5 of the time are to select the least recently rewarded, the farthest from last-visited, or least rewarded overall of the two cued ports. Dashed lines = averages across epochs. **F.** Considering for representative rat S those strategies that were in agreement with the true rule more than .5 of the time (blue, green, and pink points), the proportion of trials consistent with the rat's choices (y-axis) for each epoch (x-axis). The true rule (black points) explains behavior better than the alternative strategies. Dashed lines = averages across epochs.

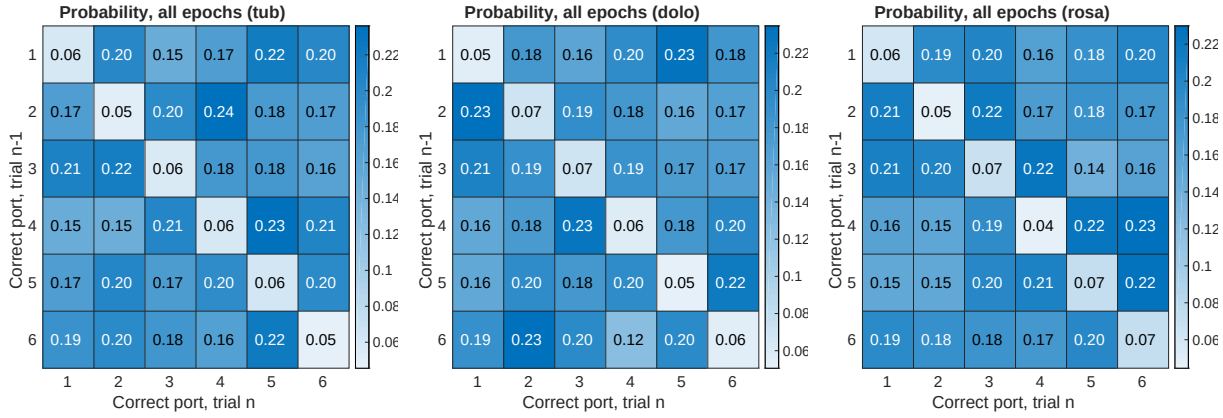
Probability, all epochs (tub)
Correct port, trial n-1
Correct port, trial n
1 0.06 0.20 0.15 0.17 0.22 0.20
2 0.17 0.05 0.20 0.24 0.18 0.17
3 0.21 0.22 0.06 0.18 0.18 0.16
4 0.15 0.15 0.21 0.06 0.23 0.21
5 0.17 0.20 0.17 0.20 0.06 0.20
6 0.19 0.20 0.18 0.16 0.22 0.05
Probability, all epochs (dolo)
Correct port, trial n-1
Correct port, trial n
1 0.05 0.18 0.16 0.20 0.23 0.18
2 0.23 0.07 0.19 0.18 0.16 0.17
3 0.21 0.19 0.07 0.19 0.17 0.17
4 0.16 0.18 0.23 0.06 0.18 0.20
5 0.16 0.20 0.18 0.20 0.05 0.22
6 0.19 0.23 0.20 0.12 0.20 0.06
Probability, all epochs (rosa)
Correct port, trial n-1
Correct port, trial n
1 0.06 0.19 0.20 0.16 0.18 0.20
2 0.21 0.05 0.22 0.17 0.18 0.17
3 0.21 0.20 0.07 0.22 0.14 0.16
4 0.16 0.15 0.19 0.04 0.22 0.23
5 0.15 0.15 0.20 0.21 0.07 0.22
6 0.19 0.18 0.18 0.17 0.20 0.07

(a) (b) (c)

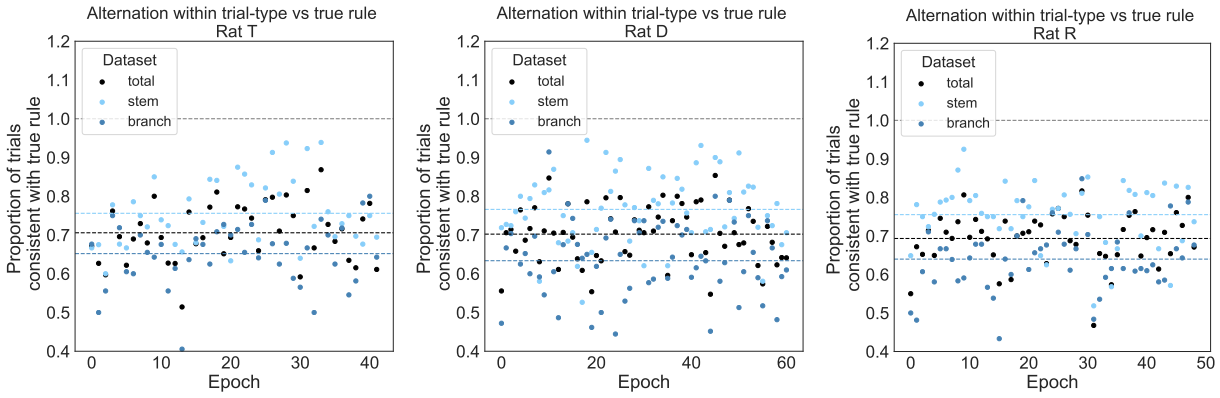
Alternation within trial-type vs true rule
Rat T
Rat D
Rat R
Dataset
total
stem
branch
Proportion of trials consistent with true rule
Epoch

(d) (e) (f)

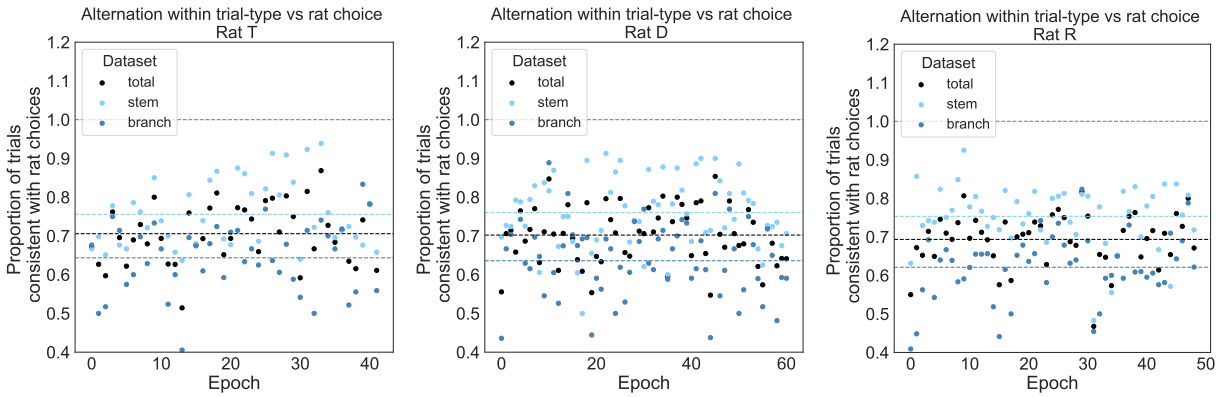
Alternation within trial-type vs rat choice
Rat T
Rat D
Rat R
Dataset
total
stem
branch
Proportion of trials consistent with rat choices
Epoch

(g) (h) (i)

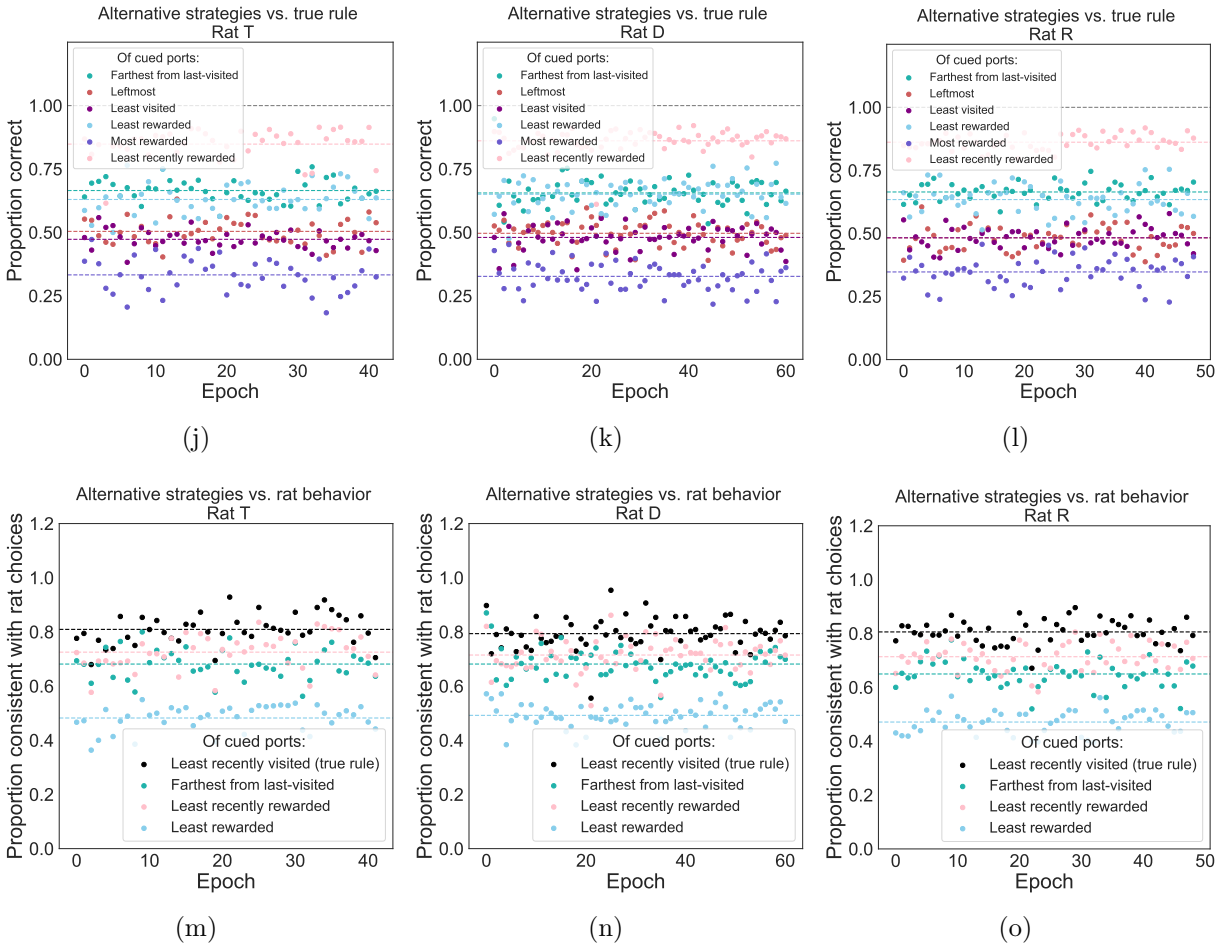

Figure 2.8: As in Figure 2.8, including data for all rats.

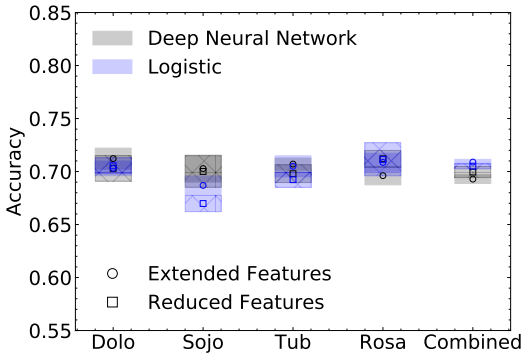

(a)

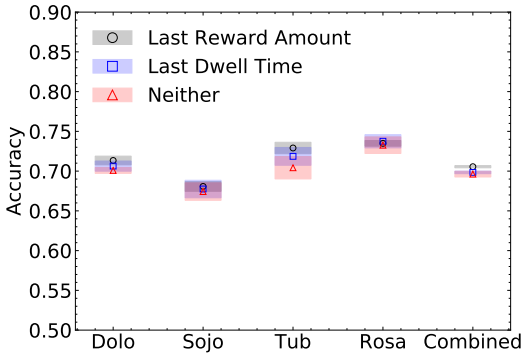

(b)

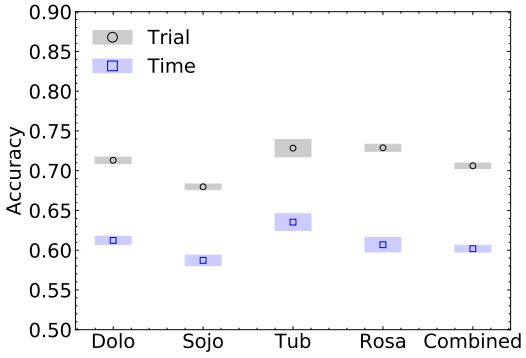

(c)

Figure 2.9: **A.** Comparison of DNN, logistic regression models for full and reduced feature sets by accuracy. Error bars = **B.** Comparison of a logistic regression model trained on feature sets including only spatial and temporal features or spatial and temporal features plus last reward amount or spatial and temporal features plus last dwell time **C.** Comparison of a logistic regression model trained on the base feature set including spatial and temporal features of time in terms of number of trials versus clock time.

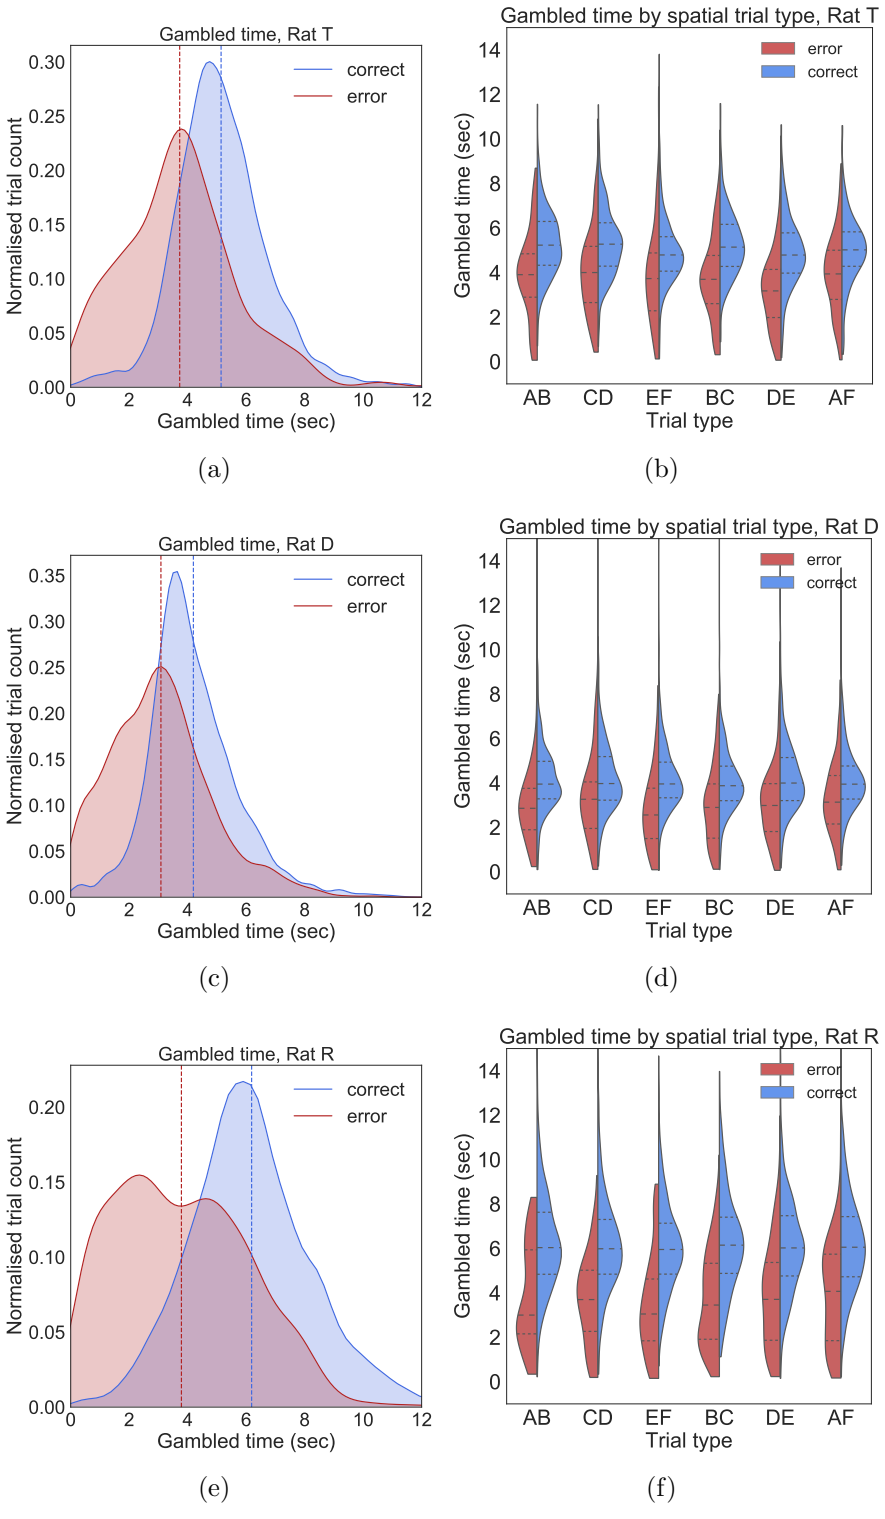

Figure 2.10: As for Figure 2.4. For rat T, D, R, N(correct trials), N(error trials) = 2400, 479; 3479, 772; 2952, 539.

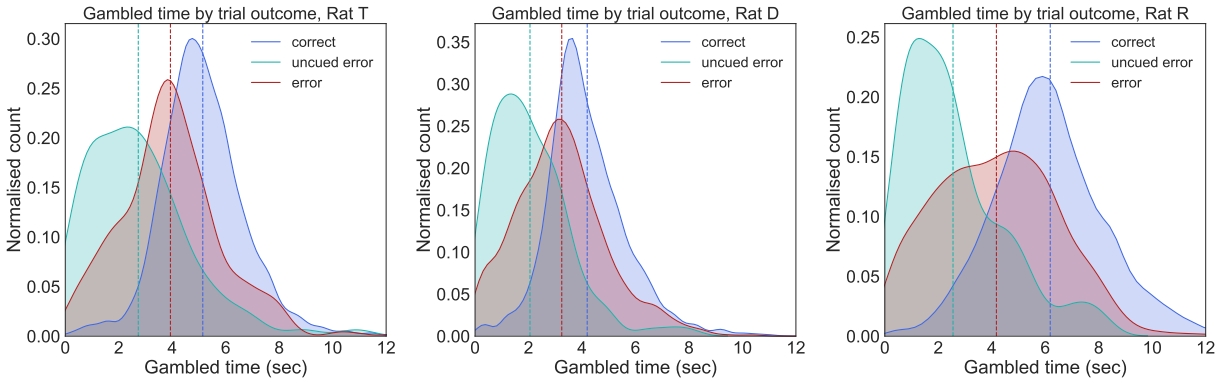

Figure 2.11: **Less time is gambled on uncued errors. A.** The majority of uncued errors occurred on trials with distractor age = 1, consistent across rats. **B.** Gambled times were significantly shorter for uncued errors relative to other errors. Excluding uncued errors, gambled times on error trials were still significantly lower than gambled times on correct trials.

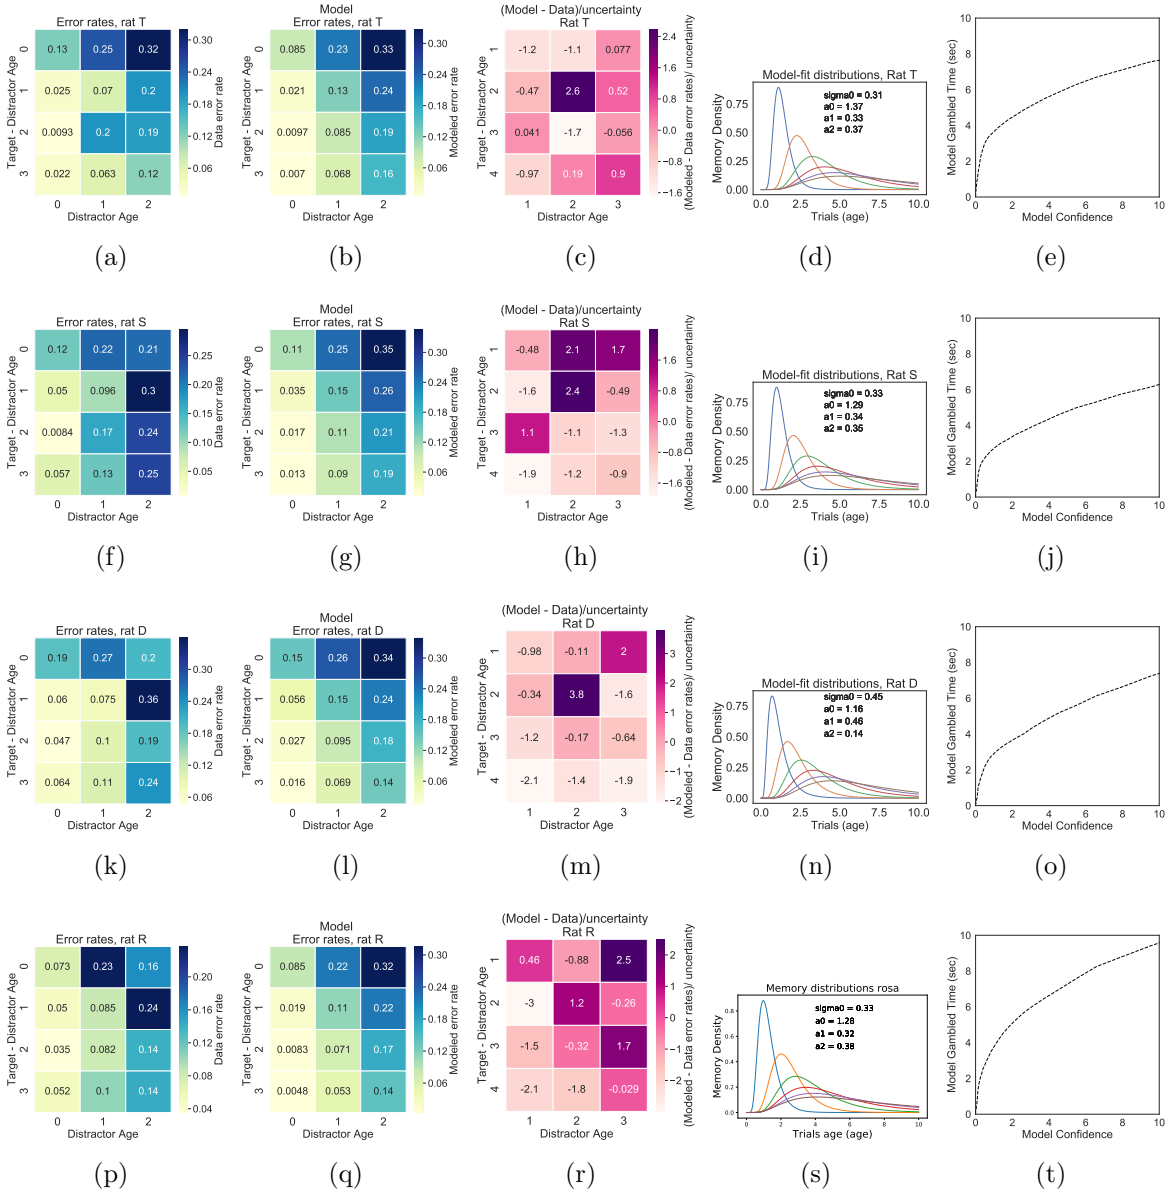

Figure 2.12: **A., F., K., P.** The measured error rate ϵ_{data} in bins of Distractor Age and Target-Distractor Age. **B., G., L., Q** fitted model predictions ϵ_{model}. **C., H., M., R.** the relative difference between the model and the fit: $(\epsilon_{\text{model}} - \epsilon_{\text{data}})/\epsilon_{\sigma_{\text{data}}}$. **D., I., N., S.** Fitted model lognormal distributions corresponding to the memory distributions (M_i') from which the rat samples in memory retrieval. Memory distributions (M_i') from which the rat samples in memory retrieval. Each distribution is a different color and corresponds to the memory distribution, from left to right, for the last trial to six trials ago. **E., J., O., T.** Mapping function from confidence to gambled times. The function $\hat{G}^{-1}(\hat{F}(c))$ that converts model confidence into model invested time for rats *T, S, D, R.*

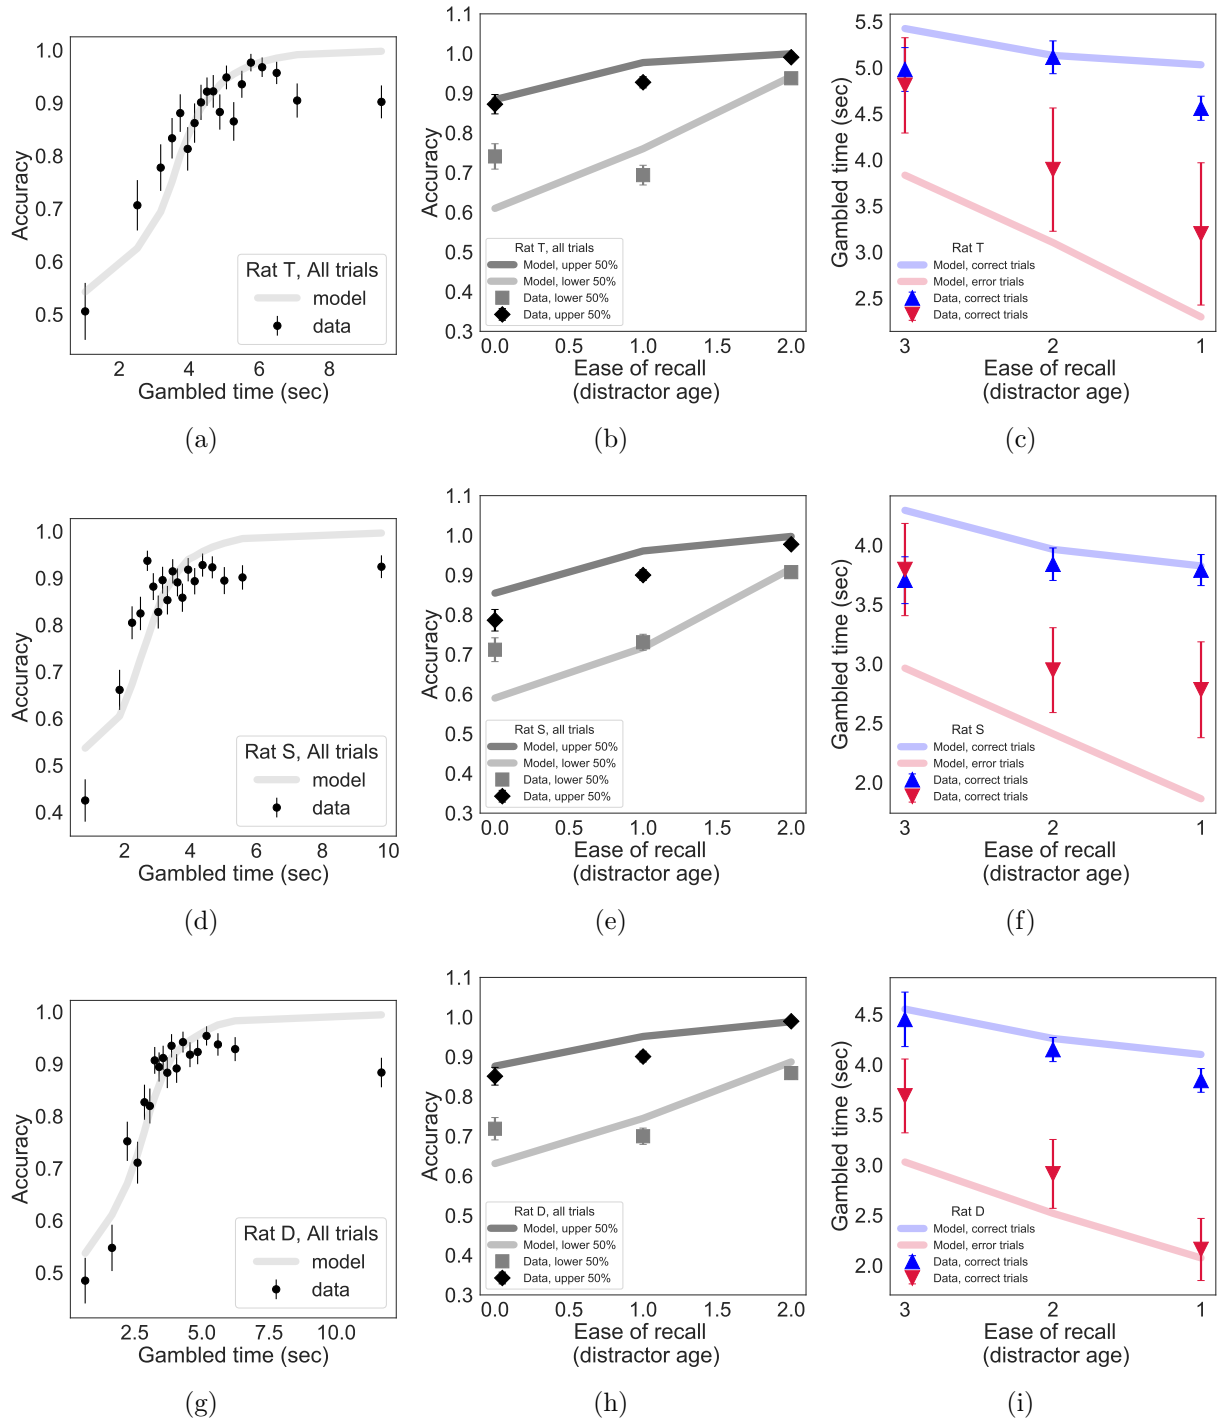

Figure 2.13: **Data from all rats for model predictions on memory difficulty axis based on distractor age.** Model-predicted gambled times (lines) with data (points) overlaid, for each rat (rows). **A., D., G. Model predictions for gambled time as a function of choice accuracy.** Excluding uncued errors and conflict trials (black points are mean ± Y-error; N=X trials), gambled time predicts choice accuracy as predicted by the model (gray line). **B., E., H. Model-predicted trends in gambled time as a function of distractor age match data.** Long (upper 50 percent, dark gray points are mean ± SEM) and short (upper 50 percent, light gray points are mean ± SEM) gambled times are predicted by the model (dark and light gray lines, respectively). **C., F., I. Trends in gambled time are predicted by the model.** Correct (blue) and error (red) trials, trends in gambled times (points, mean ± SEM) are predicted by the model (lines).

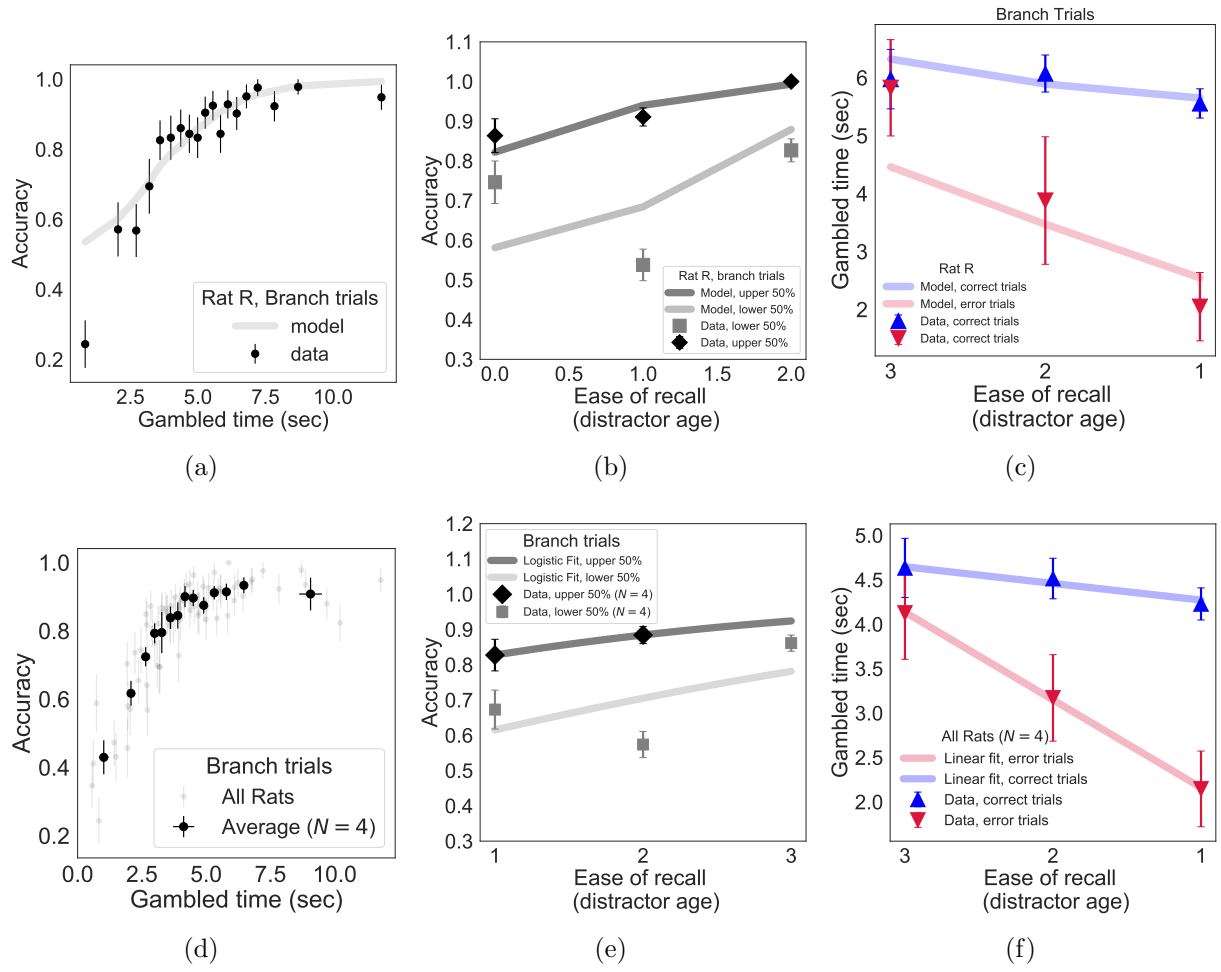

Figure 2.14: **Model predictions for distractor age, branch trials only.** Model-predicted gambled times (lines) with data (points) overlaid. **A. Model predictions for gambled time as a function of choice accuracy.** For representative rat R, excluding uncued errors and conflict trials (black points are mean ± Y-error; N=X trials), gambled time predicts choice accuracy as predicted by the model (gray line). **B. Model-predicted trends in gambled time as a function of distractor age match data.** For representative rat R, long (upper 50 percent, dark gray points are mean ± SEM) and short (upper 50 percent, light gray points are mean ± SEM) gambled times are predicted by the model (dark and light gray lines, respectively). **C. Trends in gambled time are predicted by the model.** For representative rat R, correct (blue) and error (red) trials, trends in gambled times (points, mean ± SEM) are predicted by the model (lines). **D - F Average trends across all rats are consistent with model.** The average trends depicted in subfigures D-F match the model predictions shown in A-C, respectively.

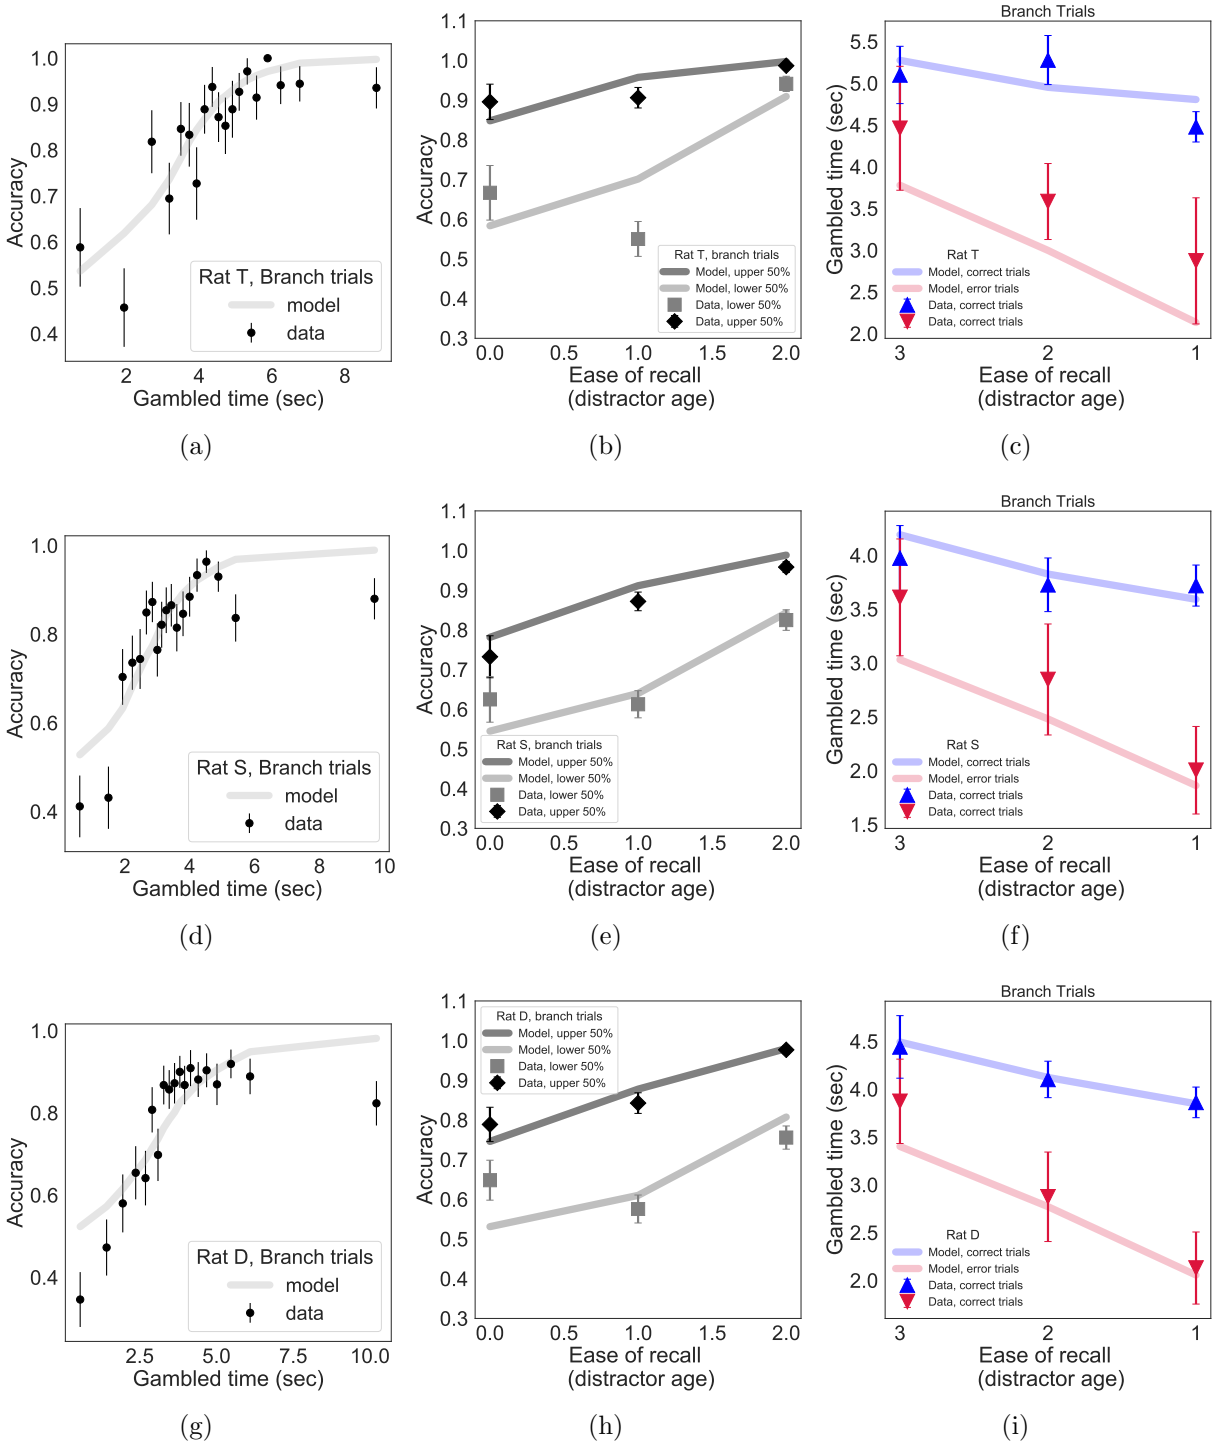

Figure 2.15: As for 2.14, data for all rats.

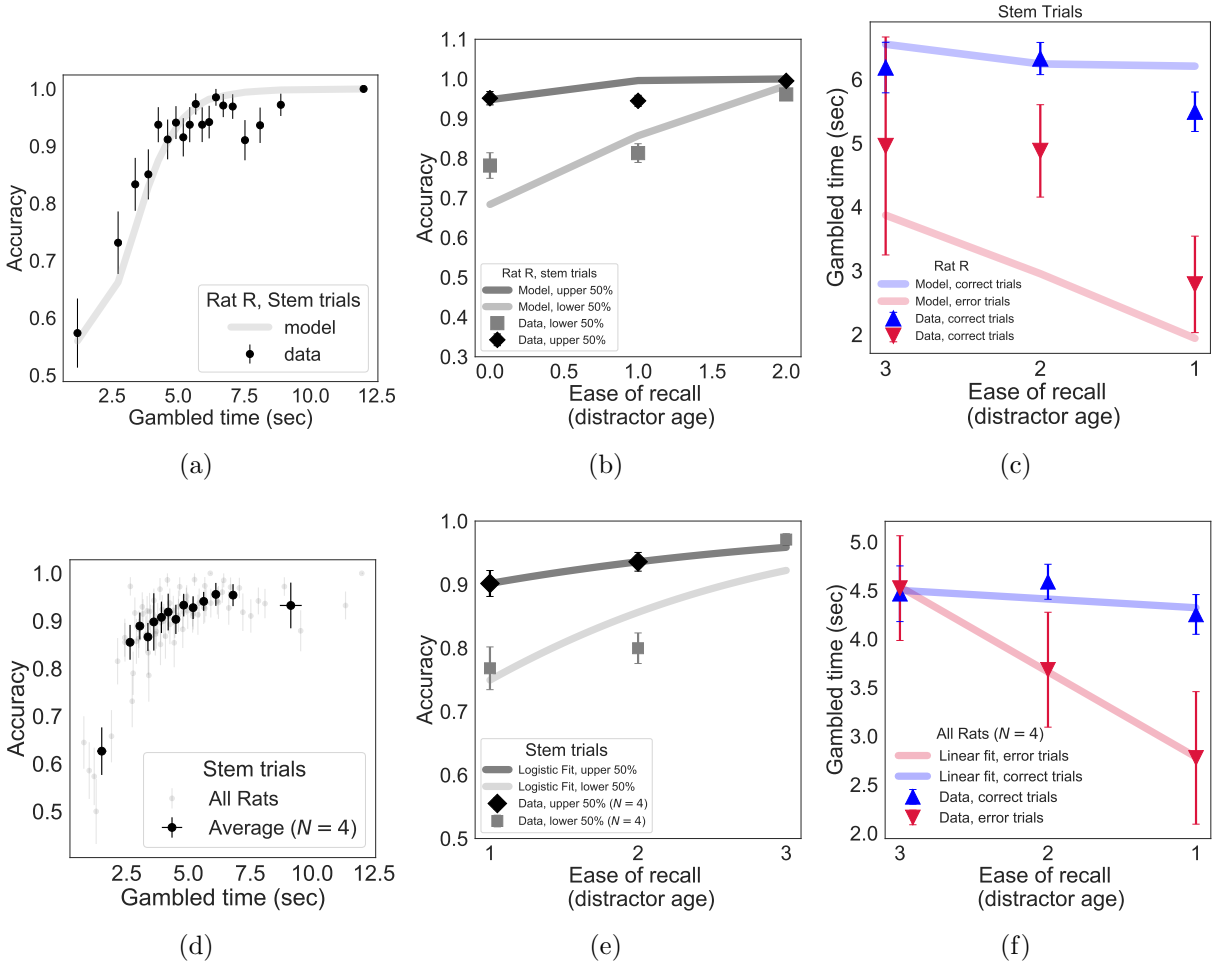

Figure 2.16: Model predictions for distractor age, stem trials only.

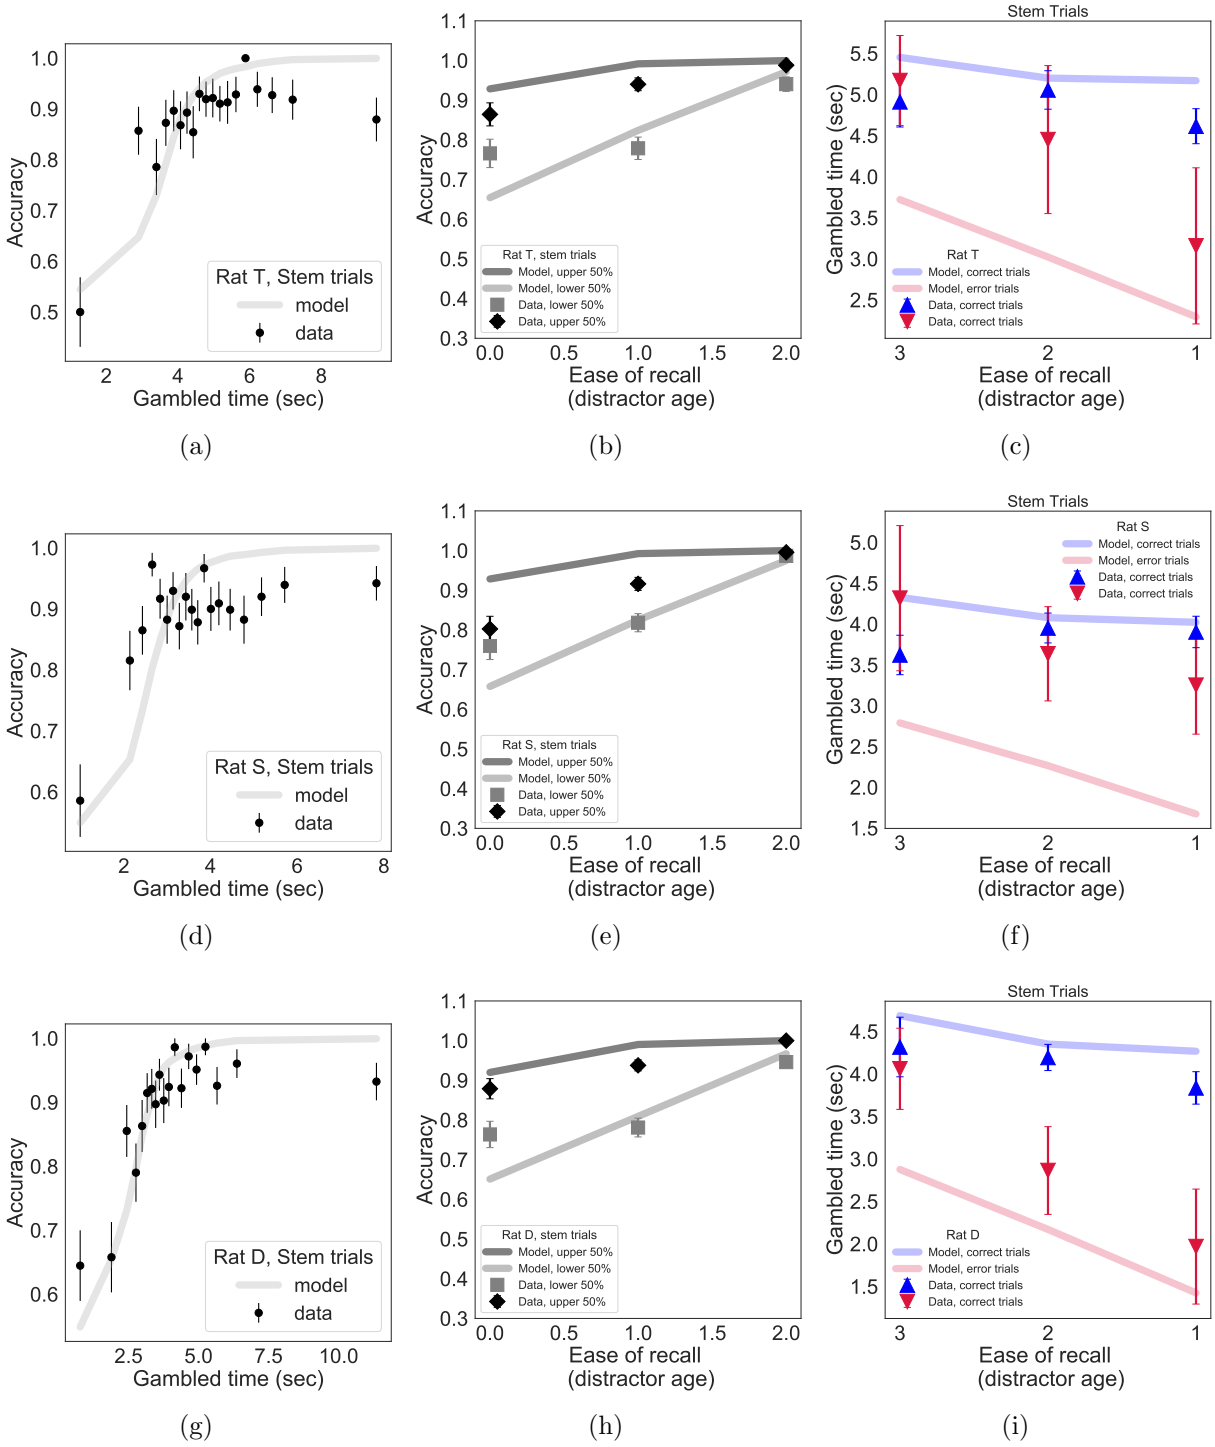

Figure 2.17: As for 2.16, data for all rats.

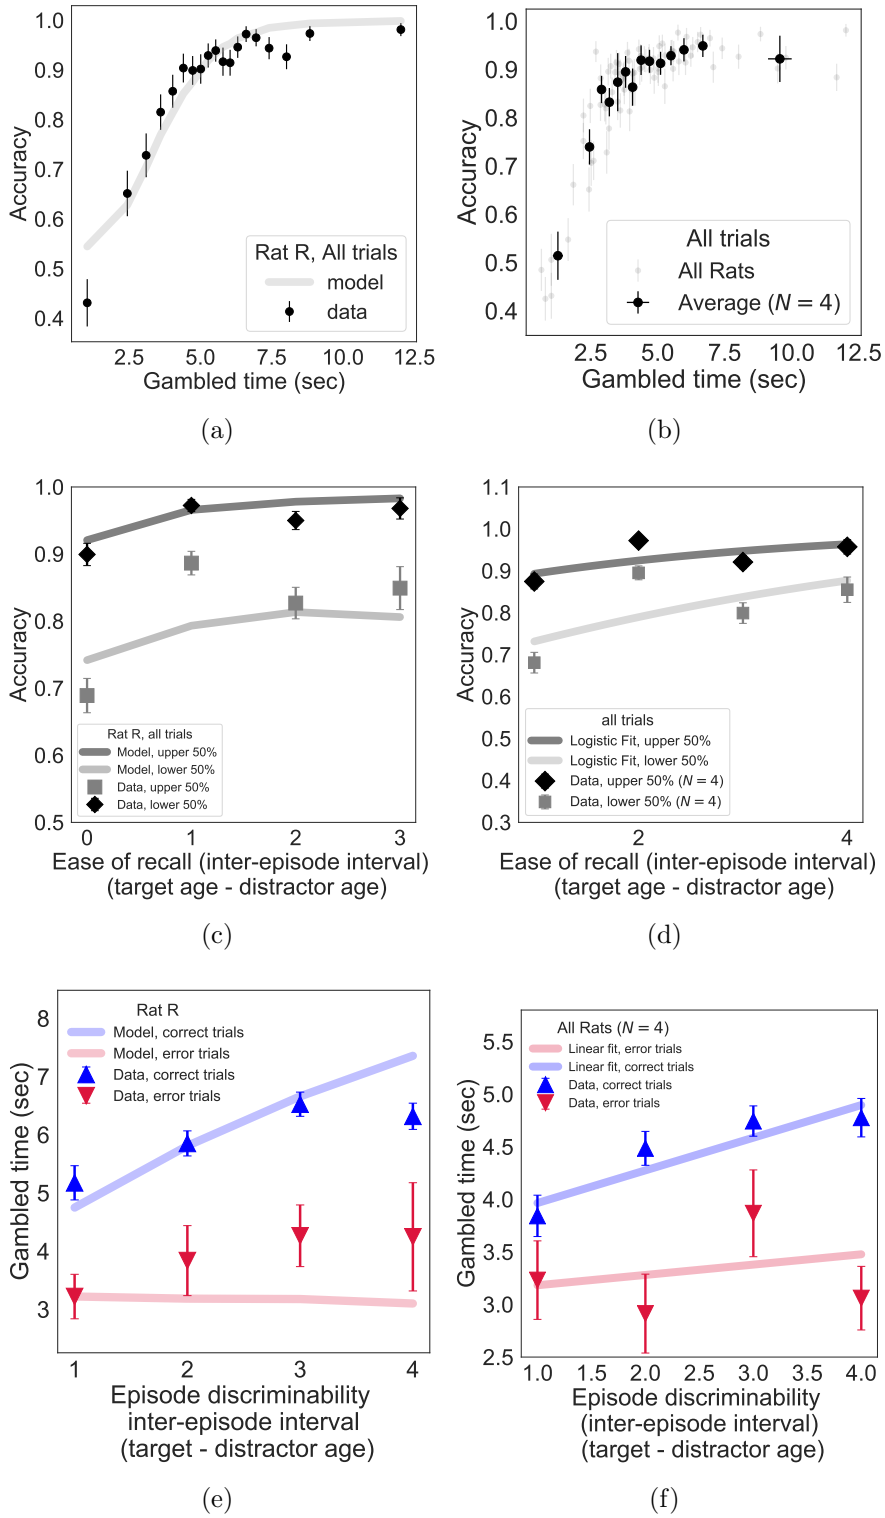

Figure 2.18: Model predictions for memory difficulty axis 'episode discriminability', in target - distractor age, for representative rat R and averages across rats.

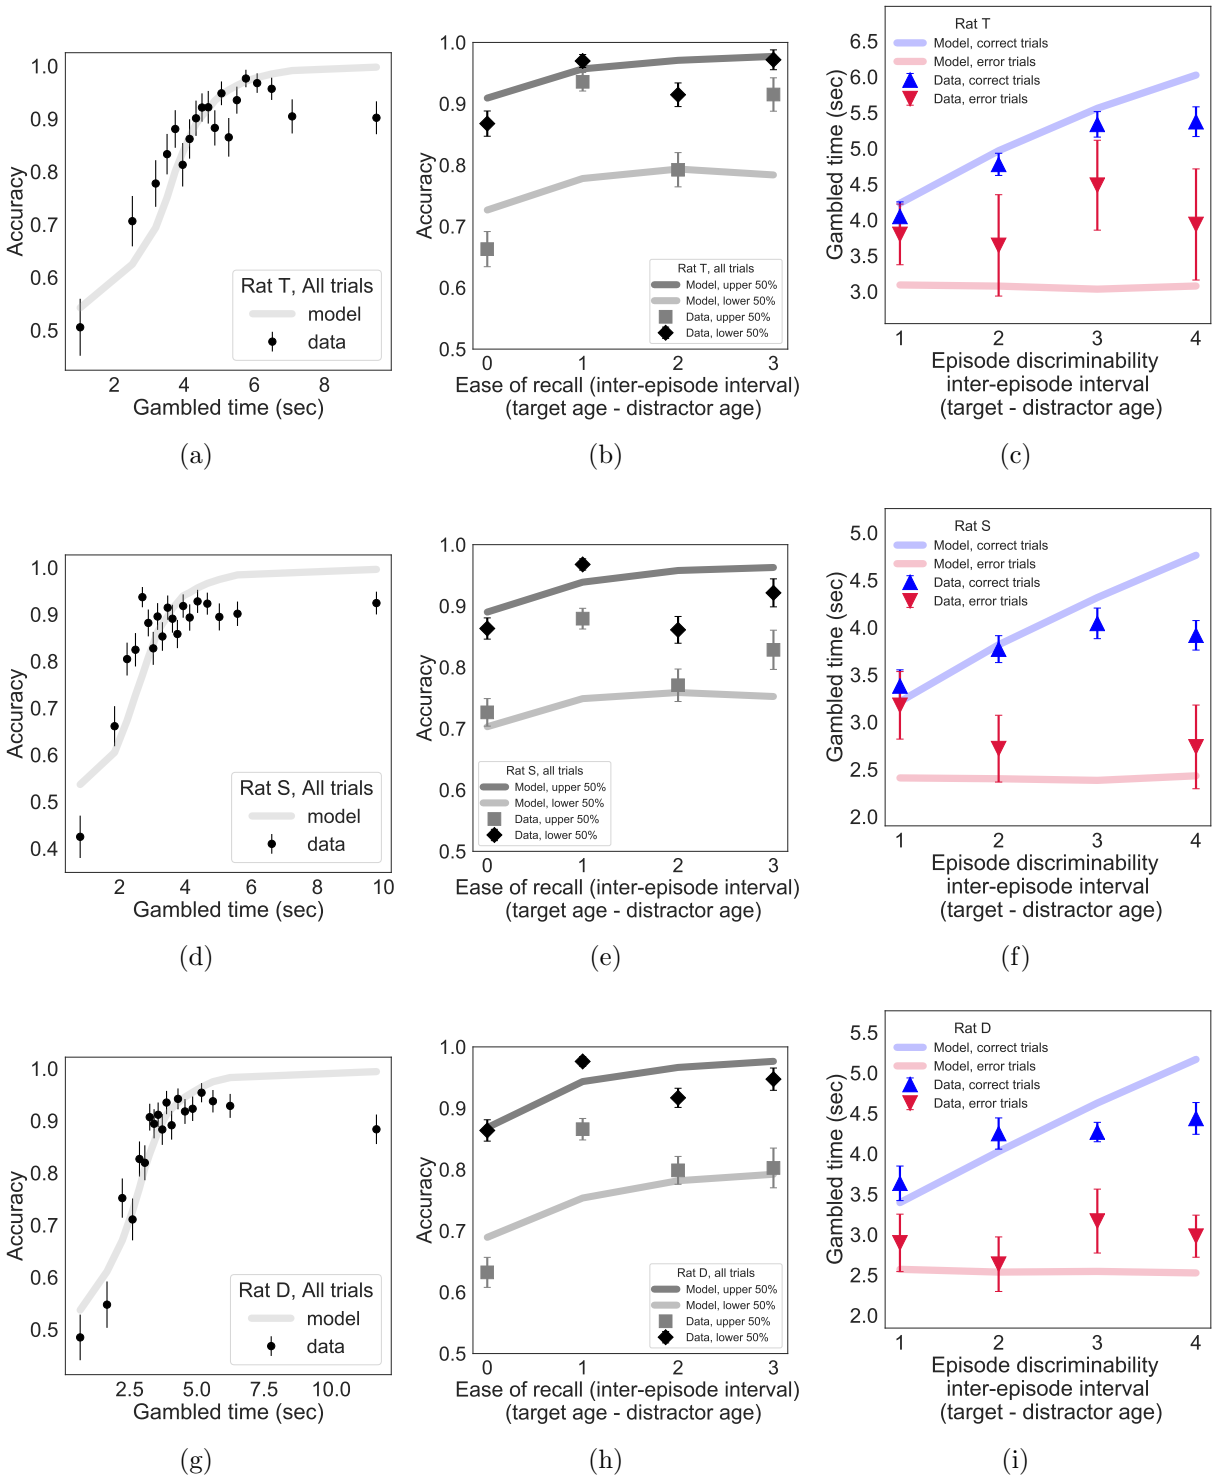

Figure 2.19: As for 2.18, all rats.

Chapter 3

A microfabricated, 3D-sharpened silicon shuttle for insertion of flexible electrode arrays through dura mater into brain

Abstract: Objective. Electrode arrays for chronic implantation in the brain are a critical technology in both neuroscience and medicine. Recently, flexible, thin-film polymer electrode arrays have shown promise in facilitating stable, single-unit recordings spanning months in rats. While array flexibility enhances integration with neural tissue, it also requires removal of the dura mater, the tough membrane surrounding the brain, and temporary bracing to penetrate the brain parenchyma. Durotomy increases brain swelling, vascular damage, and surgical time. Insertion using a bracing shuttle results in additional vascular damage and brain compression, which increase with device diameter; while a higher-diameter shuttle will have a higher critical load and more likely penetrate dura, it will damage more brain parenchyma and vasculature. One way to penetrate the intact

dura and limit tissue compression without increasing shuttle diameter is to reduce the force required for insertion by sharpening the shuttle tip. Approach. We describe a novel design and fabrication process to create silicon insertion shuttles that are sharp in three dimensions and can penetrate rat dura, for faster, easier, and less damaging implantation of polymer arrays. Sharpened profiles are obtained by reflowing patterned photoresist, then transferring its sloped profile to silicon with dry etches. Main results. We demonstrate that sharpened shuttles can reliably implant polymer probes through dura to yield high quality single unit and local field potential recordings for at least 95 days. On insertion directly through dura, tissue compression is minimal. Significance. This is the first demonstration of a rat dural-penetrating array for chronic recording. This device obviates the need for a durotomy, reducing surgical time and risk of damage to the blood-brain barrier. This is an improvement to state-of-the-art flexible polymer electrode arrays that facilitates their implantation, particularly in multi-site recording experiments. This sharpening process can also be integrated into silicon electrode array fabrication.

3.1 Introduction

Electrode arrays for implantation in the brain are a technology critical to both fundamental neuroscience and clinical treatments for diseases including epilepsy[160], retinal degeneration[161], Parkinson disease, and depression[162]. Classically, electrode arrays have been made of silicon[163,164] or other hard metal[165]. While these devices can be effective in recording single units[166], the longevity of recordings is limited[167,168]. More recent designs of silicon electrodes with smaller cross sections can record an estimated one cell per electrode, and can detect single units in mouse for at least 150 days[169], but the ability to record continuously from a given neuron over days has yet to be demonstrated. These limitations are thought to result from the stiffness of the device relative to brain tissue.

This has inspired the development of flexible neural implants. These include probes[119,167,170–172] and mesh[173] that are better matched mechanically to brain tissue and show reduced inflammation and immune responses[174–176,167,177,178], particularly if they are also small[179–181]. Chronically implanted in animal models, polymer devices can yield single-cell recordings spanning months[119,167,170,180], and many of the same individual neurons can be recorded continuously for at least ten days[119]. As penetrating neural devices are made smaller and more compliant to reduce tissue damage, surgical technique becomes more complicated. The force required to insert a flexible device into the brain typically exceeds the critical buckling load of the device. This presents a major barrier to use for many flexible devices[182]. Insertion is made more difficult as a result of the dense irregular connective tissue — the dura, arachnoid, and pia mater — that together protect the brain from mechanical and other insults[183].

Solutions to the problem of inserting a flexible electrode array into the brain involve surgical removal of the dura mater, then insertion through the remaining pia by methods that temporarily increase the critical load of the device using biodissolvable coatings[179,184–189], fluidic injection[190–193], freezing[172], or shape memory polymer (SMP)[194]. One straightforward solution is to temporarily attach the flexible array to a hard shuttle, such as silicon[187,195,196], or tungsten[197] that is shape-matched to the device itself and can be retracted once the array has been inserted to its target depth in the brain. Temporary attachment can be achieved by inserting shuttles through a small hole in the array[198] or, more commonly, by a biodissolvable adhesive. An advantage to this general method is that a very stiff material can be used to fabricate the insertion shuttle, since it will not be left in the brain. The shuttle diameter can thus be smaller, and the paired device very flexible, for reduced damage to the tissue on implantation and over the lifetime of a chronic implant[181]. Previously, we have used such a method to insert 16 independent polymer devices through rat pia, each to a different brain area, for high-quality, chronic recording[119].

If a flexible device could be inserted directly through dura with minimal brain compression, then the tissue damage and longer surgical time that result from durotomy itself would be eliminated. Furthermore, a dural-penetrating method could also avoid errors in depth targeting that result from outward swelling of the brain following durotomy. Together, these effects could yield high quality, chronic single-unit recordings. Of these previously mentioned insertion methods, only insertion using a delivery shuttle can be easily adapted for transdural delivery. For a shuttle made of a given material, with length determined by the depth of the target brain area, a lower limit on its cross-section is calculable by Euler's critical load equation. A thicker device will have a higher critical load to more likely penetrate the meninges, but is not desirable because it will compress a greater area of brain tissue on insertion and disrupt more vasculature[199]. One way to enable insertion of a device without increasing its diameter is to reduce the effective force on the device (i.e., to lower the required insertion force) by fabricating a sharper tip[200–203].

Previous work has demonstrated the benefits of sharpened silicon device tips in reducing insertion force[199,200,204] and penetrating dura[204,205]. A diamond[206] delivery shuttle for polymer devices has demonstrated successful insertion of flexible polymer devices through rat dura, as has use of an insertion guide with an SMP device[207]. However, such methods have not yet yielded chronic, single-unit recordings in freely behaving animals. We therefore aimed to develop a low-diameter, sharpened shuttle for polymer array delivery that could penetrate dura with minimal brain compression and yield high-quality single unit recordings. Here, we describe a fabrication method which sharpens silicon shuttles in 3 dimensions at any tip angle or crystal orientation and which uses readily available cleanroom tools. We demonstrate a reduction in insertion force compared to planar-sharpened silicon shuttles on insertion through rat dura. Finally, we show that compliant devices implanted by transdural shuttle delivery yield high-quality recording of local field potential (LFP) and single units over months in the awake behaving rat. This is the first in vivo demonstration of successful chronic neural recording using a method to

implant through a membrane with Young's modulus in the range 0.1-1 MPa[208], corresponding to the rat dura or primate pia. 3D-sharpened shuttles obviate the need for a durotomy, increasing the efficiency and reliability of insertion for polymer electrode arrays to that of dural-penetrating arrays while maintaining the desirable properties of polymer for high quality single-unit neural recordings. Importantly, our fabrication method can be readily adapted to provide similar benefits to silicon-based multielectrode arrays.

3.2 Materials and methods

We fabricated a sharpened shuttle for insertion of flexible, thin-film polymer probes and tested transdural insertion in vivo in the rat. In vivo tests and chronic recording were performed in 10-16 month-old, male Long Evans rats. To evaluate this method, we measured insertion force and brain compression, which have been shown to predict neural tissue damage[209]. First, we compared the insertion force through intact dura for shuttles that were either sharpened in three dimensions (sharpened) or two dimensions (planar) with otherwise identical dimensions. We first tested insertion through dura over the hippocampus, then tested insertion through the thicker dura over the OFC. Second, we calculated the brain compression in each case. Finally, we implanted polymer devices using sharpened shuttles through intact dura over the nucleus accumbens (NAc) and OFC, and recorded LFP and single units for 95 days.

3.2.1 Shuttle microfabrication

Planar and sharpened shuttles both had a 30° tip angle (Fig. 3.1), a width of 80 μm, and a thickness of 30 μm. The difference in their design was the gradual increase in thickness (the 'sharpened' profile) of the 3d-sharpened shuttle. Figure 3.2 shows SEM images of

planar and sharpened shuttles mounted on double-sided copper at magnification 300× with accelerating voltage 3kV using an APREO S Low Vacuum EDX Benchtop SEM (Fig. 3.1) or at magnification 650× with a Hitachi S-800 SEM (Fig. 3.2). Figure 3.3 is a profilometer scan, taken using a Veeco Dektak profilometer, in which the scan proceeds from the shuttle tip along the shank as it increases in thickness. For this representative shuttle, seven individual step heights, each approximately 5 μm higher than the previous, are seen over the 100 μm length of the tip. Both planar and sharpened insertion shuttles were fabricated on 4" silicon-on-insulator (SOI) wafers with 30 μm device layers (Fig. 3.4). First, 10 μm deep 'wicking' channels (Fig. 3.1, Fig. 3.2) were fabricated using standard photolithography and the Bosch Deep Reactive Ion Etch (DRIE) process in an STS ICP DRIE tool (SPTS Technologies). Wicking channels were designed to hold polyethylene glycol (PEG, molecular weight 10,000 mn), which acted as a biodissolvable adhesive to mount polymer probes on insertion shuttles, as described in Felix et al., 2012[196]. Next, 300 nm of aluminum was sputtered (Semicore Equipment Inc.), patterned photolithographically, and wet etched to create a hard mask with the insertion shuttle geometry. Sharpened shuttle fabrication involved a third photolithography step to specify the sharpened tip profile overlaid on the aluminum hard mask, followed by a 2-minute reflow bake of the photoresist (AZ4620) at 180° Celsius (Fig. 3.4). This resulted in a sloped profile in the patterned photoresist that was transferred to silicon using alternating timed silicon etch and photoresist etch steps (Fig. 3.4) in the same STS plasma etcher. The height of each 'stair' shown in Figure 3.3 is determined by a timed silicon DRIE step, which alternates SF_6 for etching and C_4F_8 for passivation. The length of each stair is determined by both a timed photoresist etching step using O_2 and the photoresist reflow step (Fig. 3.4). Uneven height and length of the stairs can arise from either imperfect photoresist reflow or imperfect timing of etching steps. Dimensions of the stairs can be adjusted to etch arbitrary 3D profiles, in contrast with anisotropic wet etching methods for which angles are restricted by silicon crystal structure. Planar shuttles underwent DRIE without the third photolithography step, resulting in a

tip profile of uniform thickness. Lastly, the aluminum hard mask was removed (Piranha etch) and individual shuttles were released from the SOI wafer using hydrofluoric acid.

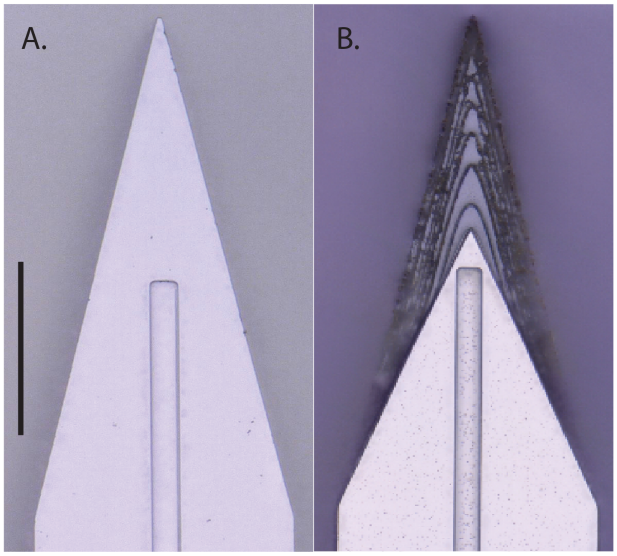

Figure 3.1: Top down view by light microscope (scale bar = 50 μm) of (A) planar and (B) sharpened profiles on 2d- and 3d-sharpened silicon shuttles, respectively, oriented with tips toward the top of the figure, and wicking channel for PEG at midline.

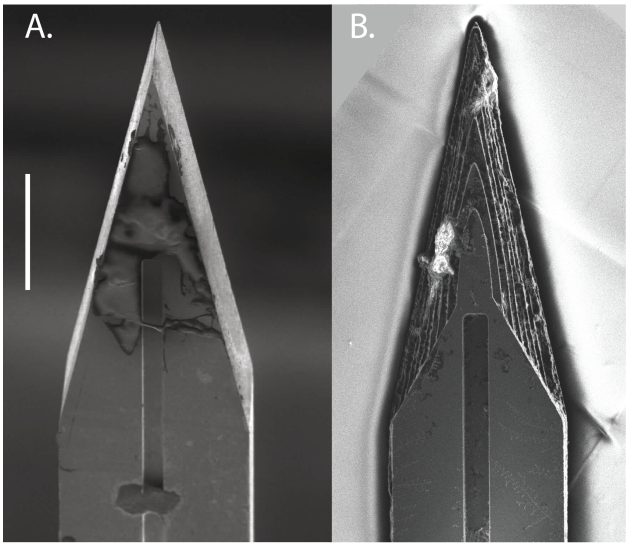

Figure 3.2: Top-down view by SEM for comparison of planar (A) and sharpened (B) shuttles (scale bar = 50 μm). Shuttles are oriented with tips toward the top of the figure and wicking channel for PEG at midline. Residue in (B) from previous insertion through neural tissue.

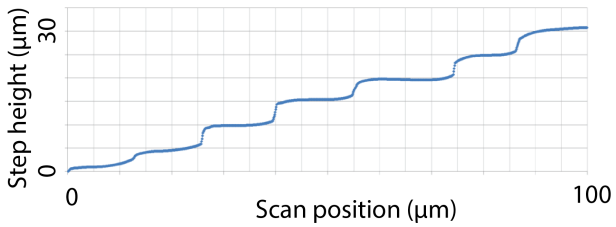

Figure 3.3: Profilometer scan of sharpened shuttle tip.

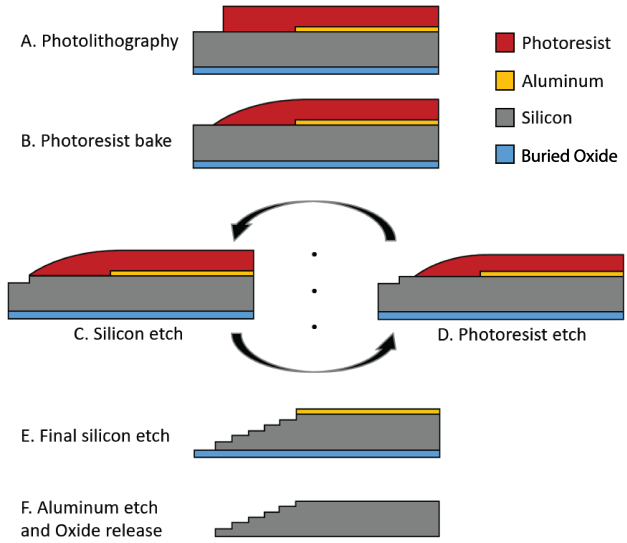

Figure 3.4: Process flow detail.

3.2.2 Polymer probe microfabrication

Polymer devices used in chronic recordings were made of polyimide with platinum electrodes electroplated with PEDOT:PSS. Each device had two shanks separated by 250 μm, each 80 μm wide and 6 or 8 mm long. As shown in Figure 3.5, each shank had 16 contacts arranged in an offset dual-line configuration of 8 contacts each. Contacts were 20 μm in diameter with 20 μm contact edge-to-edge spacing, and 6 μm from electrode edge to shank edge. The image in Figure 3.5 was taken using Keyence Model VH-Z250R at $\sim 250\times$, with illuminated lighting on auto detect for white balance and brightness.

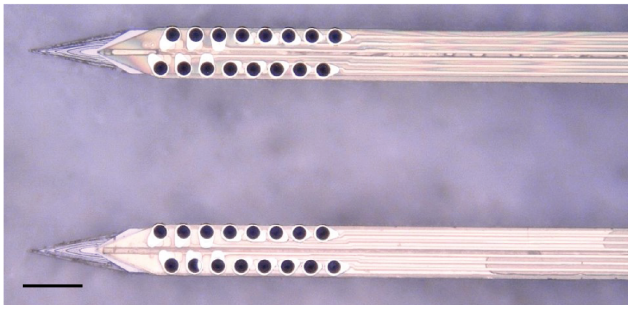

Figure 3.5: Top-down view by light microscope of a single, two-shank, 32-channel probe mounted on a sharpened, two-shank silicon shuttle. The shuttle tip extends approximately 100 μm past the tip of the probe, which appears yellow-gold atop the shuttle and has, on each shank, two offset rows of eight PEDOT-PSS plated platinum contacts. Distance between shanks is 250 μm. Scale bar = 100 μm.

3.2.3 Surgical insertion *in vivo*

We used single-shank silicon shuttles (Fig. 3.1) for in vivo insertion force tests. Single-shank shuttles were created for shuttle-only tests by cutting off one of the shanks on a two-shank shuttle. All animal-involved protocols described in this manuscript have been approved by the Institutional Animal Care and Use Committee at UCSF. Studies were performed in male Long Evans rats. For the comparison of force insertion measurements between planar and sharpened shuttles in vivo, two animals were used, ages 12 and 16 months. For sharpened shuttle-guided insertion of probes for chronic neural recordings, one animal, age 10 months, was used.

In all cases, animals were anesthetized with inhaled isoflurane plus a mixture of ketamine, xylazine, and atropine and were confirmed unresponsive to a foot pinch. The head was shaved and the animal transferred to the sterile surgery environment and head-fixed in a stereotactic frame. Body temperature was maintained throughout surgery by an isothermal pad beneath the animal. Following sterilization of the incision site and application of lidocaine as a local anaesthetic, a minimal incision was made through all skin layers at the top of the skull. The skin flaps were retracted and the tissues detached from the bone, preserving the attachment of the temporalis muscle to the temporal ridge. For the chronic

recording animal, a dental drill was used to make 12 sub-penetrating skull holes along the temporal ridge. Into each of these holes, 0–80 titanium set screws (United Titanium, OH) were partially inserted, serving to anchor the implant to the skull. These were drilled with a rotating dental drill with carbide bur (SS White carbide bur, FG 2) at 6000 rpm, at an intermittent rate, with the drill removed completely from contact with the bone between bouts of drilling. Coordinates for target insertion sites were located using bregma and stereotactic coordinates. One craniotomy of approximately 2–3 mm diameter was made to expose the meninges over each of NAc (1.5 AP, ± 1.3 ML, mm) and OFC (+3.6 AP, ± 3.4 ML, mm) bilaterally (for chronic recordings), or each of hippocampus (-3.6 AP, ± 2.5 ML, mm) and OFC bilaterally (for force insertion measurements). The craniotomy technique was the same as for skull screw holes with the exception that the final ∼500 μm of bone was removed using a smaller carbide bur (SS White carbide bur, FG 1/4). The dura mater was maintained fully intact. Following *in vivo* data collection, animals were sacrificed using an overdose of pentobarbital sodium and phenytoin sodium (Euthasol, Virbac AH, Inc.).

3.2.4 Insertion force measurements

Insertion force tests of shuttles without attached devices were performed over hippocampus and OFC bilaterally in each of two animals. We performed 16 insertions in Animal 1 (4 sharpened and 4 planar to hippocampus; 4 sharpened and 4 planar to OFC) and 11 insertions in Animal 2 (4 sharpened and 2 planar to hippocampus; 2 sharpened and 3 planar to OFC). Each test insertion was performed with a new, single-shank shuttle over a new area of dura. In both experiments, insertion force measurements were performed with a precision load cell (FUTEK FSH02534) and voltage was read out through the digitizer (FUTEK FSH03633) to a computer at 100 samples/second. The largest source of error in our force measurements is in the readout of the digitizer, which is load-dependent and equivalent to ±0.067 mN for our measured loads. Rigid adapters between the motor, load

cell, and insertion shuttle were 3D-printed with hard plastic from custom designs (PolyJetHD Blue, Stratasys Ltd.; Fig. 3.6). Insertions were conducted at a constant velocity of 50 μm/s using a micropositioner (Kopf model #2662) and insertion force was measured as the shuttle contacted dura and either inserted or failed (Fig. 3.7). To match our typical surgical conditions, we tested serial insertions through dura in the same craniotomy, with at least 300 μm separation between insertion sites. Each insertion was performed with a shuttle through intact dura. The dura and brain were hydrated for the duration of the surgeries using hand irrigation with saline.

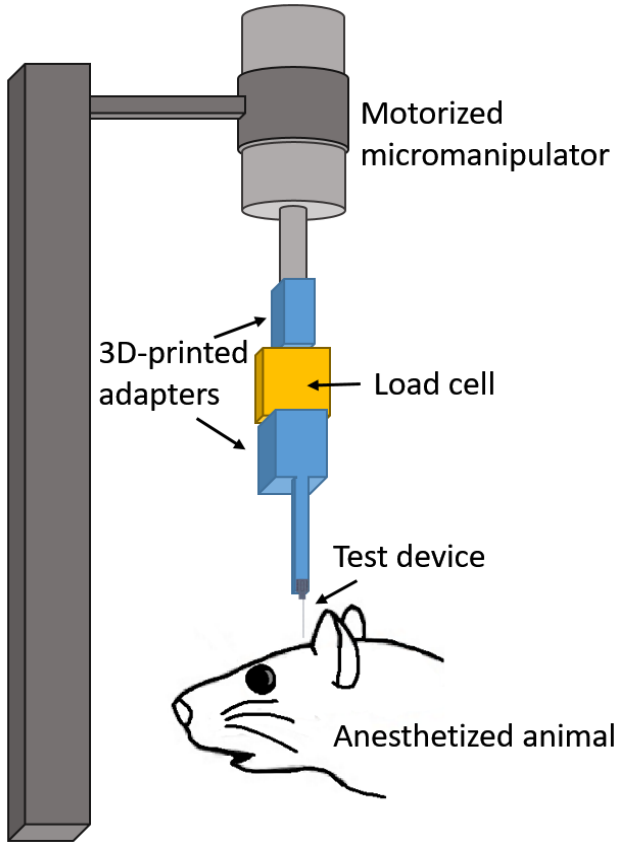

Figure 3.6: Test insertion apparatus for in vivo insertion force measurements. The test device was a single shank, 6 mm long.

Shuttle insertions were classified as successes if they insert through the dura without signs of buckling. Consequently, insertions were classified as failures if they visibly buckled, regardless of whether they eventually penetrated dura. Maximum insertion force was

measured at the time of penetration, which was taken as the first absolute maximum of the force curve that preceded the first approximately infinite-slope curve characteristic of membrane penetration[199,210]. In a few cases, a slightly higher force was measured after initial penetration. In these cases we did not use the later timepoint to determine the maximum insertion force because, by then, the device had already begun to penetrate tissue and the total force measured thus included the frictional force as well.

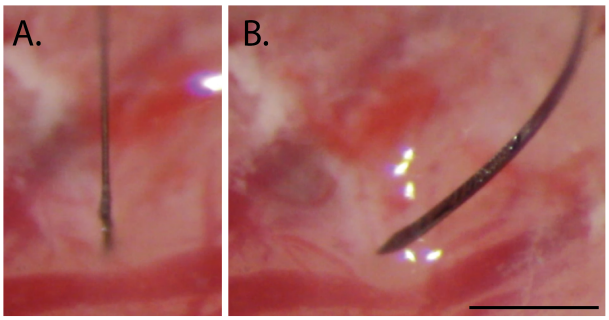

Figure 3.7: Example of a successful (A) and failed (B) insertion attempt for sharpened silicon shuttles, 6mm. Scale bar = 1 mm.

Brain deflection distances were calculated by multiplying the insertion speed (50 μm/s) by the time between dural contact and penetration. We estimated the time of dural contact by first calculating the mean and standard deviation of the (150-sample minimum) baseline period, during which the shuttle was still above tissue and contacting only air; we then determined the point at which a 5-sample sliding average of the measured force reached at least one standard deviation above the baseline mean. The time of penetration was chosen to be the first large peak of the force trace. The load cell itself compresses as a function of force applied, and the amount of compression was measured under a microscope and subtracted from the brain deflection values calculated from force measurement data as described above. This resulted in a correction of 5 to 12 μm, which is within the variance of deflection values between insertions.

3.2.5 Imaging

Shuttles were imaged using Keyence Model VH-Z250R at $\sim$1000$\times$, with illuminated lighting on auto detect for white balance and brightness. Eight planar and eight sharpened shuttles were each imaged both before and after in vivo insertions through dura to verify that no shuttle breakage occurred during insertion or retraction.

3.2.6 Chronic recordings

All procedures were in accordance with guidelines from the University of California San Francisco Institutional Animal Care and Use Committee and US National Institutes of Health. For the tests of shuttles used with chronically implanted probes, the probe and shuttle were adhered to each other by PEG on the side opposite the recording contacts, as previously described for planar silicon shuttles. The meninges were hydrated between craniotomy completion and the time of insertion using hand irrigation with a saline drip. When all craniotomies were complete, a custom designed, 3D-printed plastic base piece was secured to the skull using dental acrylic attached to the 128-channel headstage (SpikeGadgets, LLC). These insertions were performed manually, lowering with a stereotax at approximately 50 μm/s until the probe penetrated the dura mater. Once in the brain, the device was lowered using a micromanipulator (MO-10, Narshige) at 10 μm/s until it was 1 mm above target depth, then at 5 μm/s until it was 500 μm above target depth, and at 3 μm/s until it reached target depth.

The four implant targets were: OFC bilaterally, at the same coordinates for the insertion force measurements, at depth -4 mm from brain surface; and NAc at depth -7 mm from brain surface. Following insertion, probes were secured to the base piece and the electronics connected. Once the probe reached target depth, polyimide "wings" glued to the probe perpendicular to its length were adhered to the base piece with acrylic, securing

the probe at the target location. The shuttle was then detached from the probe by filling the base piece with saline, which dissolved the PEG at the probe-shuttle interface, thus allowing retraction of the shuttle without disrupting the probe. After all the probes were inserted, the base piece was drained of saline and filled with silicon elastomer (Dow Corning 3-4680) to prevent dural regrowth and scarring, brain swelling, and mechanical perturbation of the probes[211]. Kwik-sil (World Precision Instruments, LLC), dental acrylic, and a plastic case made of moldable plastic (ThermoMorph) was used to encase the devices. Post-surgically, meloxicam (Eloxiject, Henry Schein) and buprenorphine (Baytril, Reckitt Benckiser Healthcare) were given for pain, and a single dose of enrofloxacin (Baytril, Bayer Corporation) was given to prevent infection. Neural recordings were conducted in an approximately 1 square foot sleep box constructed of anti-static plastic and located in a recording room. Data were collected for 30 minutes or longer per day using the SpikeGadgets recording system (Trodes version 1.74), as previously described[119]. The recording session analysed here was 1 hour long. Typically, the animal was asleep or quietly immobile during the recording period. The animal also ran a spatial navigation behaviour in epochs distinct from the sleep epoch analysed here. Recording quality was analysed at 95 days, as we have found previously that cell count using these polymer probes stabilizes by this time[119].

3.2.7 Neural data analysis

Data pre-processing was performed using custom Python and Matlab scripts. Common-average referencing was applied across the sixteen channels of each shank. Spike sorting was performed using the MountainSort software package[39,119,212], version 4.0. An initial round of automated sorting was performed with the following sorting parameters: detect sign = -1, detect threshold = 3, clip size = 100, adjacency radius = 200 μm. The raw data were filtered between 600 and 6000 Hz. The detect interval was set to 10 samples

and the first 10 principal components were used. For units identified using these parameters, only those with cluster quality metrics above the following thresholds were included as single units: firing rate threshold = 0.01 Hz, isolation threshold = 0.96, noise overlap threshold = 0.03, peak signal-to-noise ratio threshold = 1.5. The rest were marked as multi-unit activity. All identified units, including those marked as multi-unit activity, were manually inspected and curated: in MountainView software, clusters that did not appear to be single units based on refractory period violations (*i.e.*, frequent spiking within 2 ms of the last spike) were rejected; multiple clusters that were identified as corresponding to the same unit based on a combination of firing rates, waveforms, peak channels, and temporal cross-correlograms were merged, and all cluster metrics listed above were re-calculated. Clusters that passed the curation metric thresholds were accepted.

3.3 Results

3.3.1 Planar and sharpened shuttle insertion forces

We first tested single-shank shuttles with either a flat profile (planar; $N = 6$) or a sharpened profile (sharpened; $N = 8$) but otherwise identical dimensions (6 mm × 30 μm × 80 μm, Fig. 3.1 and Fig. 3.2) in transdural insertion to the hippocampus (Fig. 3.8). Sharpened shuttles penetrated dura (8/8 shuttles, 100%) but planar shuttles did not (0/6 shuttles, 0%). Figure 3.8 shows the maximum insertion force for each insertion test over hippocampus. One shuttle characterized as failed did penetrate dura, but only after buckling to an extent that would likely have caused separation from an attached polymer probe due to the slight difference in curvature of the probe relative to the shuttle during buckling. For successfully inserted shuttles, the average insertion force was 8.9 ± 1.6 mN ($N = 8$, ± s.d.). In contrast, failed penetrations resulted in buckling and had

an average buckling force of 13.6 ± 1.8 mN ($N = 6$, ± s.e.m.; Fig. 3.8A, left).

The insertion force traces for shuttles that successfully penetrated the dura showed increasing force, followed by a characteristic series of rapid decrements between local maxima, in a sawtooth pattern (Fig. 3.9). In a few cases, a slightly higher force was measured following initial penetration (Fig. 3.9, panels 3, 9, 10, 11 from top left to bottom right), likely due to increased friction force after the device had begun to penetrate tissue. In these cases we took the first local maximum as the maximum insertion force.

Our hippocampal target was located approximately halfway between bregma and lambda, and halfway between the midline and the temporal ridge. As a test of the effectiveness of our sharpened shuttles through thicker dura, we next performed the same tests over the OFC, where the dura is thicker and more difficult to penetrate. Tests were performed for OFC as for hippocampus (Fig. 3.7). Of the sharpened profile shuttles tested, 3/6 shuttles (50%) penetrated dura, while 0/7 planar shuttles (0%) did. The worse performance over OFC relative to hippocampus is expected given the tougher dura. Figure 3.7 shows the maximum insertion force for each insertion test over OFC. Two sharpened shuttle insertions characterized as *failed*, one over left OFC and one over right OFC, successfully penetrated dura but did so only after buckling to an extent that would likely have caused separation from an attached device.

For successfully inserted shuttles over OFC, the average insertion force was 10.6 ± 1.5 mN ($N = 3$, ± s.e.m.; Fig. 3.8, right; insertion profiles also included in Fig. 3.9), which is not significantly different than for hippocampus (Welch's t-test: $p = 0.24$). In contrast, failed penetrations included both sharpened and unsharpened shuttles and had an average buckling force of 14.2 ± 0.6 mN ($N = 10$, ± s.e.m.; Fig. 3.8, left).

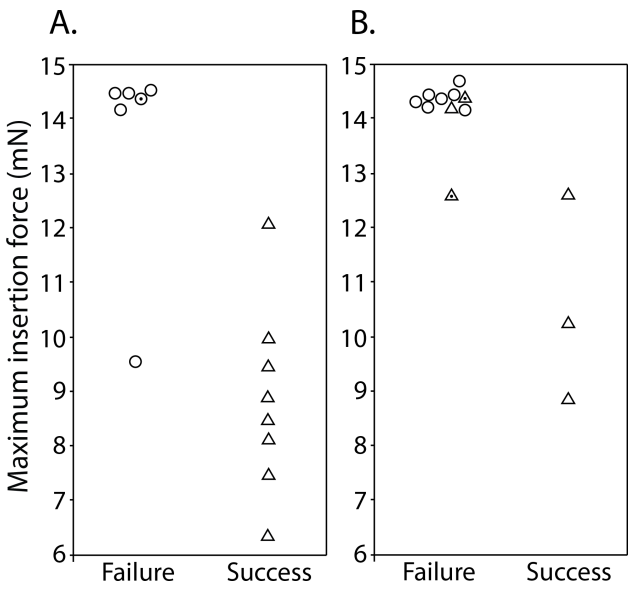

Figure 3.8: Maximum insertion force for sharpened (triangles) and planar (circles) shuttles inserted to (A) hippocampus (B) OFC that either failed (left) or succeeded (right). Y-axis = maximum insertion force in mN. Random jitter applied along x-axis for visualization of failed insertions. Points with dot centers indicate shuttles that penetrated dura but only after buckling to an extent that would likely have caused separation from an attached polymer probe.

3.3.2 Brain compression on transdural shuttle insertion

We calculated the degree of brain deflection for sharpened shuttles successfully inserted through dura without buckling (*i.e.*, those insertions included in Fig. 3.9). Deflection distances were calculated using the number of samples between dural contact and penetration (refer to Methods for details). On average, the calculated compression for sharpened shuttle insertions through dura was 389.9 $\pm$ 159.6 μm ($N = 8$, $\pm$ s.e.m.) for hippocampus and 302.5 $\pm$ 158.5 μm ($N = 3$, $\pm$ s.e.m.) for OFC.

3.3.3 Imaging before and after insertion

We imaged 8 sharpened and 8 planar shuttles before and after *in vivo* insertion. We observed no breakage or other discernible damage to any of the shuttle tips, indicating that

sharpened shuttles, despite having a tip diameter of less than 3 μm in width and thickness, maintain their structural integrity and do not break off in the brain (Fig. 3.10).

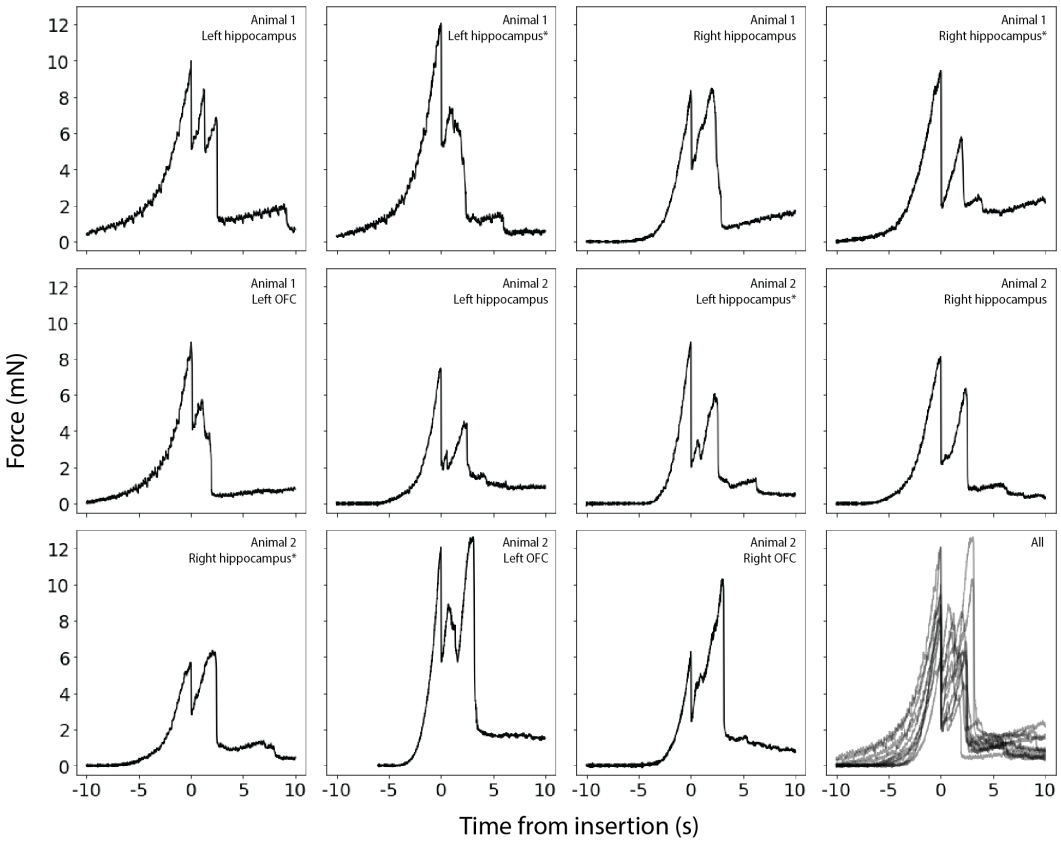

Figure 3.9: Raw force traces for all sharpened shuttles inserted through dura without buckling. X-axes in seconds, with traces aligned to 0 as the timepoint of maximum force, y-axes in mN. The animal and area of insertion are labelled in the upper right of each panel; plot at lower right shows all insertion force measurements depicted here overlaid. If the insertion was through a different area of a craniotomy where a section of dura at least 300 μm away was already used for a test insertion, the insertion site is marked with an asterisk.

3.3.4 Neural recordings using sharpened shuttles

Device-attached shuttles entered the brain with minimal compression of the dura and brain, with no observable difference relative to shuttle-only insertions. This is expected, given that devices were intentionally attached to shuttles such that their tips were shifted approximately 100 μm back along the shuttle shanks (Fig. 3.5). This enabled the bare

sharpened shuttle tip, without the device, to form the penetrating point. We observed that after the dura was punctured, tension in the membrane slightly expanded the opening to allow the probe, which itself is 14 μm thick, to enter the brain. After dural penetration using a stereotax, the device was inserted to its final depth using a micromanipulator. We observed that single units were detected on all six implanted shanks (16 channels each) for at least 90 days. We selected a single epoch on day 95 post-implant to perform spike sorting on one of the shanks targeted to the left OFC. This shank was selected for its electrode quality prior to insertion (no shorts or dead channels) and relatively low noise levels (Fig. 3.11). We identified 18 single units on this 16-channel shank, for an average of approximately one unit per electrode (Fig. 3.11). This is similar to what we have previously observed for polymer arrays inserted using planar shuttles[119]. Single units that passed curation standards were evenly distributed over the recording channels, with no apparent systematic difference in quality, from the top to bottom of the shank.

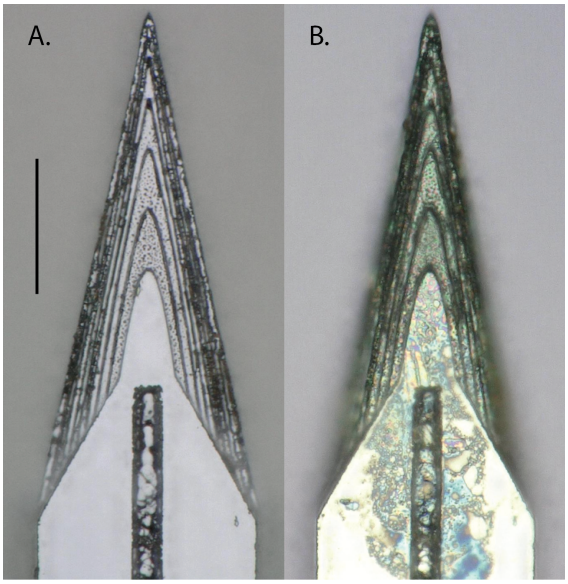

Figure 3.10: Top-down view of the same sharpened shuttle. A. Before insertion. B. After insertion through dura. Scale bar = 50 μm.

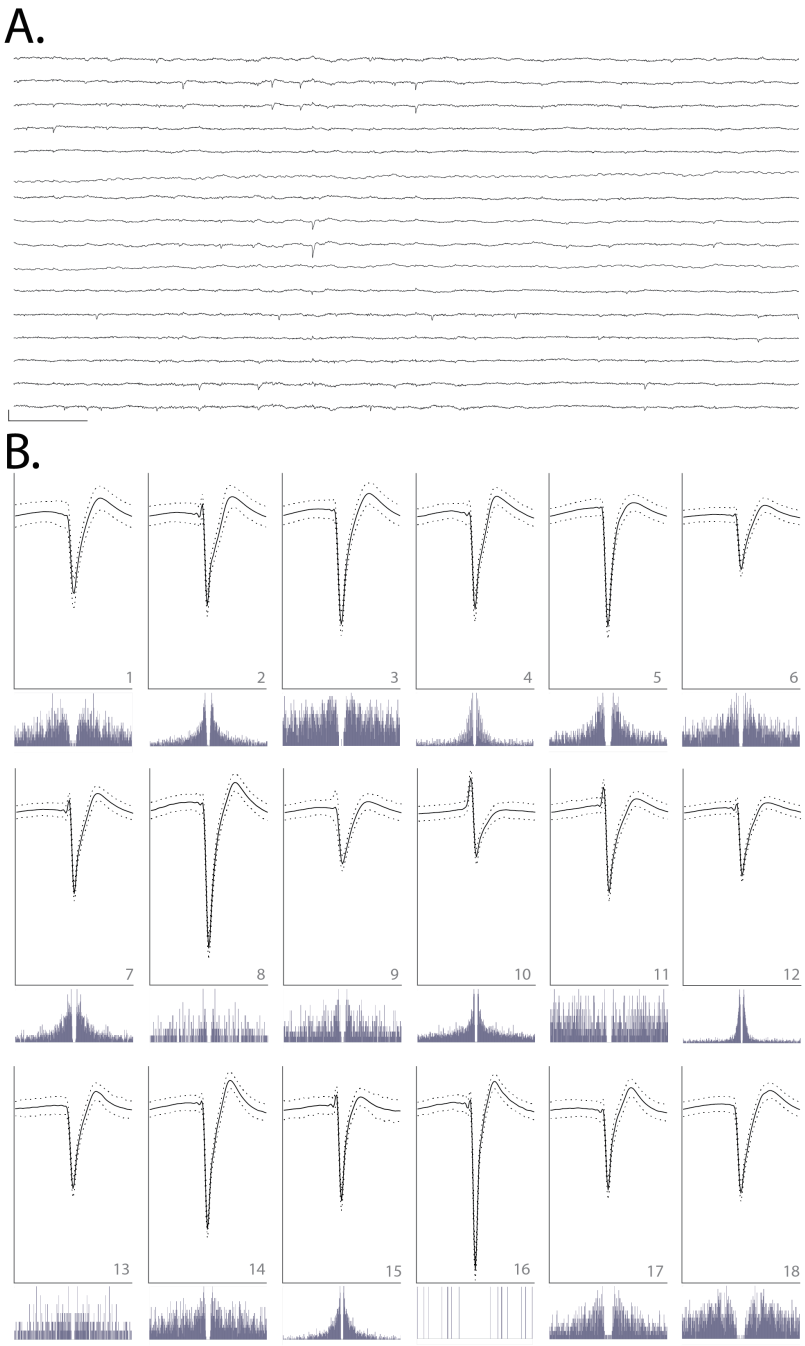

Figure 3.11: A. Raw, 150-millisecond local field potential recorded on each of the 16 channels of one shank targeted to OFC, 95 days after implant. Channels 1-16 are ordered from top to bottom. Channels 1–8 correspond to electrodes on the left side of the probe shank, from top to bottom; channels 9-16 correspond to electrodes on the right side of the device, from bottom to top (i.e., channels 1 and 16 are at similar depth). Scale bar, bottom left = 1 mV vertical, 15 milliseconds horizontal. B. Each of 18 units detected on the 16-channel shank shown in A. Odd rows: average waveforms (solid line) ± one standard deviation (dashed line) for units 1-18, numbered in bottom right of each panel. For each waveform, vertical scale bar = 2.5 millivolts, horizontal scale bar = 100 samples (at 30 kHz), or approximately 3.33 milliseconds. Even rows: spike auto-correlograms for the unit shown above, spanning 100 milliseconds, in 0.5-millisecond bins.

3.4 Discussion

Previous work has demonstrated the benefits of sharpened device tips in reducing insertion force (Sharp et al., 2009,[199,204] and penetrating dura[200,204,205]. Here, we sharpen not only in two dimensions but three, resulting in a more gradually increasing cross-section[204]. This is, to our knowledge, the first demonstration of a dural-penetrating, polymer device insertion method to yield chronic single unit recordings. It also allowed us to directly compare shuttles made of the same material (silicon) that were sharp in three versus two dimensions. We have shown that sharpened silicon shuttles result in low insertion forces and brain compression. The expected buckling force, also known as the critical load, of our silicon shuttle can be calculated with Euler's critical load equation:

$$F_{\text{crit}} = \frac{\pi^2 E I n}{L^2}, \tag{3.1}$$

where E is the Young's modulus of silicon, I is the moment of inertia for the cross-sectional area of the shuttle, n is a factor accounting for end conditions, and L is the length of the shuttle. The moment of inertia I of a rectangular cross-section is $ab^3/12$ where a and b are the longer and shorter side lengths, respectively. We take 169 GPa as the Young's modulus of our silicon shuttles, which are fabricated in the $\langle 110 \rangle$ crystal directions[214]. For translation fixed, rotation fixed and translation fixed, rotation free (brain side) end conditions, $n = 2$ and the expected critical load is:

$$F_{\text{crit}} = \frac{\pi^2 \, (169 \text{ GPa}) \times \left(80 \ \mu\text{m} \times \frac{(30 \ \mu\text{m})^3}{12}\right) \times 2}{(6 \text{ min})} \tag{3.2}$$
$$= 16.7 \text{ mN}.$$

This is close to our measured result of 13.6 mN for hippocampus and 14.2 mN for OFC. Due to the deformable nature of the brain, the boundary condition on the brain side is not completely translation fixed, with true n less than 2 and true F_{crit} expected to be slightly lower than calculated. We did not, however, visually observe the shuttle tip translate across the dural surface during insertions.

Our maximum insertion force for transdural insertion of sharpened shuttles of approximately 10 mN is lower than previous reports for planar devices (e.g., 41 mN with silicon probes;[205] and similar to the 11 $\pm$ 2 mN reported previously in a proof-of-concept study in which a less flexible, silicon-specific etch technique was used to insert but not record from silicon probes[204]. Table 1 summarizes previously published insertion force measurements made through rat dura. For comparison, previously reported peak insertion forces for rat pia-only insertions are approximately 5 mN[204,215]. In addition to those devices included in Table 3.1, other silicon arrays have been inserted transdurally in cat[216] and rat[217] without measurement of insertion force. For comparison, commercially available silicon Neuronexus probes range from 15 – 50 μm thick and the state-of-the-art silicon Neuropixels probe is 70 $\times$ 20 μm[169]. These dimensions are comparable to the smallest devices in the table. The design and fabrication technique presented here could be used for these silicon electrode arrays to enable dural penetration.

Our insertion force traces each show multiple peaks separated by approximately 100 μm in depth, which in some cases (traces 3, 9, and 11 from the upper left) include a peak that is

higher than the first. Additional force peaks over the course of insertion could be due to penetration of the pia and arachnoid mater, which lie between the dura mater and brain parenchyma. Such a sawtooth or non-monotonic pattern has been observed previously by other experimenters performing device insertions into the brain[199,210] though not in all cases[200,204,206], and is expected to depend on the shape of the device and its insertion speed. In cases where such a pattern has been observed, it has been suggested that myelinated axon bundles or other heterogeneities in the brain parenchyma itself, below the insertion site, could account for the insertion force profile[199]. Another possibility is that this pattern would be observed even in absence of the underlying brain parenchyma. The dura is a dense irregular connective tissue, made of collagen and elastin, and it is conceivable that the first peak on the insertion force trace could correspond to the tip of the shuttle breaking superficial matrix fibers, with subsequent peaks corresponding to further breakage of deeper fibers. Such a process has been observed for rat dura-only penetrations[208]. Heterogeneities in the brain parenchyma and dura exist not only across depths at a given insertion site, but at different anteroposterior and mediolateral coordinates. This may explain the variation in force traces for insertions performed at different locations within the same craniotomy, which will present different vasculature, cellular composition, and overlying dural matrix fibers.

Recently, an insertion method using a diamond shuttle tested rat transdural insertion at a variety of speeds, finding a minimum average compression of 488 μm at insertion speed 10 μm/sec with 200 Hz vibration[206]. At approximately 300–400 μm of compression without piezovibration, our method is comparable. To our knowledge, these are the lowest brain compression values for a working device through dura. It is a reasonable expectation that for surgeries where performing a durotomy is preferable, these devices can be implanted into the brain through pia with even lower brain compression. This use could be particularly valuable for superficial cortical sites in which recording quality is easily compromised by compression of the brain[218].

Force and tissue compression are important metrics to gauge neural tissue damage, but a single unit count weeks after implant is required to fully demonstrate a neural recording technology for chronic use. Recently, promising strategies for insertion of flexible polymer devices through rat dura have been demonstrated[197,207], but such methods have not yet reported chronic, single-unit recordings in freely behaving animals. Here, we evaluated our single unit count at 95 days. At that timepoint we were able to record high quality single units on 11 of 16 (69%) channels, with a total of 18 sorted units on 16 channels, or an average of 1 high-quality unit per electrode. In comparison, the other recent technology that reports a single unit count is the dural penetrating diamond shuttle for delivery of a flexible array, which yielded acute neural recordings of 20 units over 60 channels, or an average of 0.33 units per electrode[206]. Another exciting alternative to manual durotomy is laser microablation, which has been used preceding robotic insertion of thin-film polymer probes. The longest recordings reported using this method were taken at two months post-implant, with approximately 40 percent of channels recording single-unit action potentials[198].

	This paper	Hosseini *et al.*, 2007	Fekete *et al.*, 2015	Na *et al.*, 2018
Material	silicon	silicon	silicon	UNCD
Dimensions (μm)	30 x 80	100 x 120	100 x 500	11 x 65, + 27.5 x [16 ▷ 2] (tapered vertical support)
Shape	3D-sharpened tip <3μm x 3μm	planar sharpened	chemically sharpened to a tip	planar sharpened
Insertion speed (μm/sec)	50	-	50	10
Insertion force (mN, ± s.d.)	8.9 ± 1.6	41 ± 25.5	11 ± 3	6.67 ± 2.16
Insertion site	hippocampus	somatosensory cortex	various, posterior to bregma	motor cortex
Simultaneous single units : electrodes	18:16 (chronic)	-	-	20:60 (acute)

Table 3.1: Rat trans-dural insertion forces measured here and in previously published work.

Further reduction in brain compression may be achieved by fabrication of even sharper profile tips[199], which our fabrication method can achieve with minimal adjustment. Another way to reduce brain compression could be to change insertion speed. A range of speeds, including meters-per-second scale pneumatic insertion[219], have yielded dural penetration[204] and neural recordings[197,205,216], with the optimal speed dependent on tip

shape and features in the target tissue, such as vessels[201]. A recent report indicates that slower insertion of silicon probes into brain following durotomy, over the range 2 μm/sec to 1 mm/sec, improves cell yields in acute recordings[220]. Retraction speed, too, has been varied; for example, the ballistic retraction method used in combination with needle-and-thread insertion to prevent displacement of electrodes from target depth has, as previously discussed, yielded chronic neural recordings[198]. In addition to determining the optimal insertion and retraction speeds, incorporating fast axioaxial vibration with relatively slow insertion speed could also aid in insertion with minimal brain compression[206]. Additionally, application of collagenase to the dura could make it more easily penetrable[215], though this process is slow and titration to determine an appropriate concentration presents a challenge, as the dural membrane thickness varies with target location and age[204].

Another complementary insertion strategy uses an insertion guide, as recently demonstrated for insertion of SMP devices through rat dura[207]. Inspired by the labium that guides the proboscis of the female mosquito, a guide at the insertion site increases the critical buckling load of the inserted device itself. Such a strategy could be used in combination with the sharpened shuttles validated here.

This work addresses successful dural insertion with single and double-shank devices, but the fabrication method described here could be applied to larger, multi-shank arrays for recording more neural data, over a more distributed brain volume. If the spacing between shanks is much larger than the dimpling radius, we would expect each shank to experience the same force as if it were implanted independently. However, if the spacing between shanks is small, the required insertion force for the array is expected to scale with the number of shanks, as has been shown for probes from one to ten shanks[205]. The mechanics are akin to those that protect the yogi lying on a bed of nails, across which the total force is distributed to prevent injury. One array design that could maintain the ability to penetrate

dura would have shanks of varied lengths such that, at any given time, only a subset of shanks is in the process of penetrating dura. Our measurements of brain deflection can guide the design of these arrays to determine these varied shank lengths. Although we have focused on the delivery of flexible arrays, the design and fabrication technique presented here could be used for silicon electrode arrays themselves, either for dural penetration or to reduce brain compression in clinical and non-human primate implants.

3.5 Conclusion

Recording more neurons in distributed circuits will likely lead to scientific insights that are unachievable with a smaller number of neurons[221,222]. Polymer probes are among the most promising methods for chronic, large-scale neural recordings, but their insertion through the tough protective membranes of the central nervous system is challenging and currently limits their broad use and effectiveness. Here, we have validated for chronic recordings the first dural-penetrating shuttle in combination with a modular polymer probe-based recording platform. This method shows limited brain compression and obviates the need for a durotomy in rats and other model organisms with similar dural tensile strength. Maintaining intact dura will reduce post-surgical edema, likely increasing accuracy in depth-targeting of the electrode arrays. This is critical for fixed, non-drivable arrays, particularly for targets with a small dorsoventral extent. The number and quality of single units we have recorded with this system is comparable to what we have previously recorded in the OFC using planar shuttles, inserted through pia only[212]. It is our hope that this method to more efficiently implant polymer devices for high-density, chronic neural recordings will enable experimentalists to address compelling open questions in neuroscience.

Chapter 4

Sharp-wave ripples in retrieval for memory consolidation and use

Abstract: Various cognitive functions have long been known to require the hippocampus. Recently, progress has been made in identifying the hippocampal neural activity patterns that implement these functions. One such pattern is the sharp wave-ripple (SWR), an event associated with highly synchronous neural firing in the hippocampus and modulation of neural activity in distributed brain regions. Hippocampal spiking during SWRs can represent past or potential future experience, and SWR-related interventions can alter subsequent memory performance. These findings and others suggest that SWRs support both memory consolidation and memory retrieval for processes such as decision-making. In addition, studies have identified distinct types of SWR based on representational content, behavioural state and physiological features. These various findings regarding SWRs suggest that different SWR types correspond to different cognitive functions, such as retrieval and consolidation. Here, we introduce another possibility – that a single SWR may support more than one cognitive function. Taking into account classic psychological theories and recent molecular results that suggest that retrieval and consolidation share

mechanisms, we propose that the SWR mediates the retrieval of stored representations that can be utilized immediately by downstream circuits in decision-making, planning, recollection and/or imagination while simultaneously initiating memory consolidation processes.

4.1 Introduction

It is not possible for the body to go back or to leap ahead in time, but it is possible for the mind to do so, as it can store and access information about the past to conceive of the future, and thus maintain a sense of self through time. This remarkable ability depends on memory. The most general organizing framework divides memory into three phases: its initial formation, known as encoding; its ongoing storage; and its retrieval[223–226]. These phases were originally formulated by experimental psychology with reference to animal and human behaviour[226]. Following this tradition, behavioural tests designed to isolate one or more of encoding, storage and retrieval can be used in conjunction with manipulations to identify necessary brain areas, or with measurements to identify coincident neural activity patterns.

A complementary starting point is the observation of neural activity during relatively unconstrained behaviour[227–229]. In contrast to the approach from psychology, this approach does not presuppose cognitive functions, such as consolidation or retrieval. Instead, its focus is neural activity, the features and patterns in which are used to define physiological functions that support memory-associated behaviour. The mapping between the physiological functions and cognitive functions that support memory remains unclear.

The approach from psychology has established that the hippocampus supports certain types of memory-associated behaviours. Complementing these results, the approach from

physiology has identified the hippocampal sharp wave–ripple (SWR) as a specific neural activity pattern that supports various memory and cognitive functions, as described in previous reviews[230–240]. These previous reviews establish a link between SWR neural activity and two different phases of memory, consolidation and retrieval, introducing an apparent conflict between knowledge gained from the classic psychological and physiological approaches. Are there functional subtypes of SWR that preferentially subserve retrieval or consolidation? Alternatively, might a single SWR subserve both functions? Here, our aim is to extend the conclusions of earlier work in a new synthesis focused on relating SWRs with particular physiological features, which occur in specific states, to specific memory functions.

We begin with a review of hippocampal function in memory as defined by human lesion studies. Next, we establish working definitions of memory concepts including encoding, consolidation and retrieval, and we review the dependence of these processes on the hippocampus based on work in rodents and humans. We then introduce the physiological features of the SWR, which have been studied primarily in rodents. In setting out to relate the SWR to specific memory functions, we first define a series of predictions for neural mechanisms that would support consolidation or retrieval. Each of these prediction domains is the focus of a section: behavioural state dependence, necessity in memory-dependent behaviours, representational content, and relationship to activity in distributed brain regions. In each section, we summarize what is known about the SWR with respect to the predictions for consolidation and retrieval. Based on the results of this exploration, we end with an argument for consideration of an alternative conceptual approach to memory wherein each SWR mediates retrieval in support of both the immediate use of remembered information and the gradual process of memory consolidation.

4.2 The hippocampus and memory

Damage to the hippocampus, the generative structure of the SWR, has been known for at least fifty years to result in the combination of anterograde and temporally graded retrograde amnesia[223,241]. These findings, primarily based on human lesion studies (for example, of the famous patient H.M.[242]), indicated a specific deficit in what are often characterized as 'relational memories', defined as those memories that store information about complex combinations of stimuli or states[243]. In humans, these include episodic memories, which comprise the subset of declarative memories for specific experiences[225,244]. The findings from the lesion and other studies support three main conclusions. First, the encoding of relational memories requires the hippocampus, accounting for the anterograde amnesia observed in the lesion studies. Second, the hippocampus is also required for the retrieval of these memories for some time after initial storage, accounting for the retrograde amnesia with hippocampal lesions. Third, after the initial formation of memories, during their ongoing storage, memory retrieval gradually shifts from requiring an intact hippocampus to being at least partially independent of it, accounting for the temporally graded nature of retrograde amnesia following such lesions.

During encoding, neural representations of experience must be linked together to capture information about experience as it unfolds in time[245]. Physiologically, encoding is thought to depend on hippocampal mechanisms[246–248] other than the SWR[249–252]. Encoding is distinct from other memory processes in that it happens first and only once (at most) per event or episode[225,253], although multiple encoded episodes can contribute to what is thought of as a single memory (for example, each exposure in a fear conditioning task is encoded; the resulting long-term 'fear memory' is probably the result of generalization across these experiences during consolidation[254]). By contrast, retrieval and consolidation can occur repeatedly and in alternation throughout the life of a memory[255].

The adaptive value of memory storage is in its later use, the starting point for which is memory retrieval, specifically defined as the function of accessing information stored in memory. Importantly, retrieval is distinct conceptually[256] (and probably neurally) from the conscious experience of remembering, or recollection, for which retrieval is necessary but not sufficient. In addition to conscious recollection, retrieval has many other uses. In perhaps its simplest application, retrieval of a single stimulus-response association can drive behaviour directly, as when exposure to a context associated with shock leads to freezing[254]. A more complex computation may be performed when a subject is confronted with multiple options and retrieves specific episodes of past experience for decision-making or planning. Retrieval in some form could also support imagination, which can be understood as the rearrangement or elaboration of stored information in mental simulation of future possibilities[257].

Retrieval is typically inferred from behaviour, and behavioural studies indicate that all of the aforementioned retrieval and use scenarios require the hippocampus in rodents[258], monkeys[259] and humans[260–266]. In some circumstances, the dependence of such memory retrieval-and-use behaviours on the hippocampus is time-limited[254,267,268]. The gradual shift in their dependence from the hippocampus to the neocortex, initially inferred from lesion studies, is attributed to a hypothetical process known as systems consolidation. The standard model of systems consolidation and its alternatives[269] describe a process spanning weeks to years that renders memories less liable to disruption (although they are still mutable)[255,270,271].

Importantly, systems consolidation itself is thought to depend on retrieval[226,232,256,272–274]. Mechanistically, repeatedly retrieving information that is stored in memory is hypothesized to initiate synaptic consolidation processes, including subsequent protein synthesis, that strengthen or weaken specific synapses over minutes to weeks[275]. The standard model proposes that synaptic strengthening, in particular, effects the gradual, but complete,

transfer of memories from the hippocampus to neocortex[253,276]. Alternatives to this model cite findings that some memory-dependent behaviours, particularly those requiring detailed episodic memory, are always compromised following hippocampal lesion and thus may require the hippocampus indefinitely[260,277,278]. Systems consolidation is also associated with the extraction from specific episodes of general features or rules related to stored experiences[261,279] in more semantic or schema-like representations in the neocortex[260].

Although this account of memory retrieval and consolidation has its basis in behavioural studies, the extent to which behaviour alone can provide mechanistic insight to memory is limited. A behavioural report of retrieval requires not only retrieval itself but also its successful encoding, likely consolidation, and at least one of its many possible downstream uses. For this reason, behaviour is not a highly sensitive detection method for retrieval; introspection tells us that it is possible to access information from memory without any obvious outward behavioural signs. Neither is behaviour specific to the particular use of retrieval: behavioural expression of simple associations, decision-making, imagination and planning may not be differentiable. Moreover, just as retrieval cannot be reliably inferred from behavioural output, it is not fully predictable from observable inputs: although memory retrieval can be prompted by experience with contexts or stimuli from the past[280,281], it does not always occur (often we are distracted; sometimes, we forget). Thus, we can infer from behaviour that consolidation and retrieval occur, but we cannot infer precisely when, or how. Here, the behavioural approach is complemented by the study of brain physiology, which has independently identified striking patterns of neural activity that are suggestive of memory function, such as the SWR.

4.3 Sharp wave-ripples

4.3.1 Physiology of the SWR

For at least 50 years, the extracellular local field potential (LFP) has been used to relate neural and behavioural phenomena[228,229]. In early work, large, brief deflections in the hippocampal LFP were observed during periods of rest[227,228]; this striking LFP pattern was termed large-irregular activity (LIA). The defining activity during LIA in the hippocampus is a more specific pattern known as the SWR[282], the properties of which have been studied primarily in the rodent (but see Box 4.3.1).

The sharp wave component of the SWR is an extracellularly recorded event that corresponds to the summed, synchronous depolarization of a large fraction of the neurons in the CA1 subregion of the hippocampus[228,282–285]. In permissive network states[230], CA1 activity can be driven by activity in upstream CA3 that is independent of external inputs[282,286]. Such activity can be modulated by activity in CA2[287] and the dentate gyrus[288]. The strong recurrent connectivity in CA3[289,290] is thought to allow the increased activity of relatively few pyramidal cells to spread rapidly through the region. The same activity from CA3 that excites a large subset of CA1 pyramidal cells[291] also excites interneurons, resulting in the oscillatory excitation and inhibition of interneuron-coordinated pyramidal cell ensembles that manifest as the coincident ripple, a high amplitude 150–250 Hz oscillation[228,285,292–296].

The presence of high power in the ripple band is itself often used as a marker for SWRs (usually using a threshold between 3 and 9 standard deviations above the mean)[297,298]. The distribution of ripple band power is not, however, bimodal. Rather, it is approximately log-normal with a long tail towards high values[298]. The application of a threshold to ripple power, therefore, should not be understood as discriminating SWR and non-SWR events

with perfect accuracy. Nonetheless, the properties of spiking and LFP activity observed during SWRs are consistent across a variety of studies that use different thresholds[230].

Based on features that we review below, SWR-associated neural activity has been proposed to support memory consolidation, retrieval, planning and imagination[230,231,237–239,299]. In parallel, different types of SWRs have been identified based on representational and physiological characteristics. Whether there is a simple mapping from memory function to SWR type remains unclear, however. This question is particularly important because of its implications for our understanding of the relationship between retrieval and consolidation. Consolidation requires retrieval; does retrieval necessarily lead to consolidation? An absence of SWR functional types would suggest that, for this particular candidate memory mechanism, retrieval and consolidation may be effected together.

4.3.2 SWR occurrence depends on experience and behavioural state

The SWR is a physiological event of sub-second duration. To the extent that a SWR supports systems consolidation, it would likely be by initiation of slower synaptic consolidation processes[323]. In this Review, we differentiate these hypothetical consolidation-promoting events (consolidation events) from consolidation itself, the process of strengthening a memory that is thought to correspond to synaptic changes. Consolidation events are expected to occur with greater strength or frequency following any experience that would be adaptive to remember, such as a novel, rewarding, punishing or otherwise instructive experience[324]. Because sleep has been shown to have a memory-strengthening effect[279], we expect consolidation events to occur during sleep. It is possible, however, that they also occur during wake.

Because consolidation is thought to require retrieval, any period when consolidation events

The majority of sharp wave–ripple (SWR) and replay studies have been performed in mice or rats, but SWRs have also been described in vivo in cats[300,301], bats[302], rabbits[303,304], monkeys[305–309] and humans[310–313]. SWR activity has also been described in the Australian bearded dragon (*Pogona vitticeps*), but in a region that is not considered to be a hippocampal analogue[314], and whole zebrafish brains in vitro show SWR-like activity[315]. Although sequential replay has only been observed in rodents, the relationship in rodents between SWRs and memory-dependent behaviour has also been reported in rabbits[303,304], non-human primates[306] and humans[311].

Consistent with rodent findings[230], human SWRs are most frequent during slow-wave sleep and immobility[316,317] and are correlated with widespread changes in activity throughout the brain. In macaque monkeys, a series of fMRI studies identified activation of the neocortex and inhibition of subcortical structures at the time of SWRs[305,318], as well as elevated activity following SWRs of the default mode network[319], which in humans is linked to memory processes, including imagination and prospection[262,319]. Differences between SWRs in rodents and in primates include a lower SWR rate in humans and monkeys[308,311] (possibly due to challenges inherent in primate SWR detection) and, in monkeys, the observation of SWRs during visual search, which is considered an analogue of active exploration[306]. This is at odds with thinking that, in rodents, exploration is associated with a theta rhythm and memory encoding[229,320] more than with the SWR and its proposed memory functions.

A cross-species comparison of memory ability and features of the SWR, including state dependence and replay content, would potentially be informative for determining SWR function. Such a comparison will require further testing of the memory abilities of those species with SWRs[321], and causal studies of the SWR in species other than rodents[322].

Box 4.3.1: Sharp wave–ripples across species.

are expected based on state or behaviour is also expected to contain retrieval events; like consolidation events, retrieval events are expected to occur following novelty, reward and punishment, as well as in sleep and probably also in wake. These timing-based predictions therefore cannot discriminate SWR function in consolidation from function in the retrieval that supports it.

For retrieval in uses other than consolidation, similar dependencies on behavioural state are expected. Because retrieval is necessary for processes that occur in the awake state, such as

decision-making, retrieval events are expected to occur then. Retrieval events for decision-making might also be expected to decline with novelty, occurring more often in situations early in learning before general rules have been learned[325,326], as this is when specific memories may be used to guide decisions. Likewise, the use of retrieval in recollection may be more frequent for events following reward[327]. Retrieval events may also occur during sleep, as when elements of waking experiences recur in dreams, and thus would support the previously reported phenomenon of sleep insight[328]. These possibilities demonstrate that predictions based on behavioural state or experience alone cannot definitively differentiate a mechanism that supports retrieval for consolidation from retrieval for other uses – just as such predictions cannot distinguish retrieval events in consolidation from consolidation events themselves. Furthermore, because consolidation and retrieval events are expected to occur in the same behavioural states, a single mechanism could exist to support both simultaneously.

Work in rodents has demonstrated that SWRs occur most frequently during slow-wave sleep, least frequently during running, and at an intermediate rate, occurring once every few seconds, during periods of quiet rest[298,329]. This pattern of occurrence is regulated by modulatory factors[330] including cholinergic tone[232,331], which tends to be higher during movement[332]. Cholinergic modulation has also been speculated to explain a recent report that SWR occurrence is entrained by breathing[333]. SWRs typically occur more frequently during and after novel experiences[298,334–336], although the SWR rate has been reported in some tasks to increase over the course of multiple traversals of a familiar path[337]. An increase in SWR rate is also seen immediately after receipt of a reward, particularly if it occurs in an unfamiliar location[338] (Fig. 4.1).

Together, these findings indicate that the SWR rate is at its highest in the contexts of novelty and reward, consistent with functions in both consolidation and retrieval. The increase in SWR rate immediately following reward is consistent with theoretical

predictions for a consolidation mechanism that would be particularly effective in linking extended actions to their outcomes, a specialized information storage problem known as credit assignment[338,339]. However, a higher SWR rate after reward could also represent retrieval for recollection, and a higher SWR rate during novel experience could also indicate retrieval for decision-making.

4.3.3 SWRs are necessary for memory performance and stable representations

If the SWR supports consolidation, either as a retrieval mechanism or in some other capacity, it is expected to be necessary for memory-dependent behaviours and changes in synaptic strength, the established molecular correlates of learning. Under the standard model, a mechanism for systems consolidation is also expected to be necessary for the renormalization of synapses in the hippocampus and, in parallel, the alteration and stabilization of synapses in extra-hippocampal structures, particularly the neocortex[323,340,341]. Alternatively, if the SWR is a mechanism for memory retrieval for use in functions other than consolidation, it is likewise expected to be necessary for memory-dependent behaviours but not plasticity.

Studies that have interrupted or disrupted the structure of SWRs during sleep have demonstrated their necessity for memory-dependent behaviours and have been interpreted as evidence of a consolidation function. The first of these studies truncated SWRs by electrical stimulation of CA3–CA3 connections in the ventral hippocampal commissure (VHC). Across many days of learning, SWRs were truncated during hour-long sleep or rest periods following experience. This resulted in slower learning in hippocampus-dependent spatial memory tasks[342,343]. SWR disruption by other methods, including suppression of CA3 output to CA1[344], and optogenetic activation of the locus coeruleus[345] or median

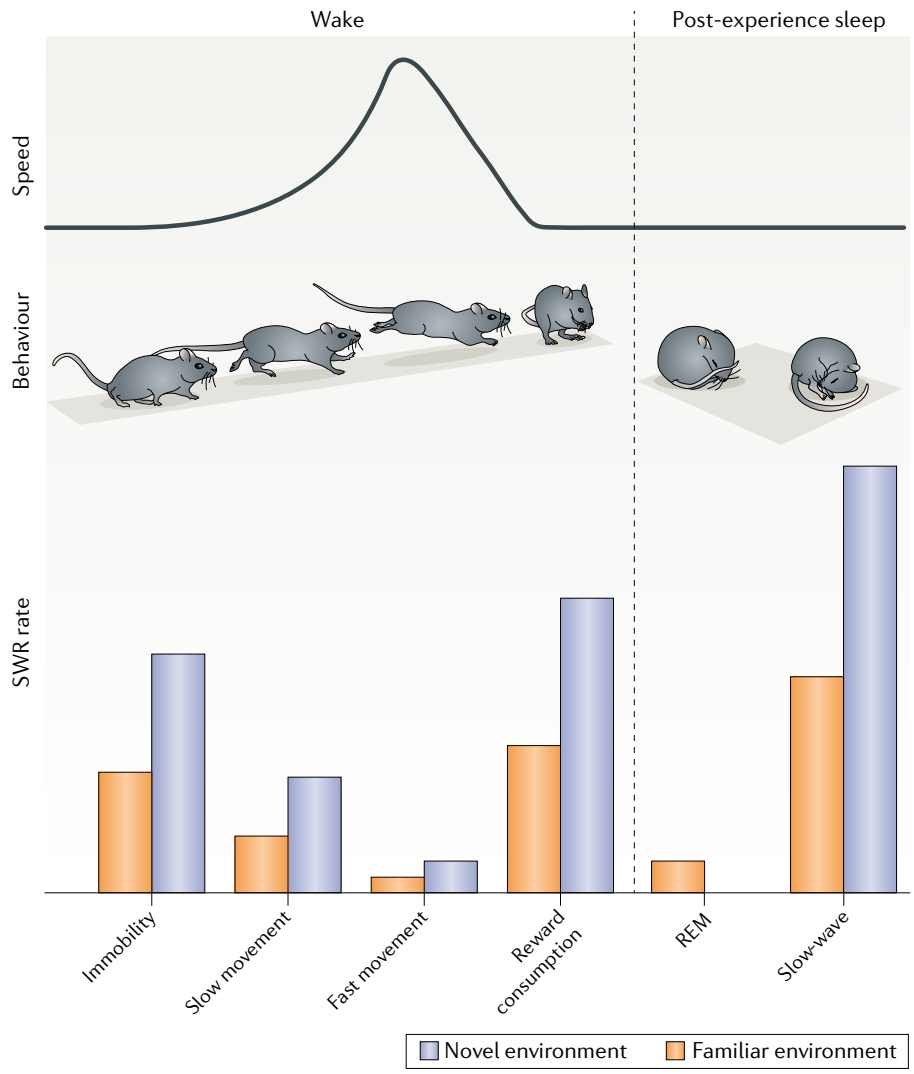

Figure 4.1: Schematic of sharp wave–ripple rate across brain state and with movement speed. In the awake state, the sharp wave–ripple (SWR) rate varies as a function of movement speed, novelty and receipt of reward. SWRs are most common during periods of immobility and become increasingly rare at higher movement speeds. Their rate is higher across all speeds in novel versus familiar environments, and in both novel and familiar environments their rate is highest following receipt of reward. During sleep, SWRs occur only rarely during rapid eye movement (REM) sleep and occur most often during slow-wave sleep. The SWR rate during slow-wave sleep (in a familiar sleep box) following exploration of a novel environment is higher than it is following exploration of a familiar environment. The SWR rate in REM sleep has not been studied following experience in a novel environment, therefore it is omitted from the figure. These patterns of modulation are consistent with an increased SWR rate during and after learning and indicate that SWR-mediated retrieval is utilized primarily at lower speeds.

raphe[330] during post-behaviour sleep, has a similar effect. Interestingly, when SWRs are disrupted in sleep after learning, but not after a random foraging task, there is a subsequent increase in their rate, suggesting the existence of homeostatic-like control of SWRs that is set based on learning[346]. A recent gain-of-function study also demonstrated that electrical stimulation of the medial prefrontal cortex (mPFC) immediately after each SWR during sleep led to an increase in coordinated activity between mPFC and hippocampus, as well as an improvement in memory for a briefly experienced set of objects in a context[347].

Disruption of SWRs during awake behaviour also impairs learning and performance in spatial memory tasks[303,348]. SWRs were demonstrated to be necessary for a normal rate of initial learning of a spatial alternation rule and, in a separate group of animals already trained in the task, for continued performance[348,349]. Similarly, delivery of a strong light stimulus after SWRs disrupted learning in a trace eye blink conditioning task, suggesting that the period extending hundreds of milliseconds following an SWR is important for memory processes[303]. In further support of this hypothesis, gain-of-function experiments showed that presentation of conditioned stimulus – unconditioned stimulus pairings specifically following SWR events led to an acceleration in learning (but also slowed extinction)[304]. The disruption of SWRs that is observed in both sleep and wake in models of diseases with memory symptoms also suggests SWRs contribute to memory (Box 4.3.2).

There is complementary evidence that SWRs contribute to stabilization of representations that are formed during experience. In awake behaving rodents, SWRs can contribute to stabilization of place fields, the spatial representations typical of principal cells in the hippocampus ('place cells')[350,351]. During awake behaviour for mice performing a spatial memory task, optogenetic silencing of principal neurons in CA1 during SWRs reduced the stability of their place representations[352], and following this period of silencing, active hippocampal neurons were more likely to have altered place fields[350] when the mice were re-exposed to the environment. Similar manipulations during sleep have suggested that the

subset of place cell ensembles that are not yet fully stable at the end of an initial novel experience can be destabilized by optogenetic suppression of neural activity during SWRs[353]. By contrast, ensembles that were stable by the end of a novel experience were not affected, perhaps explaining why another group found that optogenetic suppression of CA1 principal neurons during sleep SWRs had no effect on spatial representations[354].

The stabilization of these representations probably results from changes in synaptic strength, but studies have differed in their reported effects of SWRs on synaptic plasticity in the hippocampus. SWR-like activity in vitro can promote intrahippocampal synaptic potentiation[355,356], but recent findings have indicated that sleep SWRs can also contribute to the widespread downscaling of hippocampal synapses[357], which substantiates a previously developed model[358] predicting that SWRs induce downscaling of intrahippocampal synapses and potentiation of extrahippocampal synapses. The possibility that SWRs could drive local synaptic renormalization (perhaps to reset hippocampal synapses so that new learning could occur[357–359] is also consistent with the clearance of a hippocampal memory trace and simultaneous consolidation in neocortex predicted by the standard model of systems consolidation. Whether awake SWRs can have the same effect is unknown, as wake and sleep are different neuromodulatory states and SWRs probably differ, and/or have different effects, in each of these states[232,240]. For both sleep and wake, it is unknown whether some synapses might be maintained or strengthened while others are weakened.

The effects of SWR disruption on behaviour and hippocampal place representation constitute strong support for a memory function of SWRs in both wake and sleep. SWR disruption and augmentation effects on learning and synaptic plasticity indicate a role in a consolidation function in particular, for at least a subset of SWRs. It is not possible, however, to rule out the possibility that these or a different subpopulation of SWRs have an additional function in retrieval for other uses. For example, to date, published awake

SWR disruption experiments have interrupted all SWRs. The detriment in performance resulting from this disruption could therefore be explained instead by an effect on retrieval in support of decision making, particularly for those studies in which SWR disruption occurred after rule learning was complete and resulted in a decline in otherwise stable performance[348]. Thus, although SWR disruption studies indicate that SWRs support memory in both sleep and wake, these studies do not specifically indicate retrieval for consolidation versus other use, or a function other than retrieval in support of consolidation.

An urgent goal of neuroscience is to understand diseases of the human brain, many of which share the symptom of debilitating memory loss. One explanation for this symptomatic overlap is that myriad cellular- and molecular-level changes can cause dysfunction in the network-level activity patterns that support memory[360]. Identifying such emergent activity patterns presents the possibility of a therapeutic shortcut that is broadly effective: if activity patterns can be restored, even if the circuitry that naturally supports them cannot, memory could be restored.

The sharp wave–ripple (SWR) is a reasonable candidate for that approach. The first report of disrupted SWRs in disease was in human epilepsy[230,313,361]. Disrupted SWRs have since been reported in animal models of epilepsy[291,362], Alzheimer disease[363], dementia[364], schizophrenia[365] and normal ageing[366,367]. Recent studies have demonstrated up- or down-regulation of SWRs through neuromodulatory[330,331,345,368] or other control[369], suggesting possible therapeutic strategies.

Given the diversity of findings in disease models for simple metrics such as the SWR rate[370–372], however, success with this approach is likely to require identification of more specific patterns of disordered activity[291,365,367,373]. As an example, genetic and developmental mouse models of schizophrenia show disrupted sequential spiking activity during the SWR[366,373], or decreased coordination of SWRs with cortical sleep spindles[373]. Also, a study of a knock-in mouse expressing the human APOE4 variant known to cause late-onset disease found learning and memory deficits that were associated with a deficiency in the normal slow gamma[364] rhythm known to organize the activity of CA1 cells during the SWR[374,375]. A subsequent study found that rescuing the gamma deficit reduced the level of Alzheimer disease-related amyloid-β isoforms[376], indicating that addressing network-level dysfunction could also resolve other disease symptoms.

Box 4.3.2: Sharp wave–ripples in disease and ageing.

4.3.4 Hippocampal spiking during an SWR can represent previous experience

Mechanistically, the most readily recognizable form that retrieval could take is the precise repetition of the activity pattern that was observed during the experience itself[274]. Indeed, retrieval is thought to occur by reactivation of neural activity patterns in the hippocampus that correspond to those that occurred during a previous experience, a possibility foundational to modern engram theory[280,377]. This activity, in turn, is thought to reactivate hippocampal–cortical and subcortical activity patterns to represent the multisensory features of a memory. In synaptic and systems consolidation, it is likewise the repeated reactivation of the hippocampal–cortical patterns stored during an experience[253,323] that is hypothesized to create and strengthen the intrahippocampal and hippocampo–cortical synapses that constitute the memory trace. This correspondence further highlights that the hypothetical consolidation event and the retrieval event are effectively identical, and that the requirement for retrieval in consolidation is satisfied by the repetition of such events. The evidence indicates that SWR activity may constitute such an event.

During SWRs, more than during any other period of activity, sequences of neural spiking activity recapitulate those seen during prior experience. These striking reactivations were originally observed at the level of single cells[378] and cell pairs[379], but also occur at the level of ensembles[380,381]. Individual reactivation events can represent either specific locations in space[381,382] or can 'replay' long sequences of place cell activity that recapitulate entire spatial trajectories[299,339,383–387]. Replay events can also represent long, extended experiences, with events spanning multiple SWRs[387,388]. Such replay of past experience is seen in a subset of SWRs in both waking and sleep, although replay in sleep is a less accurate recapitulation of awake activity patterns[385,389]. The representation of past experience by replay activity suggests that it is a critical component of the SWR contribution to memory, and that replay variants may correspond to different SWR

memory functions.

Replay events can reactivate representations corresponding to either local trajectories beginning at the animal's current location or remote trajectories, defined as those that begin far from the animal?s location[385,386] or in an entirely different environment[383,385,390,391]. When spatial sequences correspond to the current environment, ∼80% of identified replay events begin with representations of the animal's current position[339,384,385,387]. These events then extend towards locations farther from the animal, with a bias for representation of future goal locations[299]. Remote replay occurs in both wake and sleep[385,389]. Behavioural state influences replay content, in that events that occur in close temporal proximity to movement more often originate at the animal?s current location (that is, they are more often local)[385,392].

Replay events also vary in the represented direction of movement (Fig 4.2). In linear environments, hippocampal place cells gradually develop directionally biased firing patterns, with a higher firing rate when the animal traverses the place field in one direction of motion than in the reverse[393,394]. These biases in firing make it possible to use ensemble spiking during replay events to infer not only the represented location of the animal but also the direction of movement. These analyses have revealed that awake replay can occur in both the same direction as the original traversal (forward replay) and the opposite direction, which may never have occurred during behaviour (reverse replay)[339,384,387].

A functional difference for forward and reverse replay in wake is suggested by their independent modulation: increases in reward magnitude increase the rate of local, reverse replay events and decreases in reward magnitude reduce it, with no effect on forward replay events[395]. This result is consistent with the hypothesis that local, reverse replay following an outcome can function as a consolidation mechanism that is specialized to the problem of credit assignment, wherein outcomes must be linked to the actions that led to them. During replay, earlier reactivation of cells with place representations that are physically

closer to the outcome location is hypothesized to facilitate the strengthening of synapses with coactive cells representing reward or punishment, while the specificity for movement towards the outcome location preserves directionality[339]. Consistent with credit assignment or consolidation in general, the intensity of reactivation during SWRs can be related to the reorganization of spatial representations and to memory for previously rewarded locations[382].

In contrast to reverse replay, local replay events in the forward direction during wake have been correlated with subsequent behaviour, suggesting that they function in retrieval of previous experiences that occurred in the same context, for immediate decision-making or planning[299]. These events could also function in retrieval for conscious recollection, or 'mental time travel', the ability to be in one location and simultaneously remember a past experience that may have occurred in another[244,299,396,397] (but see[230]). Consistent with this possibility, replays are enriched for representation of immediate future choices[299], and more intense activity during these events can be predictive of a subsequent correct choice[398]. Further, the awake replay of an upcoming location associated with shock is predictive of a subsequent change in movement direction[399].

Although reverse and forward replay events are both common during wake, reverse events are seen only infrequently during sleep[383,400]. The consolidation function of sleep is therefore expected to be carried out by forward replay events. The specific nature of this hypothesized consolidation remains unclear, however. Consistent with standard consolidation, the intensity of reactivation during sleep SWRs is related to the reorganization of spatial representations and to memory for previously rewarded locations[381]. However, the fidelity of replay events during sleep is, on average, much lower than of those in wake[385], prompting speculation that sleep events represent sequences including elements of multiple experiences to support consolidation, in the form of generalization across them[240,385].

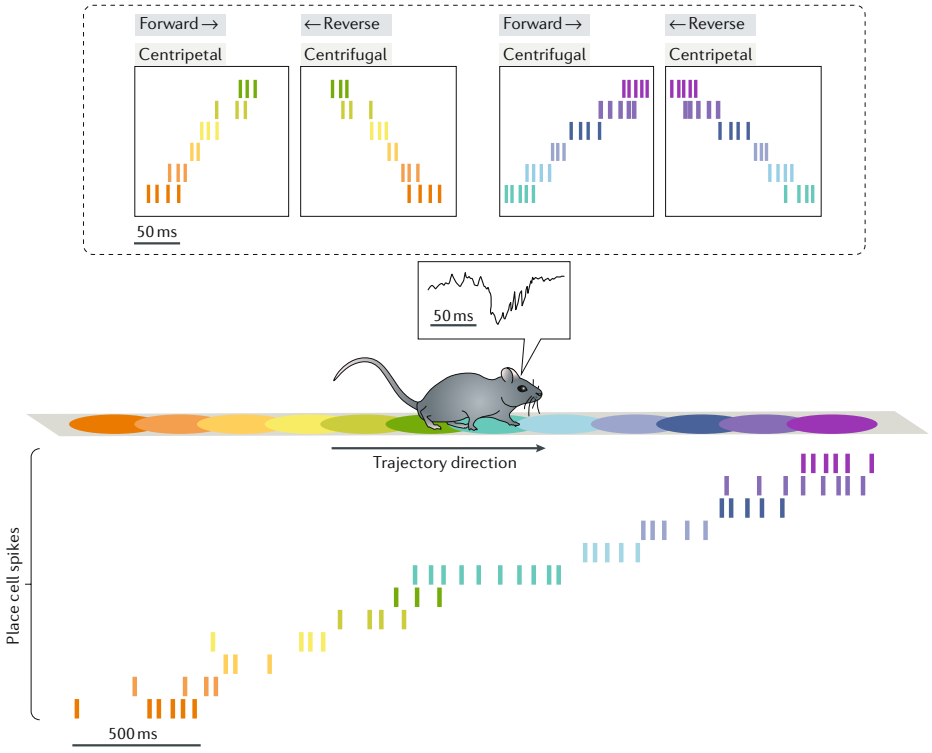

Figure 4.2: Schematic of possible local replays: in the forward and reverse directions, centrifugally and centripetally. Place cells increase their firing rates as a rat traverses the cells' respective place fields on a linear track from left to right (orange to purple). When the rat pauses, immobile, at the center of the track, a sharp wave-ripple (SWR) occurs. Place cell activity during the SWR recapitulates recent experience, firing in the same order on a compressed timescale. Relative to the order of place cell firing during actual experience, these sequences can represent trajectories that are: forward and centripetal (toward the rat); reverse and centrifugal (away from the rat); forward and centrifugal; or reverse and centripetal. Forward sequences are indicated by arrows beside the label 'forward' that are oriented in the same direction as the trajectory arrow (in this case rightward); reverse sequences, by contrast, are indicated by arrows oriented toward the rat. For centrifugal sequences, these arrows are oriented away from the rat; for centripetal sequences, the arrows are oriented toward the rat. If any of these sequences occurred prior to the rat actually traversing that track segment (as occurs more often in a less constrained environment with more path options), they would represent novel sequences. A second set of place cells with overlapping place fields, but that are preferentially active when the rat moves in the opposite direction (right to left), would participate in the same four replay types. If the rat ran on the track (middle panel), then the replays (top row) occurred when the rat was no longer on the track, they would be classified as remote.

4.3.5 Spiking during SWRs can represent actual or alternative future actions

Mechanistically, the imagination function thought to be supported by the hippocampus[401] could begin with a retrieval process wherein stored information is accessed, and proceed with the rearrangement of that information in novel combinations. These imagined scenarios might correspond to potential future choices, such as when planning new routes to a goal location. As in veridical retrieval, this neural activity would likewise be expected to reactivate patterns of activity brain-wide.

Consistent with these predictions for an imagination mechanism, SWRs can contain sequences corresponding to trajectory events that represent novel paths that have not been previously traversed by the animal[299,386]. Such prospective trajectory events suggest a function beyond veridical retrieval for decision-making or consolidation, wherein retrieval and rearrangement of previously stored representations supports a process like imagination. In addition to novel sequences pertaining to the local environment, there have been reports of what can be understood as remote prospective events, or 'preplay', in which activity sequences that will occur in a subsequent novel experience are seen during sleep before the experience[402–404], although this remains a topic of debate in the field[405]. We return to the issue of imagination in the final section.

4.3.6 SWR activity engages extra-hippocampal areas

A key prediction shared by mechanisms for consolidation and for retrieval in other uses is the modulation of neural activity in distributed brain regions. In systems consolidation, the reactivation of hippocampal–cortical patterns for synaptic stabilization necessarily requires correlated activity. Retrieval is likewise thought to depend on hippocampal

coordination of cortical (and likely subcortical) networks to represent the various aspects of experience, including different sensory modalities and extracted features[406], a possibility consistent with functional MRI (fMRI) studies of human retrieval[407]. Both retrieval and consolidation can be influenced by external stimuli, including during sleep[279]. Mechanistically, this is thought to depend on inputs to the hippocampus, such as those from the entorhinal cortex, that can bias the local network to reactivate specific representations (potentially by a pattern-completion process)[341,408]. Distributed activity in consolidation is expected to result in strengthened synapses, whereas such activity in retrieval for other uses is not expected to have such an effect. With respect to SWRs, for their typology based on replay content to be meaningful, they should be associated with different downstream effects, and only a subset of SWRs (for example, those with reverse replays in wake) should be associated with synaptic consolidation. For an SWR typology based on replay content to be meaningful, different subsets of SWRs should be associated with different downstream effects, with only one subset of SWRs (for example, those with reverse replays in wake) associated with synaptic consolidation.

Consistent with these predictions, changes in neural activity throughout the brain have been observed during SWRs. Electrophysiological studies have identified coordination between hippocampal SWRs and neocortical sleep spindle events[310,409] as well as coordinated hippocampal and extrahippocampal modulation of spiking activity at the time of SWRs in the dentate gyrus[288,410,411], deep (but not superficial[412,413]) entorhinal cortex (EC)[414,415], orbitofrontal cortex[416], mPFC[389,409,417,418], anterior cingulate cortex (ACC)[382,419,420], auditory cortex[421], parietal cortex[422], ventral striatum[423,424] and ventral tegmental area (VTA)[425].

In some extra-hippocampal structures, stronger coactivity with hippocampal cells has been reported in wake than in sleep, despite the association of SWRs with sleep spindles[389]. Patterns of hippocampal–PFC coactivity seen during behaviour were more strongly

re-expressed during awake than sleep SWRs[389], mirroring the finding of more veridical hippocampal replay in awake SWRs[385]. Similarly, SWR-associated reactivation of VTA neurons was present during both wake and sleep[426], although it was more prevalent during wake[425].

A recent study also found learning-related coordination of SWRs and high-frequency activity specifically in association cortices during sleep[427]. Additional studies have identified patterns of coordinated hippocampal–cortical activity that are consistent with SWR activity, but these studies did not detect SWR events[390,428]. All of these studies report that hippocampal and cortical or sub-cortical neurons that fired together during waking experience also fired together during subsequent SWRs or reactivation events, as expected in retrieval for consolidation (or for any other application). Complementing this result, combined hippocampal electrophysiological recordings and fMRI in primates revealed changes in blood oxygen level-dependent (BOLD) activity around the time of SWRs across virtually all cortical areas and many subcortical areas[305].

The specific subset of cells engaged in SWR-coordinated activity also changes with learning. Early in learning, hippocampal–PFC coactivity patterns during SWRs are correlated with coactivity during behaviour, suggesting a simple Hebbian association mechanism[418]. Later, once the environment and task are familiar, this relationship becomes much weaker[389], with a subset of PFC neurons that encode general features of the environment and task showing more specific engagement during SWRs[429]. The resulting hippocampal–PFC coactivity preferentially links hippocampal activity patterns representing specific locations with cortical activity patterns that generalize across a set of locations.

Studies of the precise timing of cortical activity relative to hippocampal activity during a SWR have suggested that cortical activity could influence subsequent hippocampal SWR activity. During sleep, hippocampal SWRs often occur immediately after transitions from cortical down states to cortical up states[430], and increases in hippocampal spiking can

occur as much as ∼200 milliseconds after increases in spiking in sensory cortical areas[390,421,431,432]. By contrast, although some PFC and ACC neurons appear to increase in activity before SWRs in both wake[382] and sleep[389,420], SWR-related PFC activity has most often been reported to follow, rather than precede, SWR-related hippocampal activity during both wake[418] and sleep[305,417,433]. The overall temporal offsets for PFC are small (∼15 milliseconds), however, compared with those for sensory cortex (∼200 milliseconds).

One possible explanation for these inconsistent reports arises from the recent discovery of a cortical–hippocampal–cortical loop of information transmission. It was previously known that hippocampal activity could be biased by sound presentations during sleep[434], but the mechanism was unknown. Recently, it was discovered that in auditory cortex, patterns of activity before SWRs can be used to predict subsequent hippocampal SWR activity, and that hippocampal SWR activity in turn predicts post-SWR cortical activity[421]. These findings suggest a cortical–hippocampal–cortical loop of information transmission around the time of sleep SWRs, in which cortical activity can cue hippocampal SWR activity, which in turn drives broad activation of cortical areas. This loop has been proposed to support cortical consolidation[236]. It is also possible that such a loop could explain differences in neocortical activity relative to SWRs, if reactivation of cortex by the hippocampus recruits association areas that were not initially active (a possibility reminiscent of one element of hippocampal indexing theory[435]).

In the same regions where spiking rates increase in coordination with SWRs, separate populations of cells simultaneously decrease firing rate[389,418,430,436]. These decreases occur specifically for cells that are most active immediately before SWRs, and in both the hippocampus and the PFC these neurons encode information related to the animal?s current location. This is consistent with the expectation that, during retrieval, processing of current sensory information might be suppressed.

4.4 Outstanding issues

Evidence supports the current view that the SWR is a principal network-level mechanism for reinstatement of stored representations to support both awake and sleep memory processes[229,273,437]. However, in the context of awake retrieval, issues related to its timing indicate that it is not the sole mechanism. SWRs take up only a relatively small fraction of the total time of awake experience (at most a few percent)[298], and occur very infrequently during movement. There is no evidence that memory and cognition are limited to periods of stillness, and thus awake retrieval probably engages additional mechanisms[229]. The ordered activations of place cells also observed during theta ('theta sequences'[438]) are an excellent candidate activity pattern that could support awake retrieval and use processes, and indeed the content of these sequences can predict subsequent choice[439,440]. This is also consistent with observations of intact SWR sequences following manipulations that impair theta sequences and performance in a memory task[441].

Even during immobility, the SWR may not be strictly necessary for learning and memory. Animals can still learn a spatial task when awake SWRs are disrupted, albeit at a slower pace[348] (but see[303]). Similarly, the observation of correct trials during which no SWRs occur suggests that the SWR is not required by trial-by-trial decision-making[306]. These results could be explained by incomplete SWR interruption or detection[442], however, or by compensation for an absence of SWRs by other mechanisms[325,326].

The evocative pairing of replay activity with the plasticity potential of the SWR also makes it easy to leap to the conclusion that that the SWR is a privileged period for replay, and that every SWR contains a single replay event that is retrieved for decision-making and/or systems consolidation. Although many SWRs are associated with reactivation of activity patterns representing past experience, studies typically report that only 10–40% meet statistical criteria for replay[299,339,383–387], and that reactivation events can occur outside of

identified SWRs[299]. It is also possible that some events cannot be decoded spatially because they correspond to trajectories that were not measured by the experimenter or are not spatial in nature[443,444]. Regardless, it seems likely that many SWRs do not contain spiking sequences consistent with a single, discrete past or potential future experience.

Another possibility is that SWRs without detectable replay content correspond to retrieval events for content with representations that have mutated over time to the point they no longer match the activity that was recorded during the original experience. Indeed, place representations can change over days[445]. Such replays would be undetectable but still functional, and could explain how replay, which has been observed to decline significantly in rate during the 18 hours following experience[385], could still contribute to a systems consolidation process extending weeks to years after an experience.

Finally, although there is evidence of the brain-wide coordination of activity expected during retrieval and consolidation at the time of the SWR, there is no evidence to suggest that different replay types correspond to different patterns of neural activity in areas downstream of the hippocampus. Recently, work in non-human primates has identified four SWR event subtypes, defined by the timing of the ripple relative to the sharp wave, that are associated with different patterns of cortical and subcortical activity[307]. A rodent study has identified subsets of SWR in which hippocampal–cortical patterns corresponding to immobility versus immobility are reactivated separately[382]. Regarding the question of consolidation versus retrieval, however, we are aware of no strong evidence to suggest that any specific subtype of SWR is exclusively associated with the plasticity expected in consolidation, or with behaviour indicating the planning and decision-making associated with awake use of retrieved memories.

4.5 SWRs, retrieval and consolidation

In modern work, memory retrieval and consolidation are often conceptualized as distinct processes that occur on different timescales: a retrieval event is understood to take place in milliseconds[274,446,447], whereas consolidation may take hours to years (but see[448]); furthermore, the effect of retrieval is thought to be a transient change in activity for immediate use, whereas consolidation is thought to effect lasting change. However, the two processes share a fundamental similarity that has long been hypothesized in experimental psychology[224,256,449]: the neural activity representing a previous experience that is reinstated in retrieval is thought to be reinstated repeatedly in consolidation to strengthen associations and synapses (for more on the essential function of repetition in biological and artificial learning systems, see Box 4.5.1).

SWR-associated replay is a prime candidate for that reinstatement[232,272,273]; the basic function of the SWR seems to be retrieval. Behavioural results establish that such retrieval can be used to support consolidation in sleep[342,343,381] and in wake[237,348,352,386], and to support other uses including decision-making in wake[299,348,399]. Replay during sleep could support a more creative variant of planning, like imagination, that might activate elements of multiple experiences in novel conjunctions. Previous work has not clearly addressed whether these varied retrieval-based functions are achieved by distinct functional types of SWR that would, for example, subserve exclusively retrieval for consolidation versus retrieval for decision-making[237,238,240,450]. Although such type splitting is possible given the diversity in representational content and physiological properties of the SWR, as well as the behavioural states in which they occur[230,287,307,382,387,451–453], we conclude that there is not yet evidence to suggest a correspondence between SWR types and different cognitive functions. In particular, there is no evidence to suggest that some SWRs are better suited to consolidation than others. Indeed, slice electrophysiology and modelling indicate that SWRs with forward and reverse replays may equally support synaptic strengthening or

downscaling and could therefore both subserve consolidation[231,234,355,454].

We therefore propose the working hypothesis that each SWR simultaneously retrieves a memory and, in doing so, drives that memory's consolidation (Fig. 4.3). This is consistent with recent work indicating that memory retrieval shares many molecular mechanisms with consolidation, including protein synthesis and NMDA receptor-mediated AMPA receptor trafficking[274,455,456]. Furthermore, the proposal that the two processes typically occur together to support normal behaviour might explain findings in the reconsolidation field that retrieval with blockade of synaptic strengthening (that is, consolidation) degrades a memory[457–462]. Such a relationship would also guarantee the adaptive solution of strengthening and transformation through consolidation of regularly retrieved and used memories. This notion is similar to the old idea that each act of retrieval forms a new, composite memory for the retrieved information in its current context, to extract generalities[226,281].

What, then, would happen when SWRs include activity representing partially or fully novel sequences? Specifically, is there a mechanism by which these are prevented from being consolidated? Alternatively, it is possible that they are consolidated, but in such a way as to be differentiable from memories of experience[463,464]. Such consolidation of novel (but realistic) sequences is a potentially adaptive method of creating and maintaining a cognitive map of the environment[228,231,325,465].

Learning is supported by the repetition of experience, which has many names, including practice, study and training. Internally driven replay of neural activity representing an experience – a form of training without repetition of experience itself – can efficiently promote further learning, because it enables easily adjusting the amount and the timing of training.

Such flexibility is an important point of similarity between artificial and biological learning systems; the brain's capacity for fast, flexible experience replay enables it to keep up with artificial systems, and justifies the sharing of solutions between them. For instance, a foundational problem faced by any learning system, biological or artificial, is that of balancing stability and plasticity[466,467]. An early description of this problem was given in the context of connectionist systems, now known as neural networks, in which learned information is stored in the form of altered synaptic weights, as in the nervous system[468]. In a highly stable system, a single instance of repetition will not drastically alter synaptic weights; here new learning requires multiple exposures, and information cannot be acquired quickly. A more plastic system is the opposite, and thus risks overweighting recent experience to result in 'catastrophic forgetting', the total erasure of previously stored information.

Presented as the brain's solution to this problem, the complementary learning systems (CLS) model[469] described the updating of a stable neocortex by a plastic hippocampus through 'interleaved learning', wherein new information is incorporated gradually to existing knowledge through spaced repetition, now attributed mechanistically to hippocampal SWR-replay[470]. Similar dual-network architectures were developed in artificial systems[466]. Hippocampal replay has since inspired other machine learning algorithms[471], including the successful 'prioritized experience replay' used in training the Deep Q Network, in which rewarded events replay more often[472,473].

In general, machine learning relies on repetition in the form of exposure to many events and, in some cases, multiple passes over the full training set, possibly with prioritization[474,475]. These two forms of repetition are reminiscent of the brain's capacity to learn both from multiple examples of direct experience and from their replay, potentially complementary abilities that could provide insight into learning systems in general.

Box 4.5.1: Repetition, sharp wave–ripple replay and machine learning.

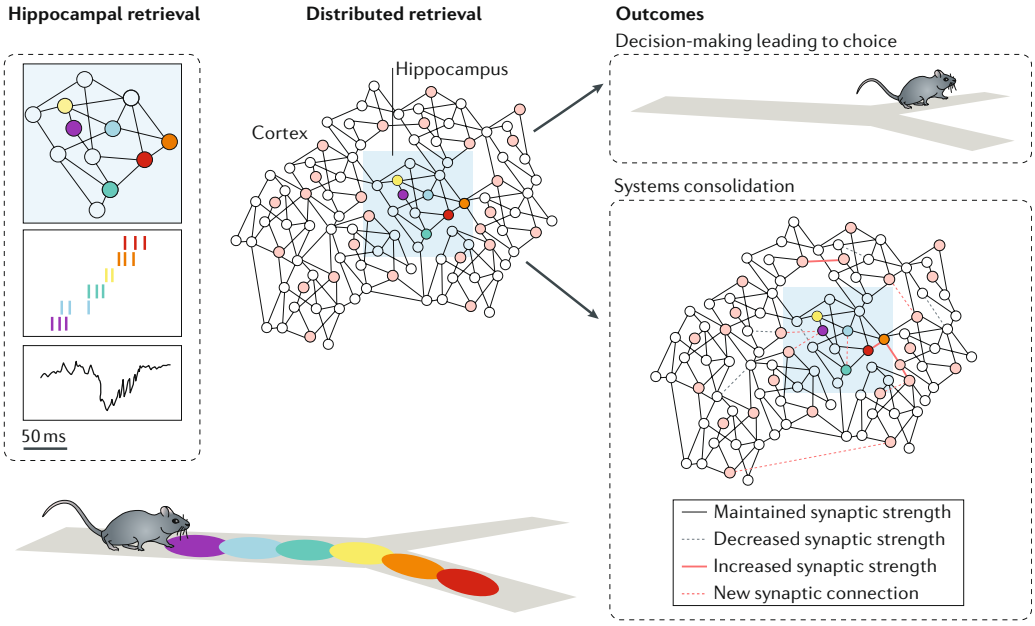

Figure 4.3: Hypothesized function for sharp wave–ripples in retrieval of information from memory for immediate use and consolidation. We hypothesize that retrieval, as it occurs here as the rat pauses on approach to a choice point, can be mediated by the sharp wave–ripple (SWR) (left panel, lower box) during which the ordered reactivation of place cells can represent trajectories previously experienced by the animal. Here, a centrifugal forward replay composed of activity from place cells with fields shown in purple to red, is depicted (left panel, middle box). The effect of replay activity in the hippocampus (left panel, top box; coloured nodes represent place cells that spiked during the replay) is the reactivation of activity in distributed networks outside the hippocampus (for example, the cortex; middle; red nodes indicate active neurons). The immediate effect of this on behaviour (top right) is to enable computations for decision-making leading to action, in this case selection of the trajectory option that was not replayed. We propose that another long-term effect (bottom right) is the initiation of consolidation processes that can maintain (solid black lines), form (dashed red lines), strengthen (solid red lines), or renormalize (dashed black lines) synapses within the hippocampus, between the hippocampus and the cortex, or within the cortex. The effect of this process, which constitutes systems consolidation, is to facilitate future retrieval events. It is possible that strengthened intra-cortical synapses could also eventually support memory retrieval independent of the hippocampus.

4.6 Conclusions

The evidence suggests that SWR activity is a general mechanism for the retrieval of information gained through past experience that provides an adaptive advantage to future behaviour on multiple timescales: in decision-making and planning in the short term; and in consolidation, facilitating future instances of retrieval and use in the long term. We hypothesize that any given SWR can mediate retrieval for decision-making, planning, imagination or recollection, depending on behavioural demands and internal brain states (for example, following reward[395])[346]. Although these different demands and brain states probably induce different types of SWR events (for example, forward versus reverse, or centrifugal versus centripetal replays (Fig. 4.2)) for different immediate uses, our working hypothesis is that, in every case, the retrieved activity pattern also contributes to a consolidation process (Fig. 4.3). Although it is less clear what the immediate uses of retrieval in sleep are[240,328], we would expect the same multiplicity of function in that state. In this formulation, the SWR is a general mechanism for ongoing consolidation and retrieval processes that support a memory at every point following its encoding.

An ideal test of this hypothesis would isolate specific periods during behaviour when retrieval or consolidation are known to occur, and, within these periods, classify subtypes of SWR based on their replay content. Observation and disruption of this activity[476,477] with concordant measurement of established synaptic consolidation processes and neural activity in other brain regions would test whether there is a specific subset of SWRs that functions in consolidation. Although emerging similarities between retrieval and consolidation have been noted previously, here we have found it valuable to discuss them explicitly and in relation to the SWR as a potentially shared mechanism. Further study of SWR activity has the potential to refine these memory concepts and uncover their mechanisms.

Bibliography

[1] L. R. Squire, B. Knowlton, and G. Musen. *The structure and organization of memory*. Annu. Rev. Psychol. **44** (1993) 453. (page 1).

[2] T. den Heijer, M. I. Geerlings, F. E. Hoebeek, A. Hofman, P. J. Koudstaal, and M. M. Breteler. *Use of hippocampal and amygdalar volumes on magnetic resonance imaging to predict dementia in cognitively intact elderly people*. Arch. Gen. Psychiatry **63** (2006) 57. (page 1).

[3] et al MacQueen G.M. *Corse of illness, hippocampal function, and hippocampal volume in major depression*. Proc. Natl. Acad. Sci. USA **100** (2003) 1387. No citations.

[4] C. H. Salmond, J. Ashburner, A. Connelly, K. J. Friston, D. G. Gadian, and F. Vargha-Khadem. *The role of the medial temporal lobe in autistic spectrum disorders*. Eur. J. Neurosci. **22** (2005) 764. No citations.

[5] L. Santarelli, M. Saxe, C. Gross, A. Surget, F. Battaglia, S. Dulawa, N. Weisstaub, J. Lee, R. Duman, and O. Arancio. *Requirement of hippocampal neurogenesis for the behavioral effects of antidepressants*. Science **301** (2003) 805. No citations.

[6] A.J. et al Saykin. *Neuropsychological function in schizophrenia. selective impairment in memory and learning*. Arch. Gen. Psychiatry **48** (1991) 618. No citations.

[7] M. E. Shenton, C. C. Dickey, M. Frumin, and R. W. McCarley. *A review of mri findings in schizophrenia*. Schizophr. Res. **49** (2001) 1. (page 1).

[8] Y. Dudai. *The neurobiology of consolidations, or, how stable is the engram?* Annu. Rev. Psychol. **55** (2004) 51. (pages 1, 5).

[9] D. J. Lewis. *Psychobiology of active and inactive memory*. Psychol. Bull. **86** (1979) 1054. No citations.

[10] K. Nader. *Memory traces unbound*. Trends Neurosci. **26** (2003) 65. No citations.

[11] N. C. Tronson and J. R. Taylor. *Molecular mechanisms of memory reconsolidation*. Nat. Rev. Neurosci. **8** (2007) 262. (page 1).

[12] J. I. Gold and M. N. Shadlen. *The neural basis of decision making*. Annu. Rev. Neurosci. **30** (2007) 535. (page 1).

[13] P. A. MacDonald, M. M. Antony, C. M. Macleod, and M. A. Richter. *Memory and confidence in memory judgements among individuals with obsessive compulsive disorder and non-clinical controls*. Behav. Res. Ther. **35** (1997) 497. (pages 2, 46).

[14] A. S. Radomsky, M. J. Dugas, G. M. Alcolado, and S. L. Lavoie. *When more is less: doubt, repetition, memory, metamemory, and compulsive checking in ocd*. Behav. Res. Ther. **59** (2014) 30. (page 2).

[15] S. Eisenacher and M. Zink. *The importance of metamemory functioning to the pathogenesis of psychosis*. Front. Psychol. **8** (2017) 304. (pages 2, 46).

[16] S. Moritz and T. S. Woodward. *The contribution of metamemory deficits to schizophrenia*. J. Abnorm. Psychol. **115** (2006) 15. No citations.

[17] S. Moritz, T. S. Woodward, J. C. Whitman, and C. Cuttler. *Confidence in errors as a possible basis for delusions in schizophrenia*. J. Nerv. Ment. Dis. **193** (2005) 9. (page 2).

[18] J. K. Pannu and A. W. Kaszniak. *Metamemory experiments in neurological populations: a review*. Neuropsychol. Rev. **15** (2005) 105. (pages 2, 3, 8, 46).

[19] G. Buzsáki. *The brain from inside out*. Oxford University Press (2019). (page 2).

[20] A. L. Foote and J. D. Crystal. *Metacognition in the rat*. Curr. Biol. **17** (2007) 551. (pages 2, 7).

[21] S. Yuki and K. Okanoya. *Rats show adaptive choice in a metacognitive task with high uncertainty*. J. Exp. Psychol. Anim. Learn Cogn. **43** (2017) 109. (page 2).

[22] R. R. Hampton. *Rhesus monkeys know when they remember*. Proc. Natl. Acad. Sci. USA **98** (2001) 5359. (pages 2, 7, 42).

[23] W. E. Shields, J. D. Smith, K. Guttmannova, and D. A. Washburn. *Confidence judgments by humans and rhesus monkeys*. J. Gen. Psychol. **132** (2005) 165. No citations.

[24] D. A. Washburn, J. D. Smith, and W. E. Shields. *Rhesus monkeys (macaca mulatta) immediately generalize the uncertain response*. J. Exp. Psychol. Anim. Behav. Process **32** (2006) 185. (pages 2, 7).

[25] J. D. Smith, J. Schull, J. Strote, K. McGee, R. Egnor, and L. Erb. *The uncertain response in the bottlenosed dolphin (tursiops truncatus)*. J. Exp. Psychol. Gen. *124* (4) (1995) 391. (pages 2, 7).

[26] J. D. Smith, W. E. Shields, and D. A. Washburn. *The comparative psychology of uncertainty monitoring and metacognition*. Behav. Brain Sci. **26** (2003) 317. (page 2).

[27] D. A. Moore and P. J. Healy. *The trouble with overconfidence*. Psychol. Rev. **115** (2008) 502. (page 2).

[28] W. M. Petrusic and J. V. Baranski. *Judging confidence influences decision processing in comparative judgments*. Psychon. Bull. Rev. **10** (2003) 177. (page 13).

[29] N. Shea, A. Boldt, D. Bang, N. Yeung, C. Heyes, and C. D. Frith. *Supra-personal cognitive control and metacognition.* Trends Cogn. Sci. **18** (2014) 186. No citations.

[30] N. Yeung and C. Summerfield. *Metacognition in human decision-making: confidence and error monitoring.* Philos. Trans. R. Soc. Lond. B Biol. Sci. **367** (2012) 1310. (pages 2, 44).

[31] H. Eichenbaum and N. J. Cohen. *From conditioning to conscious recollection.* Oxford University Press (2001). (page 3).

[32] W. B. Scoville and B. Milner. *Loss of recent memory after bilateral hippocampal lesions.* J. Neurol. Neurosurg. Psychiatry **20** (1957) 11. (page 3).

[33] G. Buzsáki. *Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning.* Hippocampus **25** (2015) 1073. (pages 3, 4, 9).

[34] A. Kepecs, N. Uchida, H. A. Zariwala, and Z. F. Mainen. *Neural correlates, computation and behavioural impact of decision confidence.* Nature **455** (2008) 227. (pages 3, 4, 7, 10).

[35] A. Lak, G. M. Costa, E. Romberg, A. A. Koulakov, Z. F. Mainen, and A. Kepecs. *Orbitofrontal cortex is required for optimal waiting based on decision confidence.* Neuron **84** (2014) 190. (pages 3, 7, 10, 13).

[36] K. Miyamoto, T. Osada, R. Setsuie, M. Takeda, K. Tamura, Y. Adachi, and Y. Miyashita. *Causal neural network of metamemory for retrospection in primates.* Science **355** (2017) 188. (pages 3, 7, 8, 8, 9, 43).

[37] U. Rutishauser, S. Ye, M. Koroma, O. Tudusciuc, I. B. Ross, J. M. Chung, and A. N. Mamelak. *Representation of retrieval confidence by single neurons in the human medial temporal lobe.* Nat. Neurosci. **18** (2015) 1041. (pages 3, 9, 43).

[38] Marielena Sosa, Hannah R. Joo, and Loren M. Frank. *Dorsal and Ventral Hippocampal Sharp-Wave Ripples Activate Distinct Nucleus Accumbens Networks.* Neuron **105** (2020) 725. (pages 3, 4, 206).

[39] J. E. Chung et al. *A fully automated approach to spike sorting.* Neuron **95** (2017) 1381. (pages 4, 85).

[40] B. Hangya, J. I. Sanders, and A. Kepecs. *A mathematical framework for statistical decision confidence.* Neural. Comput. **28** (2016) 1840. (pages 4, 7, 42).

[41] S. P. Jadhav, G. Rothschild, D. K. Roumis, and L. M. Frank. *Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events.* Neuron **90** (2016) 113. (page 4).

[42] D. Ji and M. A. Wilson. *Coordinated memory replay in the visual cortex and hippocampus during sleep.* Nat. Neurosci. **10** (2007) 100. No citations.

[43] N. K. Logothetis, O. Eschenko, Y. Murayama, M. Augath, T. Steudel, H. C. Evrard, M. Besserve, and A. Oeltermann. *Hippocampal-cortical interaction during periods of subcortical silence.* Nature **491** (2012) 547. No citations.

[44] M. Remondes and M. A. Wilson. *Slow-gamma rhythms coordinate cingulate cortical responses to hippocampal sharp-wave ripples during wakefulness.* Cell Rep. **13** (2015) 1327. No citations.

[45] G. Rothschild, E. Eban, and L. M. Frank. *A cortical-hippocampal-cortical loop of information processing during memory consolidation.* Nat. Neurosci. *20*(2) (2017) 251. No citations.

[46] A. G. Siapas and M. A. Wilson. *Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep.* Neuron **21** (1998) 1123. No citations.

[47] A. Sirota, J. Csicsvari, D. Buhl, and G. Buzsáki. *Communication between neocortex and hippocampus during sleep in rodents.* Proc. Natl. Acad. Sci. USA **100** (2003) 2065. No citations.

[48] C. T. Wu, D. Haggerty, C. Kemere, and D. Ji. *Hippocampal awake replay in fear memory retrieval.* Nat. Neurosci. **20** (2017) 571. (pages 4, 5).

[49] G Buzsáki. *The hippocampo-neocortical dialogue.* Cerebral Cortex **6** (1996) 81. (page 5).

[50] G. Buzsáki. *Two-stage model of memory trace formation: a role for "noisy" brain states.* Neuroscience **31** (1989) 551. No citations.

[51] J. J. Kim and M. S. Fanselow. *Modality-specific retrograde amnesia of fear.* Science **256** (1992) 675. No citations.

[52] D. Marr. *Simple memory: a theory for archicortex.* Philos. Trans. R. Soc. Lond. B Biol. Sci. **262** (1971) 23. No citations.

[53] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly. *Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory.* Psychol. Rev. **102** (1995) 419. No citations.

[54] L. R. Squire and S. Zola-Morgan. *The medial temporal lobe memory system.* Science **253** (1991) 1380. (page 5).

[55] R. E. Ambrose, B. E. Pfeiffer, and D. J. Foster. *Reverse replay of hippocampal place cells is uniquely modulated by changing reward.* Neuron **91** (2016) 1124. (page 5).

[56] A. C. Singer and L. M. Frank. *Rewarded outcomes enhance reactivation of experience in the hippocampus.* Neuron **64** (2009) 910. (page 5).

[57] M. P. Karlsson and L. M. Frank. *Awake replay of remote experiences in the hippocampus*. Nat. Neurosci. **12** (2009) 913. (page 5).

[58] M. F. Carr, S. P. Jadhav, and L. M. Frank. *Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval*. Nat. Neurosci. **14** (2011) 147. (page 5).

[59] A. Johnson and A. D. Redish. *Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point*. J. Neurosci. **27** (2007) 12176. No citations.

[60] B. E. Pfeiffer and D. J. Foster. *Hippocampal place-cell sequences depict future paths to remembered goals*. Nature **497** (2013) 74. No citations.

[61] A. C. Singer, M. F. Carr, M. P. Karlsson, and L. M. Frank. *Hippocampal swr activity predicts correct decisions during the initial learning of an alternation task*. Neuron **77** (2013) 1163. No citations.

[62] J. Y. Yu and L. M. Frank. *Hippocampal-cortical interaction in decision making*. Neurobiol. Learn. Mem. **117** (2014). (page 5).

[63] S. P. Jadhav, C. Kemere, P. W. German, and L. M. Frank. *Awake hippocampal sharp-wave ripples support spatial memory*. Science **336** (2012) 1454. (page 5).

[64] M. S. Nokia, J. E. Mikkonen, M. Penttonen, and J. Wikgren. *Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning*. Front. Behav. Neurosci. **6** (2012) 84. (page 5).

[65] Endel Tulving. *Elements of episodic memory*. Oxford psychology series (1983) 351. (pages 7, 9).

[66] J. I. Gold and A. A. Stocker. *Visual decision-making in an uncertain and dynamic world*. Annu. Rev. Vis. Sci. **3** (2017) 227. (page 7).

[67] C. Summerfield and F. P. de Lange. *Expectation in perceptual decision making: neural and computational mechanisms*. Nat. Rev. Neurosci. **15** (2014) 745. (page 7).

[68] D. R. Bach and R. J. Dolan. *Knowing how much you don't know: a neural organization of uncertainty estimates*. Nat. Rev. Neurosci. **13** (2012) 572. (page 7).

[69] P. Grimaldi, H. Lau, and M. A. Basso. *There are things that we know that we know, and there are things that we do not know we do not know: Confidence in decision-making*. Neurosci. Biobehav. Rev. **55** (2015) 88. No citations.

[70] A. Pouget, J. Drugowitsch, and A. Kepecs. *Confidence and certainty: distinct probabilistic quantities for different goals*. Nat. Neurosci. **19** (2016) 366. (pages 7, 8).

[71] F. Meyniel, M. Sigman, and Z. F. Mainen. *Confidence as bayesian probability: From neural origins to behavior*. Neuron **88** (2015) 78. (page 7).

[72] A. Pouget, J. M. Beck, W. J. Ma, and P. E. Latham. *Probabilistic brains: knowns and unknowns*. Nat. Neurosci. **16** (2013) 1170. (page 7).

[73] F. Meyniel and S. Dehaene. *Brain networks for confidence weighting and hierarchical inference during probabilistic learning*. Proc. Natl. Acad. Sci. USA **114** (2017) E3859. (page 7).

[74] F. Meyniel, D. Schlunegger, and S. Dehaene. *The sense of confidence during probabilistic learning: A normative account*. PLoS Comput. Biol. **11** (2015) e1004305. No citations.

[75] A. Soltani and A. Izquierdo. *Adaptive learning under expected and unexpected uncertainty*. Nat. Rev. Neurosci. **20** (2019) 635. (page 7).

[76] A. Kepecs and Z. F. Mainen. *A computational framework for the study of confidence in humans and animals*. Philos. Trans. R. Soc. Lond. B Biol. Sci. **367** (2012) 1322. (pages 7, 10).

[77] J. I. Sanders, B. Hangya, and A. Kepecs. *Signatures of a statistical computation in the human sense of confidence*. Neuron **90** (2016) 499. (pages 7, 42).

[78] Ariel Zylberberg, Christopher R. Fetsch, and Michael N. Shadlen. *The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision*. eLife **5** (2016) e17688. (page 7, 7).

[79] A. Lak, K. Nomoto, M. Keramati, M. Sakagami, and A. Kepecs. *Midbrain dopamine neurons signal belief in choice accuracy during a perceptual decision*. Curr. Biol. **27** (2017) 821. (page 7).

[80] K. Miyamoto, R. Setsuie, T. Osada, and Y. Miyashita. *Reversible silencing of the frontopolar cortex selectively impairs metacognitive judgment on non-experience in primates*. Neuron **97** (2018) 980. (page 9).

[81] R. Kiani and M. N. Shadlen. *Representation of confidence associated with a decision by neurons in the parietal cortex*. Science **324** (2009) 759. (pages 7, 8).

[82] P. G. Middlebrooks and M. A. Sommer. *Metacognition in monkeys during an oculomotor task*. J. Exp. Psychol. Learn Mem. Cogn. **37** (2011) 325. (page 7).

[83] P. G. Middlebrooks and M. A. Sommer. *Neuronal correlates of metacognition in primate frontal cortex*. Neuron **75** (2012) 517. (page 7).

[84] A. Q. Vining and H. L. Marsh. *Information seeking in capuchins (cebus apella): a rudimentary form of metacognition?* Anim. Cogn. **18** (2015) 667. (page 7).

[85] C. J. Perry and A. B. Barron. *Honey bees selectively avoid difficult choices*. Proc. Natl. Acad. Sci. USA **110** (2013) 19155. (page 7).

[86] A. Stolyarova, M. Rakhshan, E. E. Hart, T. J. O'Dell, M. A. K. Peters, H. Lau, A. Soltani, and A. Izquierdo. *Contributions of anterior cingulate cortex and*

basolateral amygdala to decision confidence and learning under uncertainty. Nat. Commun. **10** (2019) 4704. (pages 7, 8).

[87] N. A. A. Atiya, I. Rano, G. Prasad, and K. Wong-Lin. *A neural circuit model of decision uncertainty and change-of-mind*. Nat. Commun. **10** (2019) 2287. (page 7).

[88] C. R. Fetsch, R. Kiani, W. T. Newsome, and M. N. Shadlen. *Effects of cortical microstimulation on confidence in a perceptual decision*. Neuron **84** (2014) 239. (page 7).

[89] Thomas O. Nelson. *Metamemory: A theoretical framework and new findings*. Psychology of Learning and Motivation **26** (1990) 125. (page 8).

[90] D. Bang and S. M. Fleming. *Distinct encoding of decision confidence in human medial prefrontal cortex*. Proc. Natl. Acad. Sci. USA **115** (2018) 6082. (page 8).

[91] Kobe Desender, Annika Boldt, Tom Verguts, and Tobias H. Donner. *Confidence predicts speed-accuracy tradeoff for subsequent decisions*. eLife **8** (2019) e43499. No citations.

[92] B. Odegaard, P. Grimaldi, S. H. Cho, M. A. K. Peters, H. Lau, and M. A. Basso. *Superior colliculus neuronal ensemble activity signals optimal rather than subjective confidence*. Proc. Natl. Acad. Sci. USA **115** (2018) E1588. (page 8).

[93] Y. Dudai. *The restless engram: consolidations never end*. Annu. Rev. Neurosci. **35** (2012) 227. (page 8).

[94] H. R. Joo and L. M. Frank. *The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation*. Nat. Rev. Neurosci. **19** (2018) 744. (page 8).

[95] A. G. Vaccaro and S. M. Fleming. *Thinking about thinking: A coordinate-based meta-analysis of neuroimaging studies of metacognitive judgements*. Brain Neurosci. Adv. **2** (2018) 2398212818810591. (pages 8, 9).

[96] Peter Carruthers. *Meta-cognition in animals: A skeptical look*. Mind & Language **23** (2008) 58. (page 8).

[97] Nate Kornell. *Metacognition in humans and animals*. Current Directions in Psychological Science **18** (2009) 11. No citations.

[98] J. D. Smith. *The study of animal metacognition*. Trends Cogn. Sci. **13** (2009) 389. No citations.

[99] T. Suddendorf and M. C. Corballis. *Behavioural evidence for mental time travel in nonhuman animals*. Behav. Brain Res. **215** (2010) 292. (pages 8, 9, 9).

[100] E. Tulving. *Multiple memory systems and consciousness*. Human neurobiology **6** (1987) 67. (page 8).

[101] Endel Tulving. *Memory and consciousness*. Canadian Psychology **26** (1985) 1. (pages 8, 9).

[102] J. M. Gardiner. *Episodic memory and autonoetic consciousness: a first-person approach*. Philosophical transactions of the Royal Society of London. Series B, Biological sciences **356** (2001) 1351. (page 9).

[103] M. A. Wheeler, D. T. Stuss, and E. Tulving. *Toward a theory of episodic memory: the frontal lobes and autonoetic consciousness*. Psychol. Bull. **121** (1997) 331. (page 9).

[104] Endel Tulving. *Episodic memory and autonoesis: Uniquely human?* The missing link in cognition: Origins of self-reflective consciousness. (2005) 3. (page 9).

[105] S. C. Kwok, Y. Cai, and M. J. Buckley. *Mnemonic introspection in macaques is dependent on superior dorsolateral prefrontal cortex but not orbitofrontal cortex*. J. Neurosci. **39** (2019) 5922. (page 9).

[106] C. E. Stern and Michael Hasselmo. *Recognition memory*. Encyclopedia of Neuroscience (2010) 49. (page 9).

[107] B. Baird, M. Cieslak, J. Smallwood, S. T. Grafton, and J. W. Schooler. *Regional white matter variation associated with domain-specific metacognitive accuracy.* J. Cogn. Neurosci. **27** (2015) 440. (page 9).

[108] E. F. Chua, D. L. Schacter, and R. A. Sperling. *Neural correlates of metamemory: a comparison of feeling-of-knowing and retrospective confidence judgments.* J. Cogn. Neurosci. **21** (2009) 1751. (page 9).

[109] J. A. Elman, E. C. Klostermann, D. E. Marian, A. Verstaen, and A. P. Shimamura. *Neural correlates of metacognitive monitoring during episodic and semantic retrieval.* Cogn. Affect Behav. Neurosci. **12** (2012) 599. (page 9).

[110] Y. C. Kao, E. S. Davis, and J. D. Gabrieli. *Neural correlates of actual and predicted memory formation.* Nat. Neurosci. **8** (2005) 1776. (page 9).

[111] R. N. Henson, M. D. Rugg, T. Shallice, and R. J. Dolan. *Confidence in recognition memory for words: dissociating right prefrontal roles in episodic retrieval.* J. Cogn. Neurosci. **12** (2000) 913. (page 9).

[112] N. Reggev, M. Zuckerman, and A. Maril. *Are all judgments created equal? an fmri study of semantic and episodic metamemory predictions.* Neuropsychologia **49** (2011) 1332. (page 9).

[113] U. M. Risius, A. Staniloiu, M. Piefke, S. Maderwald, F. P. Schulte, M. Brand, and H. J. Markowitsch. *Retrieval, monitoring, and control processes: a 7 tesla fmri approach to memory accuracy.* Front. Behav. Neurosci. **7** (2013) 24. (page 9).

[114] Janet Metcalfe and Lisa K. Son. *Anoetic, noetic, and autonoetic metacognition.* Foundations of metacognition (2012) 289. (page 9).

[115] N. S. Clayton and A. Dickinson. *Episodic-like memory during cache recovery by scrub jays.* Nature **395** (1998) 272. (page 9).

[116] H. Eichenbaum. *Prefrontal-hippocampal interactions in episodic memory.* Nat. Rev. Neurosci. (2017). (page 9).

[117] N. J. Fortin, K. L. Agster, and H. B. Eichenbaum. *Critical role of the hippocampus in memory for sequences of events.* Nat. Neurosci. *5*(5) (2002) 458. (page 9, 9).

[118] K. Kay and L. M. Frank. *Three brain states in the hippocampus and cortex.* Hippocampus (2018). (page 9).

[119] J. E. Chung, H. R. Joo, J. L. Fan, D. F. Liu, A. H. Barnett, S. Chen, C. Geaghan-Breiner, M. P. Karlsson, M. Karlsson, K. Y. Lee, H. Liang, J. F. Magland, J. A. Pebbles, A. C. Tooker, L. F. Greengard, V. M. Tolosa, and L. M. Frank. *High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays.* Neuron **101** (2019) 21. (pages 10, 74, 74, 74, 74, 85, 85, 85, 85, 90).

[120] J. E. Chung, H. R. Joo, C. N. Smyth, J. L. Fan, C. Geaghan-Breiner, H. Liang, D. F. Liu, D. Roumis, S. Chen, K. Y. Lee, J. A. Pebbles, A. C. Tooker, V. M. Tolosa, and L. M. Frank. *Chronic implantation of multiple flexible polymer electrode arrays.* J. Vis. Exp. **152** (2019). (page 10).

[121] N. Persaud, P. McLeod, and A. Cowey. *Post-decision wagering objectively measures awareness.* Nat. Neurosci. **10** (2007) 257. (page 10).

[122] Thomas O. Nelson. *Consciousness and metacognition.* American Psychologist (2) (1996) 102. (pages 10, 41).

[123] A. R. Powers, C. Mathys, and P. R. Corlett. *Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors.* Science **357** (2017) 596. (page 10).

[124] J. D. Smith, A. C. Zakrzewski, and B. A. Church. *Formal models in animal-metacognition research: the problem of interpreting animals' behavior.* Psychon. Bull. Rev. **23** (2016) 1341. (page 13).

[125] Jonathon D. Crystal and Allison L. Foote. *Evaluating information-seeking approaches to metacognition.* Current Zoology **57** (2011) 531. (page 13).

[126] Peter R. Murphy, Ian H. Robertson, Siobhán Harty, and Redmond G. O'Connell. *Neural evidence accumulation persists after choice to inform metacognitive judgments.* eLife **4** (2015) e11946. (page 13).

[127] Joaquin Navajas, Bahador Bahrami, and Peter E. Latham. *Post-decisional accounts of biases in confidence.* Current Opinion in Behavioral Sciences **11** (2016) 55. (page 13).

[128] L. Green and J. Myerson. *A discounting framework for choice with delayed and probabilistic rewards.* Psychol. Bull. **130** (2004) 769. (page 13).

[129] A. Vanderveldt, L. Oliveira, and L. Green. *Delay discounting: Pigeon, rat, human–does it matter?* J. Exp. Psychol. Anim. Learn. Cogn. **42** (2016) 141. (page 13).

[130] R. P. Kesner and J. M. Novak. *Serial position curve in rats: role of the dorsal hippocampus.* Science **218** (1982) 173. (page 16).

[131] A. A. Chiba, R. P. Kesner, and C. J. Gibson. *Memory for temporal order of new and familiar spatial location sequences: role of the medial prefrontal cortex.* Learn Mem. **4** (1997) 311. (page 16).

[132] V. L. Templer and R. R. Hampton. *Cognitive mechanisms of memory for order in rhesus monkeys (macaca mulatta).* Hippocampus **23** (2013) 193. (page 23).

[133] J. D. Smith, J. J. Couchman, and M. J. Beran. *Animal metacognition: a tale of two comparative psychologies*. J. Comp. Psychol. **128** (2014) 115. (page 41).

[134] Robert S. Lockhart and Bennet B. Murdock. *Memory and the theory of signal detection*. Psychological Bulletin **74** (1970) 100. (pages 41, 41, 44).

[135] K. J. Malmberg. *On the form of rocs constructed from confidence ratings*. J. Exp. Psychol. Learn Mem. Cogn. **28** (2002) 380. No citations.

[136] Paul Masset and Adam Kepecs. *Resolving the reversal paradox of memory confidence*. Conference on Cognitive Computational Neuroscience (2017). (page 41).

[137] S. M. Fleming and N. D. Daw. *Self-evaluation of decision-making: A general bayesian framework for metacognitive computation*. Psychol. Rev. **124** (2017) 91. (page 42, 42).

[138] M. Rausch and M. Zehetleitner. *The folded x-pattern is not necessarily a statistical signature of decision confidence*. PLoS Comput. Biol. **15** (2019) e1007456. (page 42).

[139] W. T. Adler and W. J. Ma. *Comparing bayesian and non-bayesian accounts of human confidence reports*. PLoS Comput. Biol. **14** (2018) e1006572. (page 42).

[140] W. T. Adler and W. J. Ma. *Limitations of proposed signatures of bayesian confidence*. Neural Comput. (2018) 1. (page 42).

[141] A. P. Yonelinas, M. Aly, W. C. Wang, and J. D. Koen. *Recollection and familiarity: examining controversial assumptions and new directions*. Hippocampus *20*(11) (2010) 1178. (page 43).

[142] Andrew Yonelinas. *The nature of recollection and familiarity: A review of 30 years of research*. Journal of Memory and Language **46** (2002) 441. (page 43).

[143] John T. Wixted and Laura Mickes. *A continuous dual-process model of remember/know judgments*. Psychological Review *117*(4) (2010) 1025. (page 44).

[144] D. A. Norman and W. A. Wickelgren. *Short-term recognition memory for single digits and pairs of digits*. J. Exp. Psychol. **70** (1965) 479. (page 44).

[145] Y. Dudai and M. Carruthers. *The janus face of mnemosyne*. Nature **434** (2005) 567. (page 45).

[146] D. L. Schacter, D. R. Addis, D. Hassabis, V. C. Martin, R. N. Spreng, and K. K. Szpunar. *The future of memory: remembering, imagining, and the brain*. Neuron **76** (2012) 677. (page 45).

[147] Asher Koriat. *The subjective confidence in one's knowledge and judgements: Some metatheoretical considerations*. Foundations of Metacognition (2013). (page 45).

[148] Sheryl Gay Stolberg and Nicholas Fandos. *Brett Kavanaugh and Christine Blasey Ford Duel With Tears and Fury*. The New York Times (2018). (page 45).

[149] M. Rouault, T. Seow, C. M. Gillan, and S. M. Fleming. *Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance*. Biol. Psychiatry **84** (2018) 443. (page 45).

[150] P. C. Fletcher and C. D. Frith. *Perceiving is believing: a bayesian approach to explaining the positive symptoms of schizophrenia*. Nat. Rev. Neurosci. **10** (2009) 48. (page 45).

[151] Philipp Sterzer, Rick A. Adams, Paul Fletcher, Chris Frith, Stephen M. Lawrie, Lars Muckli, Predrag Petrovic, Peter Uhlhaas, Martin Voss, and Philip R. Corlett. *The predictive coding account of psychosis*. Biological Psychiatry **84** (2018) 634. (page 45).

[152] I. Fradkin, C. Ludwig, E. Eldar, and J. D. Huppert. *Doubting what you already know: Uncertainty regarding state transitions is associated with obsessive compulsive symptoms*. PLoS Comput. Biol. **16** (2020) e1007634. (page 46).

[153] A. S. David, N. Bedford, B. Wiffen, and J. Gilleen. *Failures of metacognition and lack of insight in neuropsychiatric disorders.* Philos. Trans. R. Soc. Lond. B Biol. Sci. *367*(1594) (2012) 1379. (page 46).

[154] Demetris Roumis. *Rhythmic Action Synchronizes Memory Replay During Reinforcement Learning.* UCSF Thesis (2020). (page 46).

[155] Fraciois Chollet. *Keras.* GitHub repository (2017). https://github.com/fchollet/keras. (page 51).

[156] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. *Tensorflow: A system for large-scale machine learning.* 12th USENIX Symposium on Operating Systems Design and Implementation **16** (2016) 265. (page 51).

[157] Diederik Kingma and Jimmy Ba. *Adam: A method for stochastic optimization.* arXiv:1412.6980 (2014). (page 51).

[158] F. Pedregosa et al. *Scikit-learn: Machine learning in Python.* Journal of Machine Learning Research **12** (2011) 2825. (page 52).

[159] Skipper Seabold and Josef Perktold. *statsmodels: Econometric and statistical modeling with python.* 9th Python in Science Conference (2010). (page 52).

[160] E. Tóth, D. Fabó, L. Entz, I. Ulbert, and L. Erőss. *Intracranial neuronal ensemble recordings and analysis in epilepsy.* J. Neurosci. Methods **260** (2016) 261. (page 73).

[161] Y. H. Luo and L. da Cruz. *The Argus® II Retinal Prosthesis System.* Prog. Retin. Eye. Res. **50** (2016) 89. (page 73).

[162] A. M. Lozano et al. *Deep brain stimulation: current challenges and future directions.* Nat. Rev. Neurol. **15** (2019). (page 73).

[163] D. Kipke, R. Vetter, J. Williams, and J. Hetke. *Silicon-Substrate Intracortical Microelectrode Arrays for Long-Term Recording of Neuronal Spike Activity in Cerebral Cortex*. IEEE Trans. Neural Syst. Rehabil. Eng. **11** (2003) 151. (page 73).

[164] P. J. Rousche and R. A. Normann. *Chronic recording capability of the Utah Intracortical Electrode Array in cat sensory cortex*. J. Neurosci. Methods *82*(1) (1998) 1. (page 73).

[165] M. A. Nicolelis, D. Dimitrov, J. M. Carmena, R. Crist, G. Lehew, J. D. Kralik, and S. P. Wise. *Chronic, multisite, multielectrode recordings in macaque monkeys*. Proc. Natl. Acad. Sci. USA **100** (2003) 11041. (page 73).

[166] K. Mols, S. Musa, B. Nuttin, L. Lagae, and V. Bonin. *In vivo characterization of the electrophysiological and astrocytic responses to a silicon neuroprobe implanted in the mouse neocortex*. Sci. Rep. **7** (2017) 15642. (page 73).

[167] J. W. Jeong, G. Shin, S. I. Park, K. J. Yu, L. Xu, and J. A. Rogers. *Soft materials in neuroengineering for hard problems in neuroscience*. Neuron **86** (2015) 175. (pages 73, 74, 74, 74).

[168] V. S. Polikov, P. A. Tresco, and W. M. Reichert. *Response of brain tissue to chronically implanted neural electrodes*. J. Neurosci. Methods **148** (2005) 1. (page 73).

[169] J. J. Jun et al. *Fully integrated silicon probes for high-density recording of neural activity*. Nature **551** (2017) 232. (pages 73, 93).

[170] T. M. Fu, G. S. Hong, T. Zhou, T. G. Schuhmann, R. D. Viveros, and C. M. Lieber. *Stable long-term chronic brain mapping at the single-neuron level*. Nature Methods **13** (2016) 875. (page 74, 74).

[171] H. S. Sohal, A. Jackson, R. Jackson, G. J. Clowry, K. Vassilevski, A. O'Neill, and

S. N. Baker. *The sinusoidal probe: a new approach to improve electrode longevity.* Front. Neuroeng. **7** (2014) 10. No citations.

[172] C. Xie, J. Liu, X. Dai, W. Zhou, and C.M. Lieber. *Three-dimensional macroporous nanoelectronic networks as minimally invasive brain probes.* Nature Materials **14** (2015) 1286. (page 74, 74).

[173] T. Zhou, G. Hong, T. M. Fu, X. Yang, T. G. Schuhmann, R. D. Viveros, and C. M. Lieber. *Syringe-injectable mesh electronics integrate seamlessly with minimal chronic immune response in the brain.* Proc. Natl. Acad. Sci. USA **114** (2017) 5894. (page 74).

[174] J. P. Harris, J. R. Capadona, R. H. Miller, B. C. Healy, K. Shanmuganathan, S. J. Rowan, C. Weder, and D. J. Tyler. *Mechanically adaptive intracortical implants improve the proximity of neuronal cell bodies.* J. Neural Eng. **8** (2011) 066011. (page 74).

[175] H. C. Lee, F. Ejserholm, J. Gaire, S. Currlin, J. Schouenborg, L. Wallman, M. Bengtsson, K. Park, and K. J. Otto. *Histological evaluation of flexible neural implants; flexibility limit for reducing the tissue response?* J. Neural Eng. **14** (2017) 036026. No citations.

[176] D. H. Szarowski, M. D. Andersen, S. Retterer, A. J. Spence, M. Isaacson, H. G. Craighead, J. N. Turner, and W. Shain. *Brain responses to micro-machined silicon devices.* Brain Res. **983** (2003) 23. (page 74).

[177] G. Hong and C. M. Lieber. *Novel electrode technologies for neural recordings.* Nat. Rev. Neurosci. **20** (2019) 330. (page 74).

[178] S. P. Lacour, G. Courtine, and J. Guck. *Materials and technologies for soft implantable neuroprostheses.* Nature Reviews Materials **1** (2016). (page 74).

[179] T. D. Kozai, N. B. Langhals, P. R. Patel, X. Deng, H. Zhang, K. L. Smith, J. Lahann, N. A. Kotov, and D. R. Kipke. *Ultrasmall implantable composite microelectrodes with bioactive surfaces for chronic neural interfaces.* Nat. Mater. **11** (2012) 1065. (page 74, 74).

[180] L. Luan et al. *Ultraflexible nanoelectronic probes form reliable, glial scar-free neural integration.* Sci. Adv. **3** (2017) e1601966. (page 74).

[181] J. P. Seymour and D. R. Kipke. *Neural probe design for reduced tissue encapsulation in CNS.* Biomaterials **28** (2007) 3594. (page 74, 74).

[182] Aziliz Lecomte, Emeline Descamps, and Christian Bergaud. *A review on mechanical considerations for chronically-implanted neural probes.* Journal of Neural Engineering **15** (2018) 031001. (page 74).

[183] Roy O. Weller, Matthew M. Sharp, Myron Christodoulides, Roxana O. Carare, and Kjeld Møllgård. *The meninges as barriers and facilitators for the movement of fluid, cells and pathogens related to the rodent and human CNS.* Acta Neuropathologica **135** (2018) 363. (page 74).

[184] F. Wu, L. W. Tien, F. Chen, J. D. Berke, D. L. Kaplan, and E. Yoon. *Silk-Backed Structural Optimization of High-Density Flexible Intracortical Neural Probes.* Journal of Microelectromechanical Systems **24** (2015) 62. (page 74).

[185] Zhuolin Xiang, Shih-Cheng Yen, Ning Xue, Tao Sun, Wei Mong Tsang, Songsong Zhang, Lun-De Liao, Nitish V Thakor, and Chengkuo Lee. *Ultra-thin flexible polyimide neural probe embedded in a dissolvable maltose-coated microneedle.* Journal of Micromechanics and Microengineering **24** (2014) 065015. No citations.

[186] Rakesh Khilwani et al. *Ultra-miniature ultra-compliant neural probes with dissolvable delivery needles: design, fabrication and characterization.* Biomedical Microdevices **18** (2016) 97. No citations.

[187] Paras R Patel, Kyounghwan Na, Huanan Zhang, Takashi D Y Kozai, Nicholas A Kotov, Euisik Yoon, and Cynthia A Chestek. *Insertion of linear 8.4 μm diameter 16 channel carbon fiber electrode arrays for single unit recordings.* Journal of Neural Engineering *12*(4) (2015) 046009. (page 74).

[188] B J Kim, J T W Kuo, S A Hara, C D Lee, L Yu, C A Gutierrez, T Q Hoang, V Pikov, and E Meng. *3D Parylene sheath neural probe for chronic recordings.* Journal of Neural Engineering **10** (2013) 045002. No citations.

[189] Meng-chen Lo, Shuwu Wang, Sagar Singh, Vinod B. Damodaran, Hilton M. Kaplan, Joachim Kohn, David I. Shreiber, and Jeffrey D. Zahn. *Coating flexible probes with an ultra fast degrading polymer to aid in tissue insertion.* Biomedical Microdevices **17** (2015) 34. (page 74).

[190] Flavia Vitale et al. *Fluidic Microactuation of Flexible Electrodes for Neural Recording.* Nano Letters **18** (2018) 326. (page 74).

[191] Thomas G. Schuhmann, Tao Zhou, Guosong Hong, Jung Min Lee, Tian-Ming Fu, Hong-Gyu Park, and Charles M. Lieber. *Syringe-injectable Mesh Electronics for Stable Chronic Rodent Electrophysiology.* Journal of Visualized Experiments **137** (2018) 58003. No citations.

[192] Jia Liu et al. *Syringe-injectable electronics.* Nature Nanotechnology **10** (2015) 629. No citations.

[193] Shoji Takeuchi, D. Ziegler, Y. Yoshida, K. Mabuchi, and T. Suzuki. *Parylene flexible neural probes integrated with microfluidic channels.* Lab on a Chip **5** (2005) 519. (page 74).

[194] Dustin M. Simon et al. *Design and demonstration of an intracortical probe technology with tunable modulus: design and demonstration of an intracortical probe technology.* Journal of Biomedical Materials Research Part A **105** (2017) 159. (page 74).

[195] Takashi D. Yoshida Kozai and Daryl R. Kipke. *Insertion shuttle with carboxyl terminated self-assembled monolayer coatings for implanting flexible polymer neural probes in the brain.* Journal of Neuroscience Methods **184** (2009) 199. (page 74).

[196] S. Felix, K. Shah, D. George, V. Tolosa, A. Tooker, H. Sheth, T. Delima, and S. Pannu. *Removable silicon insertion stiffeners for neural probes using polyethylene glycol as a biodissolvable adhesive.* 2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (2012) 871. (pages 74, 77).

[197] Z. Zhao, X. Li, F. He, X. Wei, S. Lin, and C. Xie. *Parallel, minimally-invasive implantation of ultra-flexible neural electrode arrays.* J. Neural Eng. **16** (2019) 035001. (pages 74, 95, 95).

[198] Timothy L Hanson, Camilo A Diaz-Botia, Viktor Kharazia, Michel M Maharbiz, and Philip N Sabes. *The "sewing machine" for minimally invasive neural recording.* bioRxiv (2019) 578542. (pages 74, 95, 96).

[199] Abdulmalik M. Obaid, Yu-Wei Wu, Mina-Elraheb Hanna, William D. Nix, Jun B. Ding, and Nicholas A. Melosh. *Ultra-sensitive measurement of brain penetration with microscale probes for brain machine interface considerations.* bioRxiv (2018) 454520. (pages 75, 75, 83, 92, 94, 94, 95).

[200] A. A. Sharp, A. M. Ortega, D. Restrepo, D. Curran-Everett, and K. Gall. *In vivo penetration mechanics and mechanical properties of mouse brain tissue at micrometer scales.* IEEE Trans. Biomed. Eng. **56** (2009) 45. (pages 75, 75, 92, 94).

[201] C. S. Bjornsson, S. J. Oh, Y. A. Al-Kofahi, Y. J. Lim, K. L. Smith, J. N. Turner, S. De, B. Roysam, W. Shain, and S. J. Kim. *Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion.* J. Neural Eng. **3** (2006) 196. (page 96).

[202] D.J. Edell, V.V. Toi, V.M. McNeil, and L.D. Clark. *Factors influencing the biocompatibility of insertable silicon microshafts in cerebral cortex.* IEEE Transactions on Biomedical Engineering **39** (1992) 635. No citations.

[203] W. Jensen, K. Yoshida, and U.G. Hofmann. *In-vivo implant mechanics of flexible, silicon-based ACREO microelectrode arrays in rat cerebral cortex.* IEEE Transactions on Biomedical Engineering **53** (2006) 934. (page 75).

[204] Z. Fekete, A. Nemeth, G. Marton, I. Ulbert, and A. Pongracz. *Experimental study on the mechanical interaction between silicon neural microprobes and rat dura mater during insertion.* J. Mater Sci. Mater. Med. **26** (2015) 70. (pages 75, 75, 92, 92, 92, 93, 93, 94, 95, 96).

[205] N. H. Hosseini, R. Hoffmann, S. Kisban, T. Stieglitz, O. Paul, and P. Ruther. *Comparative study on the insertion behavior of cerebral microprobes.* Conf. Proc. IEEE Eng. Med. Biol. Soc. **2007** (2007) 4711. (pages 75, 92, 93, 95, 96).

[206] Kyounghwan Na, Zachariah J Sperry, Jiaao Lu, Mihaly Voeroeslakos, Saman S Parizi, Tim M Bruns, Euisik Yoon, and John P Seymour. *Novel diamond shuttle to deliver flexible bioelectronics with reduced tissue compression.* bioRxiv (2018) 435800. (pages 75, 94, 94, 95, 96).

[207] A. J. Shoffstall, S. Srinivasan, M. Willis, A. M. Stiller, M. Ecker, W. E. Voit, J. J. Pancrazio, and J. R. Capadona. *A mosquito inspired strategy to implant microprobes into the brain.* Sci. Rep. **8** (2018) 122. (pages 75, 95, 96).

[208] J. T. Maikos, R. A. Elias, and D. I. Shreiber. *Mechanical properties of dura mater from the rat brain and spinal cord.* J. Neurotrauma. **25** (2008) 38. (pages 76, 94).

[209] C. E. Dixon, G. L. Clifton, J. W. Lighthall, A. A. Yaghmai, and R. L. Hayes. *A controlled cortical impact model of traumatic brain injury in the rat.* J Neurosci Methods **39** (1991) 253. (page 76).

[210] A. Sridharan, S. D. Rajan, and J. Muthuswamy. *Long-term changes in the material properties of brain tissue at the implant-tissue interface.* J. Neural Eng. **10** (2013) 066001. (pages 83, 94).

[211] N. Jackson and J. Muthuswamy. *Artificial dural sealant that allows multiple penetrations of implantable brain probes.* J. Neurosci. Methods **171** (2008) 147. (page 85).

[212] Jason E Chung, Hannah R Joo, et al. *A polymer probe-based system for high density, long-lasting electrophysiological recordings across distributed neuronal circuits.* bioRxiv (2018). (pages 85, 97).

[213] ????? (2017). No citations.

[214] M. A. Hopcroft, W. D. Nix, and T. W. Kenny. *What is the young's modulus of silicon?* Journal of Microelectromechanical Systems **19** (2010) 229. (page 92).

[215] K. J. Paralikar and R. S. Clement. *Collagenase-aided intracortical microelectrode array insertion: Effects on insertion force and recording performance.* IEEE Transactions on Biomedical Engineering **55** (2008) 2258. (pages 93, 96).

[216] R. Normann P. Rousche. *A method for pneumatically inserting an array of penetrating electrodes into cortical tissue.* Annals of Biomedical Engineering **20** (1992) 413. (pages 93, 95).

[217] T. Escamilla-Mackert, N. B. Langhals, T. D. Kozai, and D. R. Kipke. *Insertion of a three dimensional silicon microelectrode assembly through a thick meningeal membrane.* Conf. Proc. IEEE Eng. Med. Biol. Soc. **2009** (2009) 1616. (page 93).

[218] R. L. Rennaker, S. Street, A. M. Ruyle, and A. M. Sloan. *A comparison of chronic multi-channel cortical implantation techniques: manual versus mechanical insertion.* J. Neurosci. Methods **142** (2005) 169. (page 94).

[219] B. J. Black, A. Kanneganti, A. Joshi-Imre, R. Rihani, B. Chakraborty, J. Abbott, J. J. Pancrazio, and S. F. Cogan. *Chronic recording and electrochemical performance of utah microelectrode arrays implanted in rat motor cortex.* J Neurophysiol. **120** (2018) 2083. (page 95).

[220] R. Fiath, A. L. Marton, F. Matyas, D. Pinke, G. Marton, K. Toth, and I. Ulbert. *Slow insertion of silicon probes improves the quality of acute neuronal recordings.* Sci. Rep. **9** (2019) 111. (page 96).

[221] G. Buzsáki. *Large-scale recording of neuronal ensembles.* Nat. Neurosci. *7*(5) (2004) 446. (page 97).

[222] I. H. Stevenson and K. P. Kording. *How advances in neural recording affect data analysis.* Nat. Neurosci. **14** (2011) 139. (page 97).

[223] H. Cohen, N. J. & Eichenbaum. *Memory, amnesia, and the hippocampal system.* MIT Press (1993). (pages 99, 101).

[224] L. R. Squire and P. Alvarez. *Retrograde amnesia and memory consolidation: a neurobiological perspective.* Curr. Opin. Neurobiol. **5** (1995) 169. (page 123).

[225] E. R. Kandel, Y. Dudai, and M. R. Mayford. *The molecular and systems biology of memory.* Cell **157** (2014) 163. (page 101, 101).

[226] D. L. Schacter. *Forgotten ideas, neglected pioneers: Richard semon and the story of memory.* Psychology Press (2012). (pages 99, 99, 102, 124).

[227] C. H. Vanderwolf. *Hippocampal electrical activity and voluntary movement in the rat.* Electroencephalogr. Clin. Neurophysiol. **26** (1969) 407. (pages 99, 104).

[228] L. O'Keefe, J. & Nadel. *The hippocampus as a cognitive map.* Oxford Univ. Press (1978). (pages 104, 104, 104, 104, 124).

[229] L. M. Kay, K. & Frank. *Three brain states in the hippocampus and cortex.* Hippocampus **29** (2018) 184. (pages 99, 104, 106, 121, 121).

[230] G. Buzsáki. *Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning.* Hippocampus **25** (2015) 1073. (pages 100, 104, 105, 105, 106, 112, 115, 123).

[231] D. J. Foster. *Replay comes of age.* Annu. Rev. Neurosci. **40** (2017) 581. (pages 105, 124, 124).

[232] L. A. Atherton, D. Dupret, and J. R. Mellor. *Memory trace replay: the shaping of memory consolidation by neuromodulation.* Trends Neurosci. **38** (2015) 560. (pages 102, 107, 111, 123).

[233] S. P. Tang, W. & Jadhav. *Sharp-wave ripples as a signature of hippocampal-prefrontal reactivation for memory during sleep and waking states.* Neurobiol. Learn. Mem. **160** (2018) 11. No citations.

[234] B. E. Pfeiffer. *The content of hippocampal "replay".* Hippocampus (2017) 6. (page 124).

[235] H. F. Ólafsdottir, D. Bush, and C. Barry. *The role of hippocampal replay in memory and planning.* Curr. Biol. **28** (2018) R37. No citations.

[236] G. Rothschild. *The transformation of multi-sensory experiences into memories during sleep.* Neurobiol. Learn. Mem. **160** (2018) 58. (page 120).

[237] M. F. Carr, S. P. Jadhav, and L. M. Frank. *Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval.* Nat. Neurosci. **14** (2011) 147. (pages 105, 123, 123).

[238] J. Y. Yu and L. M. Frank. *Hippocampal-cortical interaction in decision making.* Neurobiol. Learn. Mem. **117C** (2015) 34. (page 123).

[239] G. Girardeau and M. Zugaro. *Hippocampal ripples and memory consolidation*. Curr. Opin. Neurobiol. **21** (2011) 452. (page 105).

[240] D. K. Roumis and L. M. Frank. *Hippocampal sharp-wave ripples in waking and sleeping states*. Curr. Opin. Neurobiol. **35** (2015) 6. (pages 100, 111, 115, 123, 127).

[241] L. R. Squire. *Mechanisms of memory*. Science **232** (1986) 1612. (page 101).

[242] W. B. Scoville and B. Milner. *Loss of recent memory after bilateral hippocampal lesions*. J. Neurol. Neurosurg. Psychiatry **20** (1957) 11. (page 101).

[243] N. J. Cohen, R. A. Poldrack, and H. Eichenbaum. *Memory for items and memory for relations in the procedural/declarative memory framework*. Memory **5** (1997) 131. (page 101).

[244] E. Tulving. *Episodic memory: from mind to brain*. Annu. Rev. Psychol. **53** (2002) 1. (pages 101, 115).

[245] H. Eichenbaum, P. Dudchenko, E. Wood, M. Shapiro, and H. Tanila. *The hippocampus, memory, and place cells: is it spatial memory or a memory space?* Neuron **23** (1999) 209. (page 101).

[246] S. Wirth et al. *Single neurons in the monkey hippocampus and learning of new associations*. Science **300** (2003) 1578. (page 101).

[247] S. Wirth et al. *Trial outcome and associative learning signals in the monkey hippocampus*. Neuron **61** (2009) 930. No citations.

[248] U. Rutishauser, A. N. Mamelak, and E. M. Schuman. *Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex*. Neuron **49** (2006) 805. (page 101).

[249] U. Rutishauser, I. B. Ross, A. N. Mamelak, and E. M. Schuman. *Human memory strength is predicted by theta-frequency phase-locking of single neurons.* Nature **464** (2010) 903. (page 101).

[250] J. Winson. *Loss of hippocampal theta rhythm results in spatial memory deficit in the rat.* Science **201** (1978) 160. No citations.

[251] M. E. Hasselmo, C. Bodelon, and B. P. Wyble. *A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning.* Neural Comput. **14** (2002) 793. No citations.

[252] T. Spellman et al. *Hippocampal-prefrontal input supports spatial encoding in working memory.* Nature **522** (2015) 309. (page 101).

[253] P. W. Frankland and B. Bontempi. *The organization of recent and remote memories.* Nat. Rev. Neurosci. **6** (2005) 119. (pages 101, 103, 113).

[254] J. J. Kim and M. S. Fanselow. *Modality-specific retrograde amnesia of fear.* Science **256** (1992) 675. (pages 101, 102, 102).

[255] Y. Dudai. *The restless engram: consolidations never end.* Annu. Rev. Neurosci. **35** (2012) 227. (pages 101, 102).

[256] Roediger H. L. 3rd. & Tulving E. Dudai, Y. *Science of memory concepts.* Oxford Univ. Press (2007). (pages 102, 102, 123).

[257] Y. Dudai and M. Carruthers. *The janus face of mnemosyne.* Nature **434** (2005) 567. (page 102).

[258] R. G. Morris, P. Garrud, J. N. Rawlins, and J. O'Keefe. *Place navigation impaired in rats with hippocampal lesions.* Nature **297** (1982) 681. (page 102).

[259] M. J. Jutras and E. A. Buffalo. *Recognition memory signals in the macaque hippocampus.* Proc. Natl Acad. Sci. USA **107** (2010) 401. (page 102).

[260] G. Winocur, M. Moscovitch, and B. Bontempi. *Memory formation and long-term retention in humans and animals: convergence towards a transformation account of hippocampal-neocortical interactions*. Neuropsychologia **48** (2010) 2339. (pages 102, 103, 103).

[261] L. Nadel, A. Samsonovich, L. Ryan, and M. Moscovitch. *Multiple trace theory of human memory: computational, neuroimaging, and neuropsychological results*. Hippocampus **10** (2000) 352. (page 103).

[262] D. Hassabis, D. Kumaran, and E. A. Maguire. *Using imagination to understand the neural basis of episodic memory*. J. Neurosci. **27** (2007) 14365. (page 106).

[263] D. Hassabis, D. Kumaran, S. D. Vann, and E. A. Maguire. *Patients with hippocampal amnesia cannot imagine new experiences*. Proc. Natl Acad. Sci. USA **104** (2007) 1726. No citations.

[264] D. L. Schacter. *The future of memory: remembering, imagining, and the brain*. Neuron **76** (2012) 677. No citations.

[265] W. A. Suzuki and H. Eichenbaum. *The neurophysiology of memory*. Ann. NY Acad. Sci. **911** (2000) 175. No citations.

[266] L. R. Squire and S. Zola-Morgan. *The medial temporal lobe memory system*. Science **253** (1991) 1380. (page 102).

[267] G. Riedel. *Reversible neural inactivation reveals hippocampal participation in several memory processes*. Nat. Neurosci. **2** (1999) 898. (page 102).

[268] S. G. Anagnostaras, S. Maren, and M. S. Fanselow. *Temporally graded retrograde amnesia of contextual fear after hippocampal damage in rats: within-subjects examination*. J. Neurosci. **19** (1999) 1106. (page 102).

[269] M. J. Sekeres, G. Winocur, and M. Moscovitch. *The hippocampus and related neocortical structures in memory transformation.* Neurosci. Lett. **680** (2018) 39. (page 102).

[270] Y. Dudai. *The neurobiology of consolidations, or, how stable is the engram?* Annu. Rev. Psychol. **55** (2004) 51. (page 102).

[271] Y. Dudai, A. Karni, and J. Born. *The consolidation and transformation of memory.* Neuron **88** (2015) 20. (page 102).

[272] J. Lisman, K. Cooper, M. Sehgal, and A. J. Silva. *Memory formation depends on both synapse-specific modifications of synaptic strength and cell-specific increases in excitability.* Nat. Neurosci. **21** (2018) 309. (pages 102, 123).

[273] M. M. Poo et al. *What is memory? the present state of the engram.* BMC Biol. **14** (2016) 40. (pages 121, 123).

[274] A. Ben-Yakov, Y. Dudai, and M. R. Mayford. *Memory retrieval in mice and men.* Cold Spring Harb. Persp. Biol. **7** (2015) a021790. (pages 102, 113, 123, 124).

[275] J. L. McGaugh. *Memory—a century of consolidation.* Science **287** (2000) 248. (page 102).

[276] L. R. Squire, L. Genzel, J. T. Wixted, and R. G. Morris. *Memory consolidation.* Cold Spring Harb. Perspect. Biol. **7** (2015) a021766. (page 103).

[277] R. J. Sutherland. *Retrograde amnesia after hippocampal damage: recent versus remote memories in two tasks.* Hippocampus **11** (2001) 27. (page 103).

[278] S. J. Martin, L. Hoz, and R. G. Morris. *Retrograde amnesia: neither partial nor complete hippocampal lesions in rats result in preferential sparing of remote spatial memory, even after reminding.* Neuropsychologia **43** (2005) 609. (page 103).

[279] J. Born and I. Wilhelm. *System consolidation of memory during sleep.* Psychol. Res. **76** (2012) 192. (pages 103, 105, 118).

[280] S. A. Josselyn and P. W. Frankland. *Memory allocation: mechanisms and function.* Annu. Rev. Neurosci. **41** (2018) 389. (pages 103, 113).

[281] R. W. Semon. *Mnemic psychology* (1923). (pages 103, 124).

[282] G. Buzsáki, L. W. Leung, and C. H. Vanderwolf. *Cellular bases of hippocampal eeg in the behaving rat.* Brain Res. **287** (1983) 139. (page 104, 104, 104).

[283] G. Buzsáki. *Hippocampal sharp waves - their origin and significance.* Brain Res. **398** (1986) 242. No citations.

[284] D. Sullivan. *Relationships between hippocampal sharp waves, ripples, and fast gamma oscillation: influence of dentate and entorhinal cortical activity.* J. Neurosci. **31** (2011) 8605. No citations.

[285] S. S. Suzuki and G. K. Smith. *Spontaneous eeg spikes in the normal hippocampus. v. effects of ether, urethane, pentobarbital, atropine, diazepam and bicuculline.* Electroencephalogr. Clin. Neurophysiol. **70** (1988) 84. (page 104, 104).

[286] A. Ylinen et al. *Sharp wave-associated high-frequency oscillation (200 hz) in the intact hippocampus: network and intracellular mechanisms.* J. Neurosci. **15** (1995) 30. (page 104).

[287] A. Oliva, A. Fernandez-Ruiz, G. Buzsáki, and A. Berenyi. *Role of hippocampal ca2 region in triggering sharp-wave ripples.* Neuron **91** (2016) 1342. (pages 104, 123).

[288] T. Sasaki. *Dentate network activity is necessary for spatial working memory by supporting ca3 sharp-wave ripple generation and prospective firing of ca3 neurons.* Nat. Neurosci. **21** (2018) 258. (pages 104, 118).

[289] D. G. Amaral and M. P. Witter. *The three-dimensional organization of the hippocampal formation: a review of anatomical data.* Neuroscience **31** (1989) 571. (page 104).

[290] X. G. Li, P. Somogyi, A. Ylinen, and G. Buzsáki. *The hippocampal ca3 network: an in vivo intracellular labeling study.* J. Comp. Neurol. **339** (1994) 181. (page 104).

[291] M. Valero. *Mechanisms for selective single-cell reactivation during offline sharp-wave ripples and their distortion by fast ripples.* Neuron **94** (2017) 1234. (pages 104, 112, 112).

[292] S. S. Suzuki and G. K. Smith. *Spontaneous eeg spikes in the normal hippocampus. i. behavioral correlates, laminar profiles and bilateral synchrony.* Electroencephalogr. Clin. Neurophysiol. **67** (1987) 348. (page 104).

[293] M. Jouvet, F. Michel, and J. Courjon. *Electric activity of the rhinencephalon during sleep in cats.* C. R. Seances Soc. Biol. Fil. **153** (1959) 101. No citations.

[294] G. Buzsáki, Z. Horvath, R. Urioste, J. Hetke, and K. Wise. *High-frequency network oscillation in the hippocampus.* Science **256** (1992) 1025. No citations.

[295] D. F. English. *Excitation and inhibition compete to control spiking during hippocampal ripples: intracellular study in behaving mice.* J. Neurosci. **34** (2014) 16509. No citations.

[296] E. Stark. *Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations.* Neuron **83** (2014) 467. (page 104).

[297] J. Csicsvari, H. Hirase, A. Czurko, A. Mamiya, and G. Buzsáki. *Fast network oscillations in the hippocampal ca1 region of the behaving rat.* J. Neurosci. **19** (1999) RC20. (page 104).

[298] S. Cheng and L. M. Frank. *New experiences enhance coordinated neural activity in the hippocampus*. Neuron **57** (2008) 303. (pages 104, 104, 107, 107, 121).

[299] B. E. Pfeiffer and D. J. Foster. *Hippocampal place-cell sequences depict future paths to remembered goals*. Nature **497** (2013) 74. (pages 105, 113, 114, 115, 115, 115, 117, 121, 122, 123).

[300] K. Eguchi and T. Satoh. *Relationship between positive sharp wave bursts and unitary discharges in the cat hippocampus during slow wave sleep*. Physiol. Behav. **40** (1987) 497. (page 106).

[301] N. Kanamori. *A spindle-like wave in the cat hippocampus: a novel vigilance level-dependent electrical activity*. Brain Res. **334** (1985). (page 106).

[302] N. Ulanovsky and C. F. Moss. *Hippocampal cellular and network activity in freely moving echolocating bats*. Nat. Neurosci. **10** (2007) 224. (page 106).

[303] M. S. Nokia, J. E. Mikkonen, M. Penttonen, and J. Wikgren. *Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning*. Front. Behav. Neurosci. **6** (2012) 84. (pages 106, 106, 110, 110, 121).

[304] M. S. Nokia, M. Penttonen, and J. Wikgren. *Hippocampal ripple-contingent training accelerates trace eyeblink conditioning and retards extinction in rabbits*. J. Neurosci. **30** (2010) 11486. (pages 106, 106, 110).

[305] N. K. Logothetis. *Hippocampal-cortical interaction during periods of subcortical silence*. Nature **491** (2012) 547. (pages 106, 106, 119, 120).

[306] T. K. Leonard and K. L. Hoffman. *Sharp-wave ripples in primates are enhanced near remembered visual objects*. Curr. Biol. **27** (2017) 257. (pages 106, 106, 121).

[307] J. F. Ramirez-Villegas, N. K. Logothetis, and M. Besserve. *Diversity of*

sharp-wave-ripple lfp signatures reveals differentiated brain-wide dynamical events. Proc. Natl Acad. Sci. USA **112** (2015) E6379. (pages 122, 123).

[308] W. E. Skaggs. *Eeg sharp waves and sparse ensemble unit activity in the macaque hippocampus.* J. Neurophysiol. **98** (2007) 898. (page 106).

[309] T. K. Leonard. *Sharp wave ripples during visual exploration in the primate hippocampus.* J. Neurosci. **35** (2015) 14771. (page 106).

[310] B. P. Staresina. *Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep.* Nat. Neurosci. **18** (2015) 1679. (pages 106, 118).

[311] N. Axmacher, C. E. Elger, and J. Fell. *Ripples in the medial temporal lobe are relevant for human memory consolidation.* Brain **131** (2008) 1806. (page 106, 106).

[312] Q. M. Van. *Cell type-specific firing during ripple oscillations in the hippocampal formation of humans.* J. Neurosci. **28** (2008) 6104. No citations.

[313] A. Bragin. *High-frequency oscillations in human brain.* Hippocampus **9** (1999) 137. (pages 106, 112).

[314] M. Shein-Idelson, J. M. Ondracek, H. P. Liaw, S. Reiter, and G. Laurent. *Slow waves, sharp waves, ripples, and rem in sleeping dragons.* Science **352** (2016) 590. (page 106).

[315] R. Vargas, H. Thorsteinsson, and K. A. Karlsson. *Spontaneous neural activity of the anterodorsal lobe and entopeduncular nucleus in adult zebrafish: a putative homologue of hippocampal sharp waves.* Behav. Brain Res. **229** (2012) 10. (page 106).

[316] R. J. Staba. *High-frequency oscillations recorded in human medial temporal lobe during sleep.* Ann. Neurol. **56** (2004) 108. (page 106).

[317] Z. Clemens. *Temporal coupling of parahippocampal ripples, sleep spindles and slow oscillations in humans.* Brain **130** (2007) 2868. (page 106).

[318] N. K. Logothetis. *Neural-event-triggered fmri of large-scale neural networks*. Curr. Opin. Neurobiol. **31** (2015) 214. (page 106).

[319] R. Kaplan. *Hippocampal sharp-wave ripples influence selective activation of the default mode network*. Curr. Biol. **26** (2016) 686. (page 106, 106).

[320] G. Buzsáki and E. I. Moser. *Memory, navigation and theta rhythm in the hippocampal-entorhinal system*. Nat. Neurosci. **16** (2013) 130. (page 106).

[321] M. M. Yartsev. *The emperor's new wardrobe: rebalancing diversity of animal models in neuroscience research*. Science **358** (2017) 466. (page 106).

[322] O. Talakoub, A. Gomez Palacio Schjetnan, T. A. Valiante, M. R. Popovic, and K. L. Hoffman. *Closed-loop interruption of hippocampal ripples through fornix stimulation in the non-human primate*. Brain Stimul. **9** (2016) 911. (page 106).

[323] G. Buzsáki. *Two-stage model of memory trace formation: a role for "noisy" brain states*. Neuroscience **31** (1989) 551. (pages 105, 108, 113).

[324] E. L. Thorndike. *The fundamentals of learning*. Teachers College Bureau of Publications (1932). (page 105).

[325] A. E. Papale, M. C. Zielinski, L. M. Frank, S. P. Jadhav, and A. D. Redish. *Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making*. Neuron **92** (2016) 975. (pages 107, 121, 124).

[326] A. D. Redish. *Vicarious trial and error*. Nat. Rev. Neurosci. **17** (2016) 147. (pages 107, 121).

[327] M. E. Speer, J. P. Bhanji, and M. R. Delgado. *Savoring the past: positive memories evoke value representations in the striatum*. Neuron **84** (2014) 847. (page 107).

[328] U. Wagner, S. Gais, H. Haider, R. Verleger, and J. Born. *Sleep inspires insight*. Nature **427** (2004) 352. (pages 107, 127).

[329] J. O'Neill, T. Senior, and J. Csicsvari. *Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior*. Neuron **49** (2006) 143. (page 107).

[330] D. V. Wang. *Mesopontine median raphe regulates hippocampal ripple oscillation and memory consolidation*. Nat. Neurosci. **18** (2015) 728. (pages 107, 110, 112).

[331] M. Vandecasteele. *Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus*. Proc. Natl Acad. Sci. USA **111** (2014) 13535. (pages 107, 112).

[332] M. G. Giovannini. *Effects of novelty and habituation on acetylcholine, gaba?, and glutamate release from the frontal cortex and hippocampus of freely moving rats*. Neuroscience **106** (2001) 43. (page 107).

[333] Y. Liu, S. S. McAfee, and D. H. Heck. *Hippocampal sharp-wave ripples in awake mice are entrained by respiration*. Sci. Rep. **7** (2017) 8950. (page 107).

[334] M. P. Karlsson and L. M. Frank. *Network dynamics underlying the formation of sparse, informative representations in the hippocampus*. J. Neurosci. **28** (2008) 14271. (page 107).

[335] O. Eschenko, W. Ramadan, M. Molle, J. Born, and S. J. Sara. *Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning*. Learn. Mem. **15** (2008) 222. No citations.

[336] J. O'Neill, T. J. Senior, K. Allen, J. R. Huxter, and J. Csicsvari. *Reactivation of experience-dependent cell assembly patterns in the hippocampus*. Nat. Neurosci. **11** (2008) 209. (page 107).

[337] J. C. Jackson, A. Johnson, and A. D. Redish. *Hippocampal sharp waves and*

reactivation during awake states depend on repeated sequential experience. J. Neurosci. **26** (2006) 12415. (page 107).

[338] A. C. Singer and L. M. Frank. *Rewarded outcomes enhance reactivation of experience in the hippocampus.* Neuron **64** (2009) 910. (pages 107, 108).

[339] D. J. Foster and M. A. Wilson. *Reverse replay of behavioural sequences in hippocampal place cells during the awake state.* Nature **440** (2006) 680. (pages 108, 113, 114, 114, 115, 121).

[340] R. C. O'Reilly and J. W. Rudy. *Computational principles of learning in the neocortex and hippocampus.* Hippocampus **10** (2000) 389. (page 108).

[341] D. Marr. *Simple memory: a theory for archicortex.* Phil. Trans. R. Soc. Lond. B. **262** (1971) 23. (pages 108, 118).

[342] V. Ego-Stengel and M. A. Wilson. *Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat.* Hippocampus **20** (2010) 1. (pages 108, 123).

[343] G. Girardeau, K. Benchenane, S. I. Wiener, G. Buzsáki, and M. B. Zugaro. *Selective suppression of hippocampal ripples impairs spatial memory.* Nature Neurosci. **12** (2009) 1222. (pages 108, 123).

[344] T. Nakashiba, D. L. Buhl, T. J. McHugh, and S. Tonegawa. *Hippocampal ca3 output is crucial for ripple-associated reactivation and consolidation of memory.* Neuron **62** (2009) 781. (page 108).

[345] Y. Novitskaya, S. J. Sara, N. K. Logothetis, and O. Eschenko. *Ripple-triggered stimulation of the locus coeruleus during post-learning sleep disrupts ripple/spindle coupling and impairs memory consolidation.* Learn. Mem. **23** (2016) 238. (pages 108, 112).

[346] G. Girardeau, A. Cei, and M. Zugaro. *Learning-induced plasticity regulates hippocampal sharp wave-ripple drive.* J. Neurosci. **34** (2014) 5176. (pages 110, 127).

[347] N. Maingret, G. Girardeau, R. Todorova, M. Goutierre, and M. Zugaro. *Hippocampo-cortical coupling mediates memory consolidation during sleep.* Nat. Neurosci. **19** (2016) 959. (page 110).

[348] S. P. Jadhav, C. Kemere, P. W. German, and L. M. Frank. *Awake hippocampal sharp-wave ripples support spatial memory.* Science **336** (2012) 1454. (pages 110, 110, 112, 121, 123, 123).

[349] L. M. Jadhav, S. P. & Frank. *Memory replay in the hippocampus.* Space,Time and Memory in the Hippocampal Formation (2014) 351. (page 110).

[350] J. O'Keefe and J. Dostrovsky. *The hippocampus as a spatial map.preliminary evidence from unit activity in the freely-moving rat.* Brain Res. **34** (1971) 171. (page 110, 110).

[351] M.-.. B. Moser, D. C. Rowland, and E. I. Moser. *Place cells, grid cells, and memory.* Cold Spring Harb. Perspect. Biol. **7** (2015) a021808. (page 110).

[352] L. Roux, B. Hu, R. Eichler, E. Stark, and G. Buzsáki. *Sharp wave ripples during learning stabilize the hippocampal spatial map.* Nat. Neurosci. **20** (2017) 845. (pages 110, 123).

[353] G. M. Ven, S. Trouche, C. G. McNamara, K. Allen, and D. Dupret. *Hippocampal offline reactivation consolidates recently formed cell assembly patterns during sharp wave-ripples.* Neuron **92** (2016) 968. (page 111).

[354] K. A. Kovacs. *Optogenetically blocking sharp wave ripple events in sleep does not interfere with the formation of stable spatial representation in the ca1 area of the hippocampus.* PLoS One **11** (2016) e0164675. (page 111).

[355] J. H. Sadowski, M. W. Jones, and J. R. Mellor. *Sharp-wave ripples orchestrate the induction of synaptic plasticity during reactivation of place cell firing patterns in the hippocampus*. Cell Rep. **14** (2016) 1916. (pages 111, 124).

[356] C. J. Behrens, L. P. Boom, H. L. de, A. Friedman, and U. Heinemann. *Induction of sharp wave-ripple complexes in vitro and reorganization of hippocampal networks*. Nat. Neurosci. **8** (2005) 1560. (page 111).

[357] H. Norimoto. *Hippocampal ripples down-regulate synapses*. Science **359** (2018) 10.1126/science.aao0702. (page 111, 111).

[358] E. V. Lubenov and A. G. Siapas. *Decoupling through synchrony in neuronal circuits with propagation delays*. Neuron **58** (2008) 118. (page 111).

[359] G. Tononi and C. Cirelli. *Sleep function and synaptic homeostasis*. Sleep Med. Rev. **10** (2006) 49. (page 111).

[360] J. J. Palop and L. Mucke. *Network abnormalities and interneuron dysfunction in alzheimer disease*. Nat. Rev. Neurosci. **17** (2016) 777. (page 112).

[361] C. Alvarado-Rojas. *Different mechanisms of ripple-like oscillations in the human epileptic subiculum*. Ann. Neurol. **77** (2015) 281. (page 112).

[362] A. Bragin, I. Mody, C. L. Wilson, and J. Engel. *Local generation of fast ripples in epileptic brain*. J. Neurosci. **22** (2002) 2012. (page 112).

[363] A. K. Gillespie. *Apolipoprotein e4 causes age-dependent disruption of slow gamma oscillations during hippocampal sharp-wave ripples*. Neuron **90** (2016) 740. (page 112).

[364] J. Witton. *Disrupted hippocampal sharp-wave ripple-associated spike dynamics in a transgenic mouse model of dementia*. J. Physiol. **594** (2014) 4615. (page 112, 112).

[365] C. Altimus, J. Harrold, H. Jaaro-Peled, A. Sawa, and D. J. Foster. *Disordered ripples are a common feature of genetically distinct mouse models relevant to schizophrenia.* Mol. Neuropsychiatry **1** (2015) 52. (page 112, 112).

[366] J. P. Wiegand. *Age is associated with reduced sharp-wave ripple frequency and altered patterns of neuronal variability.* J. Neurosci. **36** (2016) 5650. (page 112, 112).

[367] J. L. Gerrard, S. N. Burke, B. L. McNaughton, and C. A. Barnes. *Sequence reactivation in the hippocampus is impaired in aged rats.* J. Neurosci. **28** (2008) 7883. (page 112, 112).

[368] R. Ul Haq. *Pretreatment with β-adrenergic receptor agonists facilitates induction of ltp and sharp wave ripple complexes in rodent hippocampus.* Hippocampus **26** (2016) 1486. (page 112).

[369] D. Ishikawa, N. Matsumoto, T. Sakaguchi, N. Matsuki, and Y. Ikegaya. *Operant conditioning of synaptic and spiking activity patterns in single hippocampal neurons.* J. Neurosci. **34** (2014) 5044. (page 112).

[370] O. Nicole. *Soluble amyloid beta oligomers block the learning-induced increase in hippocampal sharp wave-ripple rate and impair spatial memory formation.* Sci. Rep. **6** (2016) 22728. (page 112).

[371] S. M. Ciupek, J. Cheng, Y. O. Ali, H. C. Lu, and D. Ji. *Progressive functional impairments of hippocampal neurons in a tauopathy mouse model.* J. Neurosci. **35** (2015) 8118. No citations.

[372] J. Suh, D. J. Foster, H. Davoudi, M. A. Wilson, and S. Tonegawa. *Impaired hippocampal ripple-associated replay in a mouse model of schizophrenia.* Neuron **80** (2013) 484. (page 112).

[373] K. G. Phillips. *Decoupling of sleep-dependent cortical and hippocampal interactions in a neurodevelopmental model of schizophrenia*. Neuron **76** (2012) 526. (page 112, 112, 112).

[374] B. E. Pfeiffer and D. J. Foster. *Autoassociative dynamics in the generation of sequences of hippocampal place cells*. Science **349** (2015) 180. (page 112).

[375] M. F. Carr, M. P. Karlsson, and L. M. Frank. *Transient slow gamma synchrony underlies hippocampal memory replay*. Neuron **75** (2012) 700. (page 112).

[376] H. F. Iaccarino. *Gamma frequency entrainment attenuates amyloid load and modifies microglia*. Nature **540** (2016) 230. (page 112).

[377] S. Tonegawa, X. Liu, S. Ramirez, and R. Redondo. *Memory engram cells have come of age*. Neuron **87** (2015) 918. (page 113).

[378] C. Pavlides and J. Winson. *Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes*. J. Neurosci. **9** (1989) 2907. (page 113).

[379] M. A. Wilson and B. L. McNaughton. *Reactivation of hippocampal ensemble memories during sleep*. Science **265** (1994) 676. (page 113).

[380] H. S. Kudrimoti, C. A. Barnes, and B. L. McNaughton. *Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and eeg dynamics*. J. Neurosci. **19** (1999) 4090. (page 113).

[381] D. Dupret, J. O'Neill, B. Pleydell-Bouverie, and J. Csicsvari. *The reorganization and reactivation of hippocampal maps predict spatial memory performance*. Nat. Neurosci. **13** (2010) 995. (pages 113, 113, 115, 123).

[382] J. Y. Yu et al. *Distinct hippocampal-cortical memory representations for experiences*

associated with movement versus immobility. eLife **6** (2017) e27621. (pages 113, 115, 118, 120, 122, 123).

[383] A. K. Lee and M. A. Wilson. *Memory of sequential experience in the hippocampus during slow wave sleep*. Neuron **36** (2002) 1183. (pages 113, 114, 115, 121).

[384] K. Diba and G. Buzsáki. *Forward and reverse hippocampal place-cell sequences during ripples*. Nat. Neurosci. **10** (2007) 1241. (page 114, 114).

[385] M. P. Karlsson and L. M. Frank. *Awake replay of remote experiences in the hippocampus*. Nat. Neurosci. **12** (2009) 913. (pages 113, 114, 114, 114, 114, 114, 115, 115, 119, 122).

[386] A. S. Gupta, M. A. Meer, D. S. Touretzky, and A. D. Redish. *Hippocampal replay is not a simple function of experience*. Neuron **65** (2010) 695. (pages 114, 117, 123).

[387] T. J. Davidson, F. Kloosterman, and M. A. Wilson. *Hippocampal replay of extended experience*. Neuron **63** (2009) 497. (pages 113, 113, 114, 114, 121, 123).

[388] X. Wu and D. J. Foster. *Hippocampal replay captures the unique topological structure of a novel environment*. J. Neurosci. **34** (2014) 6459. (page 113).

[389] W. Tang, J. D. Shin, L. M. Frank, and S. P. Jadhav. *Hippocampal-prefrontal reactivation during learning is stronger in awake compared with sleep states*. J. Neurosci. **37** (2017) 11789. (pages 113, 114, 118, 118, 119, 119, 120, 120).

[390] D. Ji and M. A. Wilson. *Coordinated memory replay in the visual cortex and hippocampus during sleep*. Nat. Neurosci. **10** (2007) 100. (pages 114, 119, 120).

[391] J. Jackson and A. D. Redish. *Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks*. Hippocampus **17** (2007) 1209. (page 114).

[392] H. F. Olafsdottir, F. Carpenter, and C. Barry. *Task demands predict a dynamic switch in the content of awake hippocampal replay.* Neuron **96** (2017) 925. (page 114).

[393] B. L. McNaughton, C. A. Barnes, and J. O'Keefe. *The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats.* Exp. Brain Res. **52** (1983) 41. (page 114).

[394] L. M. Frank, G. B. Stanley, and E. N. Brown. *Hippocampal plasticity across multiple days of exposure to novel environments.* J. Neurosci. **24** (2004) 7681. (page 114).

[395] R. E. Ambrose, B. E. Pfeiffer, and D. J. Foster. *Reverse replay of hippocampal place cells is uniquely modulated by changing reward.* Neuron **91** (2016) 1124. (pages 114, 127).

[396] N. S. Clayton, T. J. Bussey, and A. Dickinson. *Can animals recall the past and plan for the future?* Nat. Rev. Neurosci. **4** (2003) 685. (page 115).

[397] D. Panoz-Brown. *Replay of episodic memories in the rat.* Curr. Biol. **28** (2018) 1628. (page 115).

[398] A. C. Singer, M. F. Carr, M. P. Karlsson, and L. M. Frank. *Hippocampal swr activity predicts correct decisions during the initial learning of an alternation task.* Neuron **77** (2013) 1163. (page 115).

[399] C. T. Wu, D. Haggerty, C. Kemere, and D. Ji. *Hippocampal awake replay in fear memory retrieval.* Nat. Neurosci. **20** (2017) 571. (pages 115, 123).

[400] A. M. Wikenheiser and A. D. Redish. *The balance of forward and backward hippocampal sequences shifts across behavioral states.* Hippocampus **23** (2013) 22. (page 115).

[401] R. L. Buckner. *The role of the hippocampus in prediction and imagination.* Annu. Rev. Psychol. **61** (2010) 27. (page 117).

[402] G. Dragoi and S. Tonegawa. *Selection of preconfigured cell assemblies for representation of novel spatial experiences.* Phil. Trans. R. Soc. B **369** (2014) 20120522. (page 117).

[403] G. Dragoi and S. Tonegawa. *Preplay of future place cell sequences by hippocampal cellular assemblies.* Nature **469** (2011) 397. No citations.

[404] A. D. Grosmark and G. Buzsáki. *Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences.* Science **351** (2016) 1440. (page 117).

[405] D. Silva, T. Feng, and D. J. Foster. *Trajectory events across hippocampal place cells require previous experience.* Nat. Neurosci. **18** (2015) 1772. (page 117).

[406] C. D. Schwindel and B. L. McNaughton. *Hippocampal-cortical interactions and the dynamics of memory trace reactivation.* Prog. Brain Res. **193** (2011) 163. (page 118).

[407] J. Rissman and A. D. Wagner. *Distributed representations in memory: insights from functional brain imaging.* Annu. Rev. Psychol. **63** (2012) 101. (page 118).

[408] P. Rajasethupathy. *Projections from neocortex mediate top-down control of memory retrieval.* Nature **526** (2015) 653. (page 118).

[409] A. G. Siapas and M. A. Wilson. *Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep.* Neuron **21** (1998) 1123. (page 118, 118).

[410] M. Penttonen, A. Kamondi, A. Sik, L. Acsady, and G. Buzsáki. *Feed-forward and feed-back activation of the dentate gyrus in vivo during dentate spikes and sharp wave bursts.* Hippocampus **7** (1997) 437. (page 118).

[411] A. Bragin, G. Jando, Z. Nadasdy, M. Landeghem, and G. Buzsáki. *Dentate eeg spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat.* J. Neurophysiol. **73** (1995) 1691. (page 118).

[412] J. O'Neill, C. N. Boccara, F. Stella, P. Schoenenberger, and J. Csicsvari. *Superficial layers of the medial entorhinal cortex replay independently of the hippocampus.* Science **355** (2017) 184. (page 118).

[413] J. J. Chrobak and G. Buzsáki. *Selective activation of deep layer (v-vi) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat.* J. Neurosci. **14** (1994) 6160. (page 118).

[414] J. J. Chrobak and G. Buzsáki. *High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat.* J. Neurosci. **16** (1996) 3056. (page 118).

[415] H. F. Olafsdottir, F. Carpenter, and C. Barry. *Coordinated grid and place cell replay during rest.* Nat. Neurosci. **19** (2016) 792. (page 118).

[416] J. E. Chung, H. R. Joo, et al. *High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays.* Neuron (2019) 21. (page 118).

[417] C. M. Wierzynski, E. V. Lubenov, M. Gu, and A. G. Siapas. *State-dependent spike-timing relationships between hippocampal and prefrontal circuits during sleep.* Neuron **61** (2009) 587. (pages 118, 120).

[418] S. P. Jadhav, G. Rothschild, D. K. Roumis, and L. M. Frank. *Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events.* Neuron **90** (2016) 113. (pages 118, 119, 120, 120).

[419] M. Remondes and M. A. Wilson. *Slow-γ rhythms coordinate cingulate cortical responses to hippocampal sharp-wave ripples during wakefulness.* Cell Rep. **13** (2015) 1327. (page 118).

[420] D. V. Wang and S. Ikemoto. *Coordinated interaction between hippocampal sharp-wave ripples and anterior cingulate unit activity.* J. Neurosci. **36** (2016) 10663. (pages 118, 120).

[421] G. Rothschild, E. Eban, and L. M. Frank. *A cortical-hippocampal-cortical loop of information processing during memory consolidation.* Nat. Neurosci. **20** (2017) 251. (pages 118, 120, 120).

[422] A. A. Wilber, I. Skelin, W. Wu, and B. L. McNaughton. *Laminar organization of encoding and memory reactivation in the parietal cortex.* Neuron **95** (2017) 1406. (page 118).

[423] C. M. Pennartz. *The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples.* J. Neurosci. **24** (2004) 6446. (page 118).

[424] C. S. Lansink, P. M. Goltstein, J. V. Lankelma, B. L. McNaughton, and C. M. Pennartz. *Hippocampus leads ventral striatum in replay of place-reward information.* PLOS Biol. **7** (2009) e1000173. (page 118).

[425] S. N. Gomperts, F. Kloosterman, and M. A. Wilson. *Vta neurons coordinate with the hippocampal reactivation of spatial experience.* eLife **4** (2015) e05360. (pages 118, 119).

[426] J. L. Valdés, B. L. McNaughton, and J. M. Fellous. *Offline reactivation of experience-dependent neuronal firing patterns in the rat ventral tegmental area.* J. Neurophysiol. **114** (2015) 1183. (page 119).

[427] D. Khodagholy, J. N. Gelinas, and G. Buzsáki. *Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus.* Science **358** (2017) 369. (page 119).

[428] S. Ribeiro. *Long-lasting novelty-induced neuronal reverberation during slow-wave sleep in multiple forebrain areas.* PLOS Biol. **2** (2004) E24. (page 119).

[429] J. Y. Yu, D. F. Liu, A. Loback, I. Grossrubatscher, and L. M. Frank. *Specific hippocampal representations are linked to generalized cortical representations in memory.* Nature Communications (2018) 2209. (page 119).

[430] F. P. Battaglia, G. R. Sutherland, and B. L. McNaughton. *Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions.* Learn. Mem. **11** (2004) 697. (pages 119, 120).

[431] A. Sirota, J. Csicsvari, D. Buhl, and G. Buzsáki. *Communication between neocortex and hippocampus during sleep in rodents.* Proc. Natl Acad. Sci. USA **100** (2003) 2065. (page 120).

[432] Y. Isomura. *Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations.* Neuron **52** (2006) 871. (page 120).

[433] A. Peyrache, M. Khamassi, K. Benchenane, S. I. Wiener, and F. P. Battaglia. *Replay of rule-learning related neural patterns in the prefrontal cortex during sleep.* Nat. Neurosci. **12** (2009) 919. (page 120).

[434] D. Bendor and M. A. Wilson. *Biasing the content of hippocampal replay during sleep.* Nat. Neurosci. **15** (2012) 1439. (page 120).

[435] T. J. Teyler and J. W. Rudy. *The hippocampal indexing theory and episodic memory: updating the index.* Hippocampus **17** (2007) 1158. (page 120).

[436] K. Kay et al. *A hippocampal network for spatial coding during immobility and sleep.* Nature **531** (2016) 185. (page 120).

[437] L. L. Colgin. *Rhythms of the hippocampal network.* Nat. Rev. Neurosci. **17** (2016) 239. (page 121).

[438] G. Dragoi and G. Buzsáki. *Temporal encoding of place sequences by hippocampal cell assemblies*. Neuron **50** (2006) 145. (page 121).

[439] A. M. Wikenheiser and A. D. Redish. *Hippocampal theta sequences reflect current goals*. Nat. Neurosci. **18** (2015) 289. (page 121).

[440] A. Johnson and A. D. Redish. *Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point*. J. Neurosci. **27** (2007) 12176. (page 121).

[441] Y. Wang, Z. Roth, and E. Pastalkova. *Synchronized excitability in a network enables generation of internal neuronal sequences*. eLife **5** (2016) e20697. (page 121).

[442] J. Patel, E. W. Schomburg, A. Berenyi, S. Fujisawa, and G. Buzsáki. *Local generation and propagation of ripples along the septotemporal axis of the hippocampus*. J. Neurosci. **33** (2013) 17029. (page 121).

[443] D. Aronov, R. Nevers, and D. W. Tank. *Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit*. Nature **543** (2017) 719. (page 122).

[444] J. Lisman. *Viewpoints: how the hippocampus contributes to memory?, navigation and cognition*. Nat. Neurosci. **20** (2017) 1434. (page 122).

[445] Y. Ziv. *Long-term dynamics of ca1 hippocampal place codes*. Nat. Neurosci. **16** (2013) 264. (page 122).

[446] S. Thorpe, D. Fize, and C. Marlot. *Speed of processing in the human visual system*. Nature **381** (1996) 520. (page 123).

[447] M. N. Shadlen and D. Shohamy. *Decision making and sequential sampling from memory*. Neuron **90** (2016) 927. (page 123).

[448] D. Tse. *Schema-dependent gene activation and memory encoding in neocortex*. Science **333** (2011) 891. (page 123).

[449] S. A. Josselyn, S. Kohler, and P. W. Frankland. *Heroes of the engram.* J. Neurosci. **37** (2017) 4647. (page 123).

[450] J. O'Neill, B. Pleydell-Bouverie, D. Dupret, and J. Csicsvari. *Play it again: reactivation of waking experience and memory.* Trends Neurosci. **33** (2010) 220. (page 123).

[451] J. Csicsvari, J. O'Neill, K. Allen, and T. Senior. *Place-selective firing contributes to the reverse-order reactivation of ca1 pyramidal cells during sharp waves in open-field exploration.* Eur. J. Neurosci. **26** (2007) 704. (page 123).

[452] J. Yamamoto and S. Tonegawa. *Direct medial entorhinal cortex input to hippocampal ca1 is crucial for extended quiet awake replay.* Neuron **96** (2017) 217. No citations.

[453] M. Valero. *Determinants of different deep and superficial ca1 pyramidal cell dynamics during sharp-wave ripples.* Nat. Neurosci. **18** (2015) 1281. (page 123).

[454] R. K. Mishra, S. Kim, S. J. Guzman, and P. Jonas. *Symmetric spike timing-dependent plasticity at ca3-ca3 synapses optimizes storage and recall in autoassociative networks.* Nat. Commun. **7** (2016) 11552. (page 124).

[455] J. Lopez, K. Gamache, R. Schneider, and K. Nader. *Memory retrieval requires ongoing protein synthesis and nmda receptor activity-mediated ampa receptor trafficking.* J. Neurosci. **35** (2015) 2465. (page 124).

[456] G. Szapiro. *Molecular mechanisms of memory retrieval.* Neurochem. Res. **27** (2002) 1491. (page 124).

[457] Y. Dudai. *Reconsolidation: the advantage of being refocused.* Curr. Opin. Neurobiol. **16** (2006) 174. (page 124).

[458] K. Nader, G. E. Schafe, and J. E. Doux. *Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval.* Nature **406** (2000) 722. No citations.

[459] K. Nader. *Reconsolidation and the dynamic nature of memory.* Cold Spring Harb. Perspect. Biol. **7** (2015) a021782. No citations.

[460] S. J. Sara. *Retrieval and reconsolidation: toward a neurobiology of remembering.* Learn. Memory **7** (2000) 73. No citations.

[461] K. Nader. *Memory traces unbound.* Trends Neurosci. **26** (2003) 65. No citations.

[462] J. R. Misanin, R. R. Miller, and D. J. Lewis. *Retrograde amnesia produced by electroconvulsive shock after reactivation of a consolidated memory trace.* Science **160** (1968) 554. (page 124).

[463] B. Straube. *An overview of the neuro-cognitive processes involved in the encoding, consolidation, and retrieval of true and false memories.* Behav. Brain Funct. **8** (2012) 35. (page 124).

[464] J. S. Simons, J. R. Garrison, and M. K. Johnson. *Brain mechanisms of reality monitoring.* Trends Cogn. Sci. **21** (2017) 462. (page 124).

[465] A. D. Redish. *Beyond the cognitive map: From place cells to episodic memory.* MIT Press (1999). (page 124).

[466] R. M. French. *Catastrophic forgetting in connectionist networks.* Trends Cogn. Sci. **3** (1999) 128. (page 125, 125).

[467] M. E. Hasselmo. *Avoiding catastrophic forgetting.* Trends Cogn. Sci. **21** (2017) 407. (page 125).

[468] N. J. McCloskey, M. & Cohen. The Psychology of Learning and Motivation **24** (1989) 109. (page 125).

[469] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly. *Why there are complementary learning systems in the hippocampus and neocortex: insights from the*

successes and failures of connectionist models of learning and memory. Psychol. Rev. **102** (1995) 419. (page 125).

[470] D. Kumaran, D. Hassabis, and J. L. McClelland. *What learning systems do intelligent agents need? complementary learning systems theory updated.* Trends Cogn. Sci. **20** (2016) 512. (page 125).

[471] G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. *The "wake-sleep" algorithm for unsupervised neural networks.* Science **268** (1995) 1158. (page 125).

[472] V. Mnih. *Human-level control through deep reinforcement learning.* Nature **518** (2015) 529. (page 125).

[473] Quan J. Antonoglou I. & Silver D. Schaul, T. *Prioritized experience replay.* arXiv:1511:05952 (2015). (page 125).

[474] A. Pritzel et al. *Neural episodic control.* arXiv:1703.01988 (2017). (page 125).

[475] Charles Blundell et al. *Model-free episodic control.* arXiv:1606.04460 (2016). (page 125).

[476] D. Ciliberti and F. Kloosterman. *Falcon: a highly flexible open-source software for closed-loop neuroscience.* J. Neural Eng. **14** (2017) 045004. (page 127).

[477] X. Deng, D. F. Liu, M. P. Karlsson, L. M. Frank, and U. T. Eden. *Rapid classification of hippocampal replay content for real-time applications.* J. Neurophysiol. **116** (2016) 2221. (page 127).

Appendix: Dataset characterization

For the main cohort of rats presented in this thesis, I collected simultaneous behavioral and *in vivo* recordings of local field potentials and spikes from a hybrid tetrode drive/flexible polymer probe device (366 channels) targeted bilaterally to the dorsal hippocampus, nucleus accumbens, and orbitofrontal cortex from four rats performing the episodic memory confidence task presented in chapter 2. The following section describes an overview of the pre-processing required to prepare the behavioral and neural data, with particular focus on the challenge of realtime measurement of the invested time.

A.1 Epoch and trial counts

Data were collected in hour-long epochs from four rats in two rig rooms, with behavioral data recorded in two ways: in the stateScriptLog (SSL) written out by the stateScript behavior module of the data acquisition system Trodes, and the Digital In/Out (DIO) channels recorded by the main neural data recording module of the data acquisition system Trodes. Two types of epochs exist: those for which behavioral data were recorded without simultaneous neural data, and those for which behavioral and neural data were recorded simultaneously. All epochs necessarily generated an SSL, but only those for which neural

data were recorded also had a DIO. As a result, the SSL-sourced dataset is larger for three of four animals. For animal *Dolo*, some of the SSL were corrupted, resulting in a larger usable DIO-sourced dataset.

Despite constituting a larger data source, significant data corruption in the SSLs (see below) required that behavioral data be sourced from the DIOs. The final DIO-sourced epoch and trial counts, including only the subset of epochs equal or greater in length to the animal's typical continuous performance period, are listed below for the four animals in the order they were collected:

- Tub (room 1)
  - 3119 trials / 46 epochs (SSL)
  - 2978 trials / 42 epochs (DIO parsed *post-hoc*; 20-minute epoch length threshold)
- Sojo (room 1)
  - 4113 trials / 40 epochs (SSL)
  - 4111 trials / 40 epochs (DIO parsed *post-hoc*; 40-minute epoch length threshold)
- Dolo (room 2)
  - 4212 trials / 60 epochs (SSL)
  - 4369 trials / 61 epochs (DIO parsed *post-hoc*; 40-minute epoch length threshold)
- Rosa (room 2)
  - 3798 trials / 52 epochs (SSL)
  - 3662 trials / 49 epochs (DIO parsed *post-hoc*; 45-minute epoch length threshold)

In this task, every correct trial is programmed to be rewarded. For two animals, *Tub, Sojo*, the SSLs failed to record reward delivery periods for a subset of correct trials (*unrewarded*

corrects). For animal *Tub*, there were 410 unrewarded corrects. For animal *Sojo*, there were 236 unrewarded corrects. These were distributed across epochs. Based on the SSLs alone, it was initially unknown whether reward was actually delivered at these time periods. If it was not, this would constitute a serious error in the basic task logic. Independent data extraction from the DIOs revealed that reward was indeed delivered on every correct trial. DIO-sourced data, however, was only collected for epochs during which neural data were recorded and so constitutes a subset of the SSL-sourced data. The DIO-sourced dataset is additionally reduced because of various errors in the data acquisition software and fragility in the hardware design, that resulted in irrecoverable errors (*i.e.*, core dumps and segmentation faults) that crashed the main recording module and corrupted the DIO log but were survived by the SSL stream.

A.2 Timing noise in realtime behavioral task implementation

On each trial, a rat indicates his choice by poking at a choice port, and can then maintain the nosepoke position for as long as he chooses. The reward amount he receives on correct trials is a monotonic function of the length of time for which he maintains the poke position. On error trials, he receives no reward, regardless of the invested time. Because trials take place in a fixed-length epoch, it is possible to get more reward by preferentially investing more time on correct trials, and therefore advantageous for a subject to use its prediction of whether a trial will be correct. The relationship between invested time and reward amount is crucial to the task logic. Handling of the timing of the invested time duration is done in custom code (HRJ) in a combination of the stateScript behavior module of the Trodes recording system and Python.

Throughout the time investment duration - when the animal is in the nosepoke - the photodiode beam within the port is broken when the animal's nose is low within the port,

but when it is higher within the port or briefly withdrawn, the beam can re-form. Rats bobbed their heads slightly within the port, resulting in a series of short beam breaks and unbreaks (*i.e.*, beam breaks do not correspond 1:1 with the nosepoke position). To avoid detecting only slight head movements (see video), we instituted a brief grace period (800 ms for rats *Sojo, Tub, Dolo*; 700 ms for rat *Rosa*) during which the rat could move its head up or down, allowing the photodiode beam to reform, without ending the invested time duration. That is, as long as the rat broke the beam again within the grace period, or *trigger duration*, he was determined to still be investing time. When the trigger duration occurred following the last beam unbreak, the investment period was declared to have ended.

The time investment duration had to be calculated in realtime, on the order of tens of milliseconds, to deliver the appropriate reward amount. This was performed based on the state of the beam and a timer in the stateScript module. The same beam states were recorded in the DIO log, as were the reward delivery and other task changes based on the realtime-calculated invested time. After data collection, I verified that the reward delivered, based on the realtime calculation of invested time, was equal to that expected based on the *post-hoc* calculation of invested time from the DIO log of beam states.

In the majority of cases, realtime- and *post-hoc*-calculated invested times were equivalent. In both the SSL and DIO-sourced datasets, however, I found that for a minority of trials, the reward amount received did not correspond to the *post-hoc*-calculated invested time. Approximately one percent of trials delivered a reward amount that exceeded what was expected given the invested time, and another approximately one percent delivered a reward amount less than expected (Figure 1).

The difference between the delivered reward amount and the reward amount expected based on *post-hoc* parsing of the DIO poke log could possibly be accounted for by a discrepancy in the calculation of invested time. To investigate this possibility, I compared

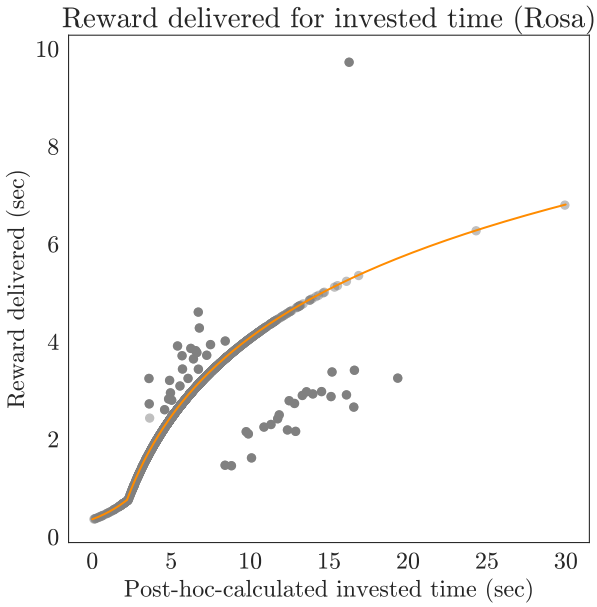

Figure 4: Reward delivered, from DIO milk pump log, as a function of *post-hoc*-calculated invested time from DIO beam break and unbreak log. Each gray point represents a trial. Orange line is the programmed reward amount to be delivered as a function of invested time. Gray points above the line correspond to trials for which the reward delivery amount exceeded what was coded, taking as input the *post-hoc*-calculated invested time. Gray points below the line correspond to trials for which the reward delivery amount was less than what was coded, taking as input the *post-hoc* calculated invested time.

the *post-hoc* and realtime-calculated invested times (Figure 5).

Trials for which the realtime-calculated invested time was long (light blue, Figure 5) received more reward than expected based on *post-hoc* parsing of invested time (light blue, above reward function in Figure 6), and trials for which the realtime-calculated invested time was short (purple, Figure 5) received less reward than expected based on *post-hoc* parsing of invested time (purple, below reward function in Figure 6). The discrepant reward amounts were accounted for fully by discrepant invested times (with exception of a few outliers due to parsing errors).

I hypothesized that slight timing errors in the realtime parsing could account for the discrepant invested time measurements or calculations. If a beam unbreak duration only slightly longer than the trigger interval passed, the *post-hoc* measurement would identify this as the end of the invested time duration, but a realtime timer running slowly would

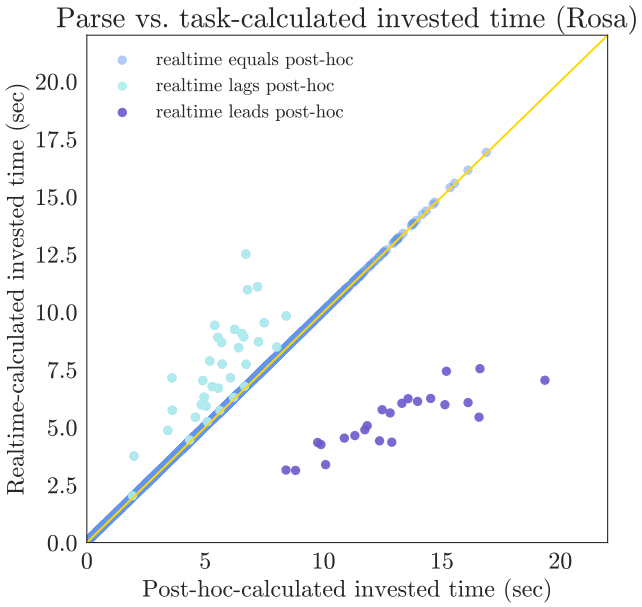

Figure 5: Realtime- versus *post-hoc*-calculated invested time. Each point represents a trial. Trials with equivalent realtime- and *post-hoc*-calculated invested time fall along the unity line (gold). Trials with realtime-calculated invested time greater than *post-hoc*-calculated invested time are light blue and above the line; trials with *post-hoc*-calculated invested time greater than realtime-calculated invested time are in purple and below the line.

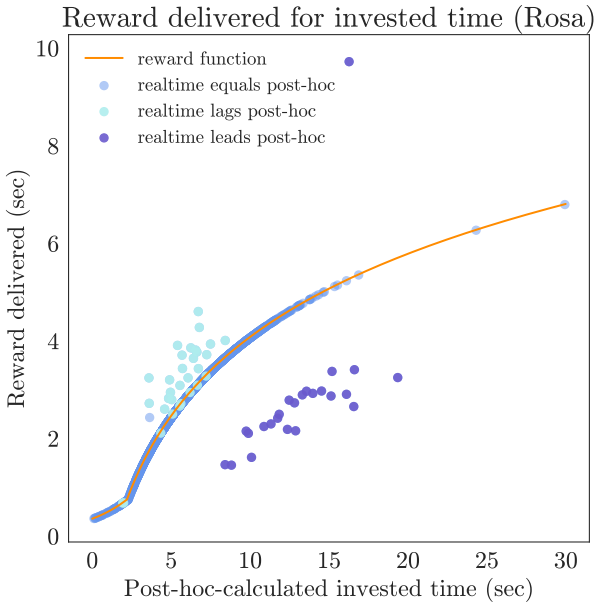

Figure 6: Reward delivered as a function of *post-hoc*-calculated invested time, as in Figure 4, with each point (trial) colored as in Figure 5. Programmed reward function in orange. Trials for which realtime-calculated invested time exceeds *post-hoc*-calculated invested time are light blue, and above the reward function line; trials for which *post-hoc*-calculated invested time exceeds realtime-invested time are purple, and below the reward function line.

not detect it. This false negative would result in an elongated invested time. Likewise, a beam un-break duration slightly shorter than the trigger interval could be measured in realtime to be equivalent to it if the stateScript timer were running slowly. This false positive would result in a truncated invested time.

The start of the invested time is always identified as the time at which the rat first pokes at the choice port. Even during the task, this timepoint is not determined until after the end of the invested time and does not rely on a stateScript timer, so is not a source of error. The realtime-calculated invested time is:

$$IT_{realttime} = end_{realtime} - start$$

For *post-hoc* calculation of invested time, the end of the invested time is detected by measuring in the DIO log the intervals between each unpoke and the subsequent poke and detecting the first interval that exceeds the trigger duration. The invested time is given by:

$$IT_{post} = end_{post} - start$$

The invested time durations thus depend on the timing precision of the realtime code, or the difference between the realtime-measured end time and the *post-hoc*-measured end time (with the latter taken as the true clock time):

$$\begin{aligned} IT_{post} - IT_{realtime} &= (end_{post} - start) - (end_{realtime} - start) \\ &= end_{post} - end_{realtime} \end{aligned}$$

After the first beam un-break interval equal to the trigger interval, the investment period is declared over, a 50-millisecond delay follows, then the choice outcome is revealed: if the choice was correct, the choice port will light; if it was in error, the back port will light. The timing of the choice or back port lighting can therefore be used to infer the $end_{realtime}$ that was measured in real time:

For a correct trial: $end_{realtime} = choiceon - delaylength - triggerinterval$

For an error trial: $end_{realtime} = backon - delaylength - triggerinterval$

Discrepant realtime- vs. *post-hoc*-parsed invested times were distributed across epochs and trial outcomes (Figure 7).

To identify false positive detections→truncated invested times, I calculated the inferred realtime invested time (T_{rt}) and, given the start time, identified the time of unpoke that resulted in the most similar invested time duration. I then determined the interval from this unpoke to the subsequent unpoke, which was measured in realtime to be equivalent to the trigger interval. These intervals were reliably slightly less than the trigger interval, suggesting that intervals slightly shorter than the trigger interval were sometimes measured in realtime as being equal to it, resulting in an earlier end to the invested time than should have occurred. The average of these durations was 4 milliseconds below the trigger interval (for representative rat *Rosa*).

To identify false negative detections→elongated invested times*, I also used the procedure above to identify the time of unpoke that resulted in the most similar invested time duration and determined the interval from this unpoke to the subsequent unpoke. As expected, these intervals were reliably greater than the trigger interval. I next used the same procedure, but taking *post-hoc* invested time duration as input, to identify the corresponding time of unpoke and the interval to the subsequent poke. Reliably, these intervals were all only slightly more than the trigger interval. This is consistent with the expectation that the realtime code may have mis-timed these intervals as being shorter than the trigger interval though they were in fact longer. The average of these durations was 3 milliseconds over the trigger interval (for representative rat *Rosa*).

Figure 8 shows the distribution of differences from the trigger interval for false positives→truncated investment times and false negatives→elongated investment times. This occurred with two different threshold values (rat *Rosa* = 700 ms, all others = 800 ms)

*NB: It is somewhat counterintuitive that: (i) mis-measurement of a headlift duration as shorter than it was (false negative) results in a longer invested time duration; (ii) mis-measurement of a headlift duration as longer than it was (false positive) results in a shorter invested time duration.

and in two different rig rooms with different hardware and software setups.

As expected, truncated invested times occurred when there was an earlier headlift duration only slightly shorter than the trigger interval that was detected in realtime to be equivalent to the trigger interval. Long invested times occurred when an earlier headlift duration only slightly longer than the trigger duration or more was not detected, the rat continued to invest time, and a later headlift duration longer than the trigger interval was instead detected (Figure 9).

The relative symmetry between false positive and false negative errors indicates a temporal resolution on the order of milliseconds for the timer itself running in stateScript. That false negative detections that occur more frequently (i) and are longer (ii) could be consistent with a slow stateScript clock due to additional delays, either inherent in the stateScript module itself or between the stateScript and python modules in handling the task logic. This is not seen for every rat, however.

The total fraction of false negative and false positive detections, and their percentage of total, is:

- Tub: 22 false positives (0.7 percent of total); 12 false negatives (0.4 percent of total)
- Sojo: 55 false positives (1.3 percent of total); 74 false negatives (1.8 percent of total)
- Dolo: 27 false positives (0.6 percent of total); 45 false negatives (1.0 percent of total)
- Rosa: 35 false positives (1.0 percent of total); 24 false negatives (0.7 percent of total)

As it is unlikely that rats can discern temporal differences on the order of milliseconds, they likely could not detect when the task treated an unpoke interval as order-milliseconds shorter or longer than the actual time. Indeed, there was no indication of this in behavior,

for which every trial was observed. All analyses were based on realtime-parsed invested times.

Individual rat noise characterization

The same plots shown for representative rat *Rosa* above are included for each rat in this section.

A.3 Rat: *Tub*

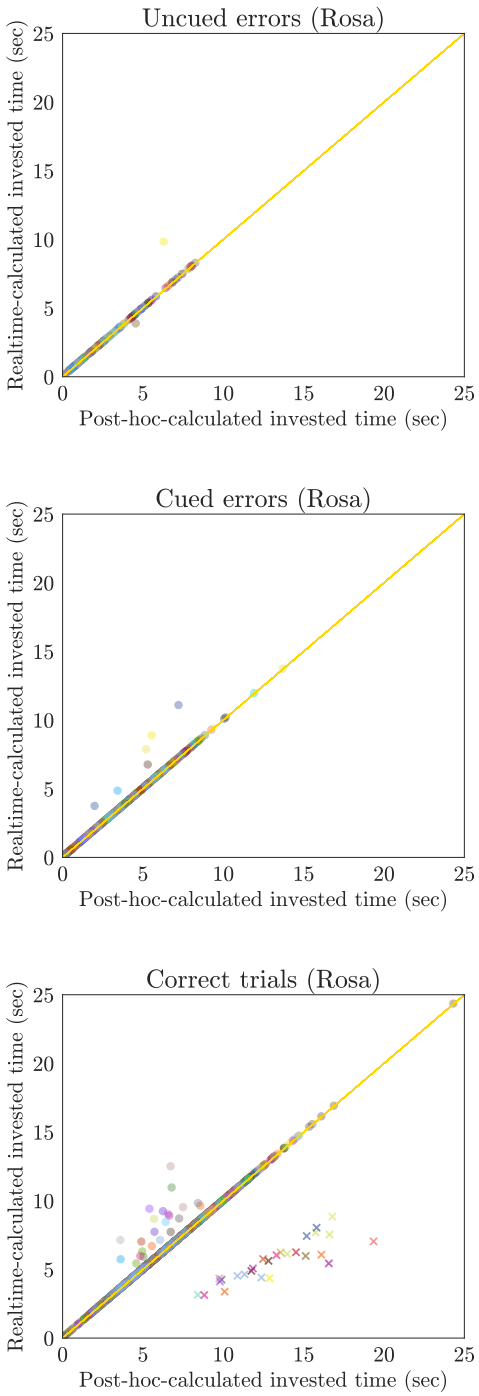

Figure 7: Discrepant realtime- vs. *post-hoc*-parsed invested times were distributed across epochs and trial outcomes. Each color represents the trials of a different epoch. For a subset of correct trials, the realtime-determined invested time was based on the *post-hoc*-parsed reward amount, and the trial is marked with an *x*; for all other trials, the realtime-determined invested time was determined based on the posthoc-parsed cue/back-on time, and the trial is marked with an *o*. Unity line in gold.

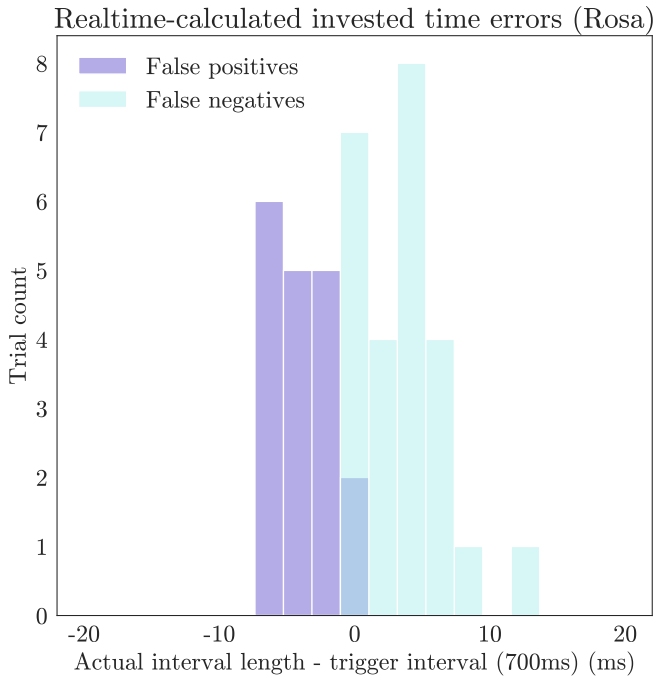

Figure 8: Distribution of differences from trigger interval for false positives (purple) and false negatives (light blue).

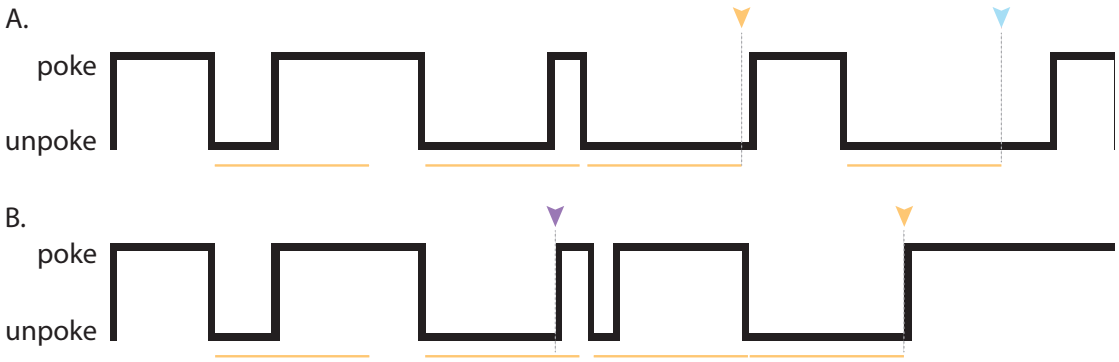

Figure 9: A. False negative detection of unpoke interval end. A series of poke (beam break) and unpoke (beam unbreak) intervals after the rat has nose-poked at a choice port. A timer begins at the start of every unpoke period, and if the length of the unpoke period exceeds that of the trigger interval (orange horizontal line), the rat is declared to be 'gone' and the investment period to have ended. The first two unpoke intervals (from left) are shorter than the trigger interval. The third is longer than the trigger interval, and *post-hoc* calculation of invested time takes this unpoke interval as the end of the invested time (orange arrow). In realtime, however, a false negative error was made, and the third unpoke interval was measured to be shorter than it actually was, and insufficient to trigger the end of the invested time. Note that this occurred because the unpoke interval was only slightly greater than the trigger interval. The fourth unpoke interval was also longer than the trigger interval, and was detected in realtime as the end of the investment period (light blue arrow). B. False positive detection of unpoke interval end. The first unpoke interval is shorter than the trigger interval. The second is also shorter than the trigger interval, but was measured in realtime as being equal to it (purple arrow) and was taken as the end of the investment period. Note that this occurred because the unpoke interval was only slightly shorter than the trigger interval. The third unpoke interval was shorter than the trigger interval; the fourth was longer than it, and was detected *post-hoc* as the end to the investment period.

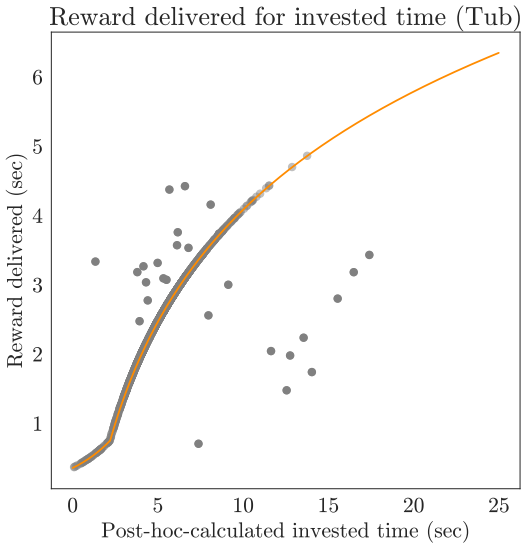

Figure 10: Reward delivered, from DIO milk pump log, as a function of *post-hoc*-calculated invested time from DIO beam break and unbreak log. Each gray point represents a trial. Orange line is the programmed reward amount to be delivered as a function of invested time. Gray points above the line correspond to trials for which the reward delivery amount exceeded what was coded, taking as input the post-hoc-calculated invested time. Gray points below the line correspond to trials for which the reward delivery amount was less than what was coded, taking as input the *post-hoc* calculated invested time.

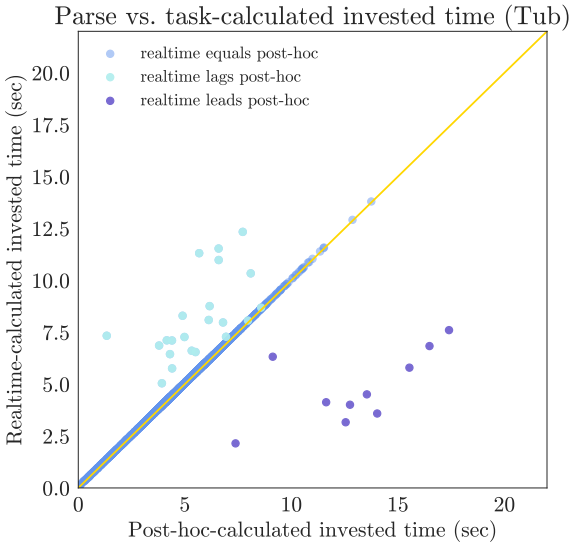

Figure 11: Realtime- versus *post-hoc*-calculated invested time. Each point represents a trial. Trials with equivalent realtime- and *post-hoc*-calculated invested time fall along the unity line (gold). Trials with realtime-calculated invested time greater than *post-hoc*-calculated invested time are light blue and above the line; trials with *post-hoc*-calculated invested time greater than realtime-calculated invested time are in purple and below the line.

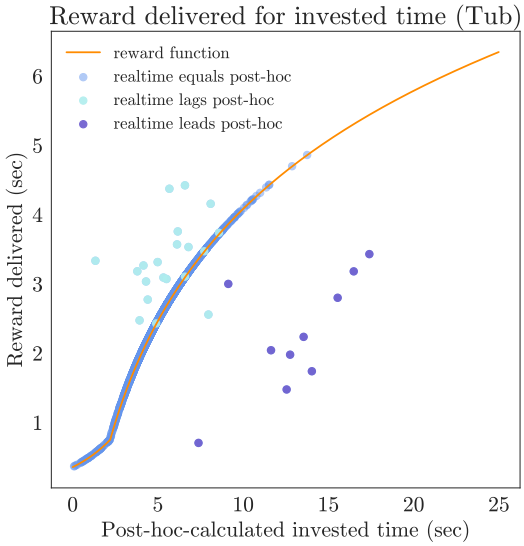

Figure 12: Reward delivered as a function of *post-hoc*-calculated invested time, as in Figure 10, with each point (trial) colored as in Figure 11. Programmed reward function in orange. Trials for which realtime-calculated invested time exceeds *post-hoc*-calculated invested time are light blue, and above the reward function line; trials for which *post-hoc*-calculated invested time exceeds realtime-invested time are purple, and below the reward function line.

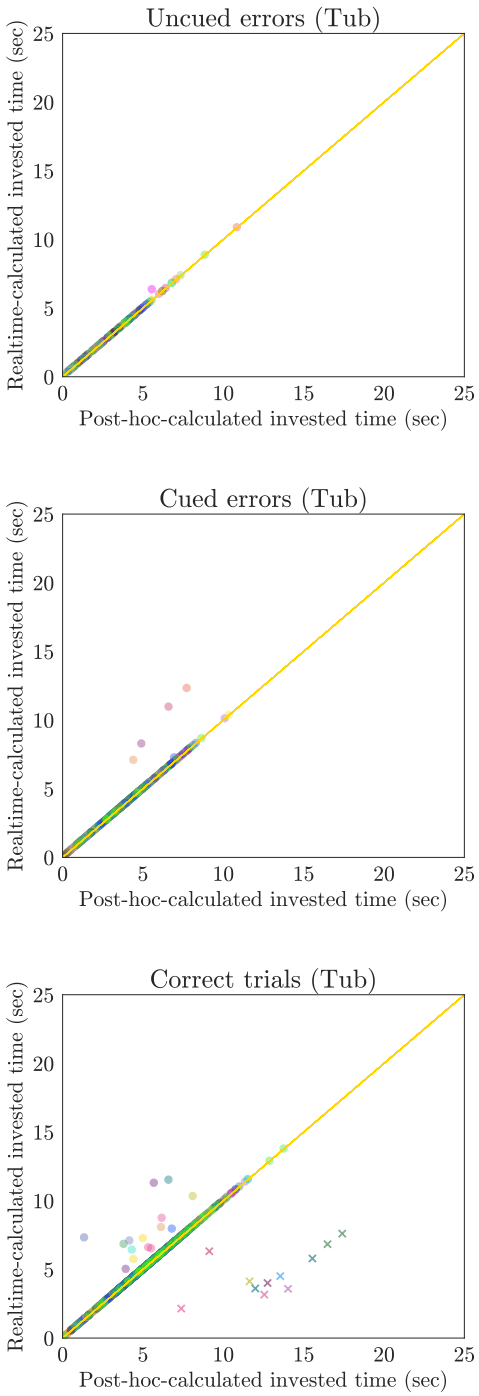

Figure 13: Discrepant realtime- vs. *post-hoc*-parsed invested times were distributed across epochs and trial outcomes. Each color represents the trials of a different epoch. For a subset of correct trials, the realtime-determined invested time was based on the *post-hoc*-parsed reward amount, and the trial is marked with an *x*; for all other trials, the realtime-determined invested time was determined based on the posthoc-parsed cue/back-on time, and the trial is marked with an *o*. Unity line in gold.

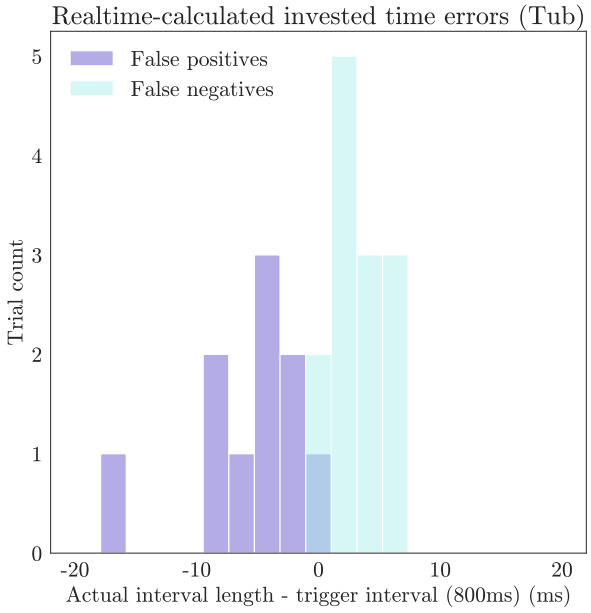

Figure 14: Distribution of differences from trigger interval for false positives (purple) and false negatives (light blue).

A.4 Rat: *Sojo*

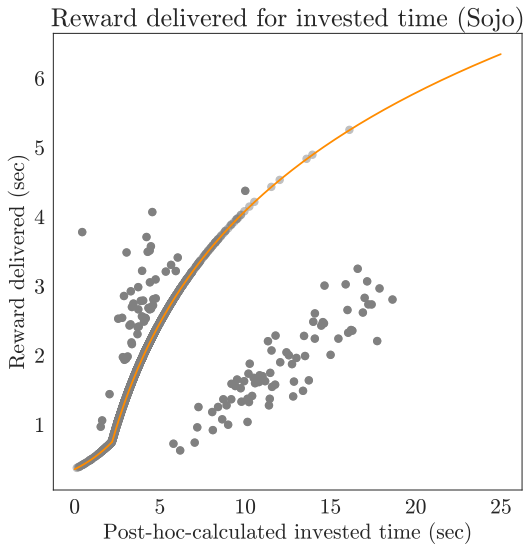

Figure 15: Reward delivered, from DIO milk pump log, as a function of *post-hoc*-calculated invested time from DIO beam break and unbreak log. Each gray point represents a trial. Orange line is the programmed reward amount to be delivered as a function of invested time. Gray points above the line correspond to trials for which the reward delivery amount exceeded what was coded, taking as input the *post-hoc*-calculated invested time. Gray points below the line correspond to trials for which the reward delivery amount was less than what was coded, taking as input the *post-hoc* calculated invested time.

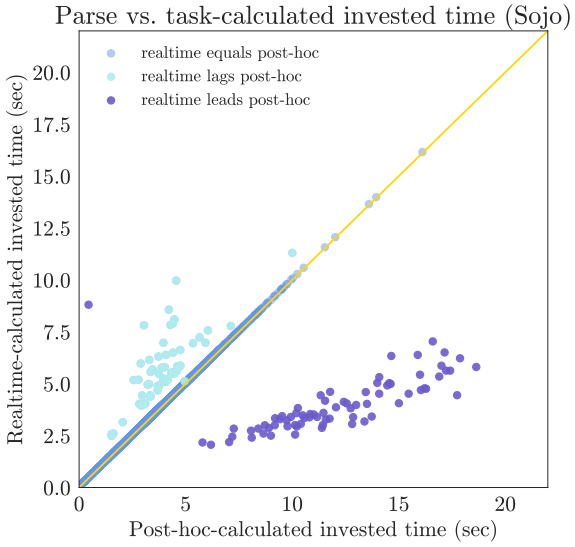

Figure 16: Realtime- versus *post-hoc*-calculated invested time. Each point represents a trial. Trials with equivalent realtime- and *post-hoc*-calculated invested time fall along the unity line (gold). Trials with realtime-calculated invested time greater than *post-hoc*-calculated invested time are light blue and above the line; trials with *post-hoc*-calculated invested time greater than realtime-calculated invested time are in purple and below the line.

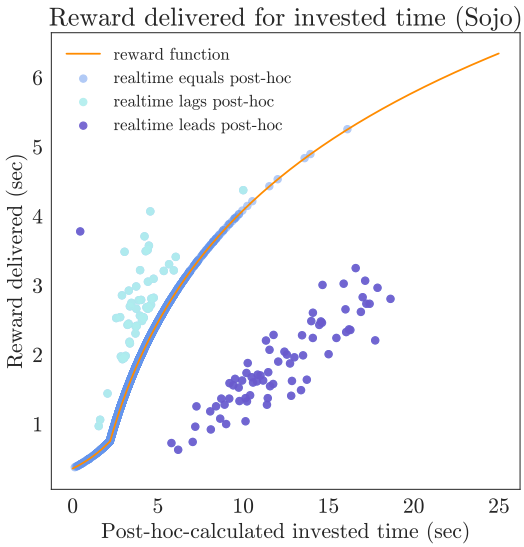

Figure 17: Reward delivered as a function of *post-hoc*-calculated invested time, as in Figure 15, with each point (trial) colored as in Figure 16. Programmed reward function in orange. Trials for which realtime-calculated invested time exceeds *post-hoc*-calculated invested time are light blue, and above the reward function line; trials for which *post-hoc*-calculated invested time exceeds realtime-invested time are purple, and below the reward function line.

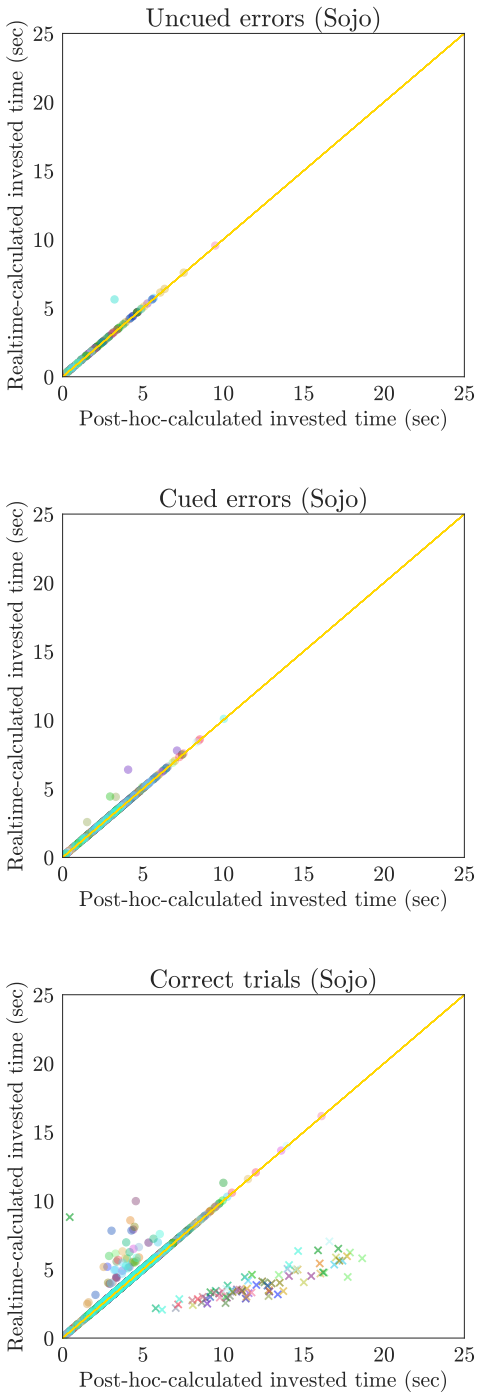

Figure 18: Discrepant realtime- vs. *post-hoc*-parsed invested times were distributed across epochs and trial outcomes. Each color represents the trials of a different epoch. For a subset of correct trials, the realtime-determined invested time was based on the *post-hoc*-parsed reward amount, and the trial is marked with an *x*; for all other trials, the realtime-determined invested time was determined based on the posthoc-parsed cue/back-on time, and the trial is marked with an *o*. Unity line in gold.

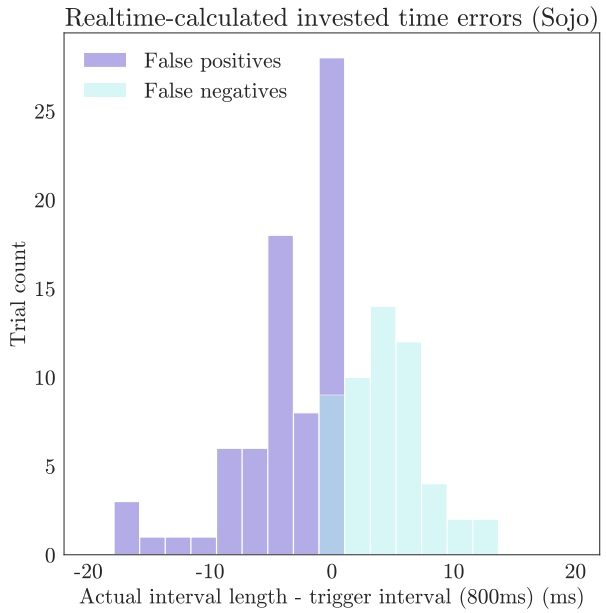

Figure 19: Distribution of differences from trigger interval for false positives (purple) and false negatives (light blue).

A.5 Rat: *Dolo*

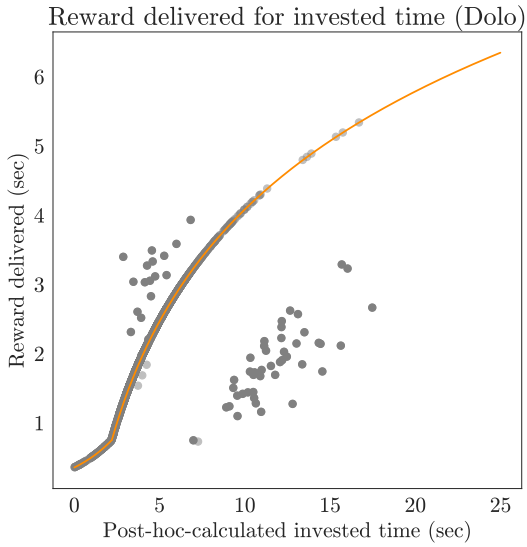

Figure 20: Reward delivered, from DIO milk pump log, as a function of *post-hoc*-calculated invested time from DIO beam break and unbreak log. Each gray point represents a trial. Orange line is the programmed reward amount to be delivered as a function of invested time. Gray points above the line correspond to trials for which the reward delivery amount exceeded what was coded, taking as input the *post-hoc*-calculated invested time. Gray points below the line correspond to trials for which the reward delivery amount was less than what was coded, taking as input the *post-hoc* calculated invested time.

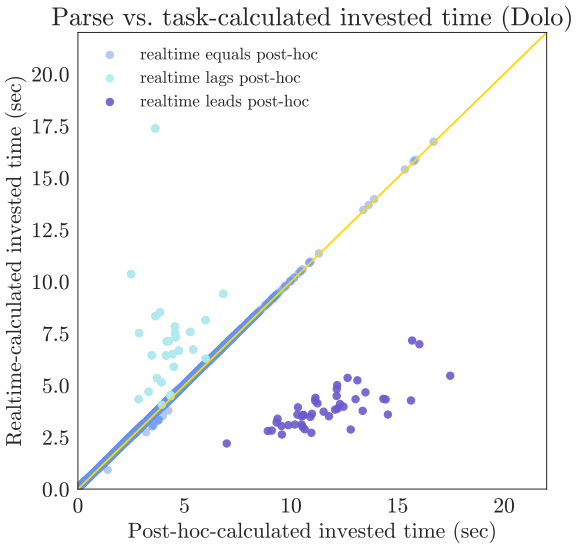

Figure 21: Realtime- versus *post-hoc*-calculated invested time. Each point represents a trial. Trials with equivalent realtime- and *post-hoc*-calculated invested time fall along the unity line (gold). Trials with realtime-calculated invested time greater than *post-hoc*-calculated invested time are light blue and above the line; trials with *post-hoc*-calculated invested time greater than realtime-calculated invested time are in purple and below the line. The six blue points just below the unity line represent a different, lower-magnitude, rarer, and less systematic parse error than that characterized here.

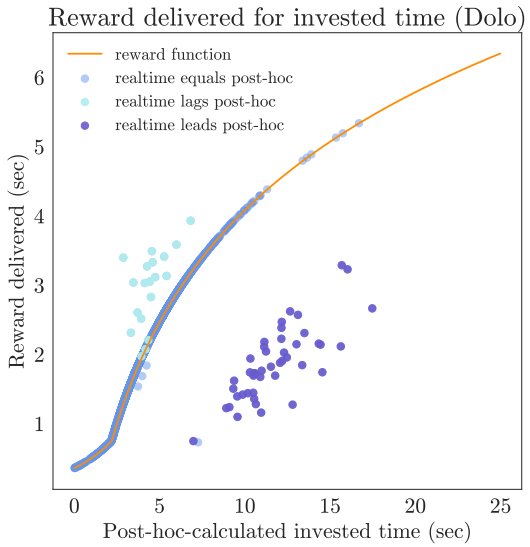

Figure 22: Reward delivered as a function of *post-hoc*-calculated invested time, as in Figure 20, with each point (trial) colored as in Figure 21. Programmed reward function in orange. Trials for which realtime-calculated invested time exceeds *post-hoc*-calculated invested time are light blue, and above the reward function line; trials for which *post-hoc*-calculated invested time exceeds realtime-invested time are purple, and below the reward function line.

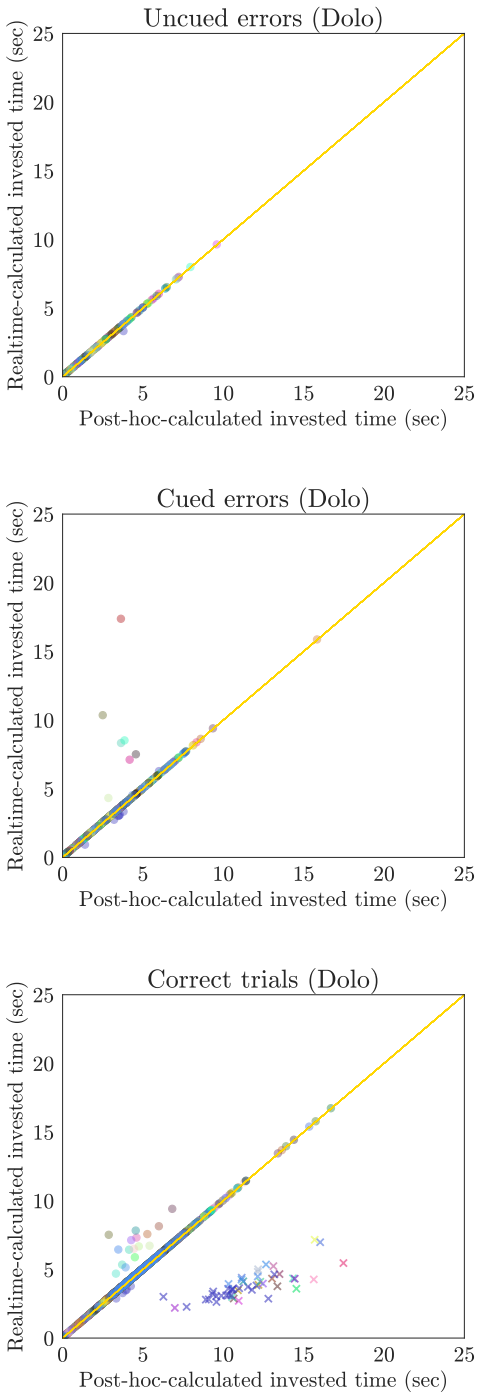

Figure 23: Discrepant realtime- vs. *post-hoc*-parsed invested times were distributed across epochs and trial outcomes. Each color represents the trials of a different epoch. For a subset of correct trials, the realtime-determined invested time was based on the *post-hoc*-parsed reward amount, and the trial is marked with an *x*; for all other trials, the realtime-determined invested time was determined based on the posthoc-parsed cue/back-on time, and the trial is marked with an *o*. Unity line in gold.

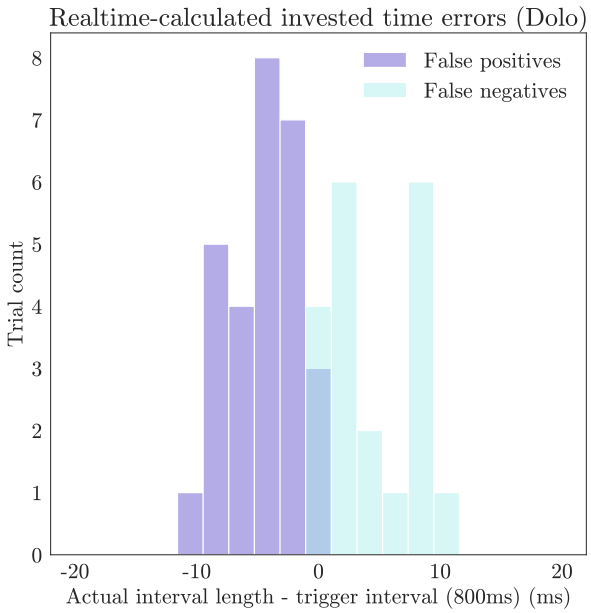

Figure 24: Distribution of differences from trigger interval for false positives (purple) and false negatives (light blue).

Appendix: Sharp wave-ripples in the episodic memory task

The design of the episodic memory confidence task (chapter 2) allowed us to investigate the potential function of SWRs in the retrieval of information on a short (minutes) timescale. High performance in the task requires rats to accurately recall and compare the timing of two recent episodes. By focusing on the period after which the task rule was well learned, we could dissociate the potential contribution of SWRs to rule learning (which should no longer be occurring) versus the learning of information necessary to successfully apply the rule. The task was designed with distinct trial phases, which enabled us to separately study the storage ('consolidation') and use ('retrieval') phases of memory. To study the potential function of SWRs in retrieval for selection of the choice port, we could examine their rate, duration, and amplitude during the current trial. To study their potential function in a storage capacity, we could examine those metrics during the two previous episodes being queried on the current trial.

We first asked which trial phase SWRs occurred during. As expected, we observed the highest rate of SWRs during periods of immobility and reward (Figure 25), corresponding to periods when the rat was at a port rather than traversing the track between them.

We next investigated whether SWRs contribute to retrieval on a short timescale. The task

can be divided into the decision phase, the confidence reporting phase, and the outcome phase. The decision phase begins when the cues are lit and ends when the rat selects a port. The confidence reporting phase begins when the rat selects a port and ends when he withdraws, and the outcome phase begins when the cue light comes back on to indicate a correct trial. The distinction between the decision and the confidence reporting phases enabled us to study the relative contribution of SWRs to these two processes.

Figure 25 shows the average relative ripple rate for correct and error trials as a function of task phase for one representative rat. For the entire trial ('homepoke - backunpoke', leftmost), the ripple rate is higher for correct than error trials. We compared the 'pre-outcome' phase that starts at homepoke and ends when the choice or back port light turns on, indicating a correct or error outcome respectively, to the 'post-outcome' phase that starts when the choice or back port light comes on and ends at back port unpoke (rightmost). We saw that there was a higher ripple rate for correct trials in both pre- and post-outcome phases, and that the rate was higher in the post-outcome than pre-outcome period for correct trials. Other differences have emerged from this line of investigation.

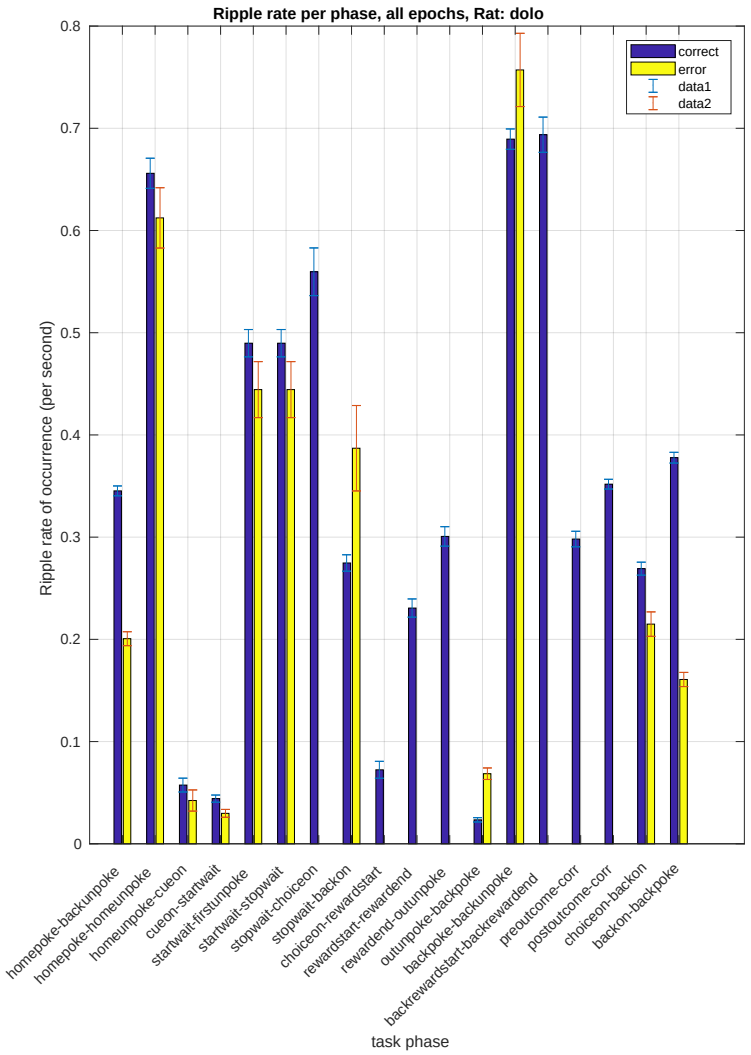

Figure 25: SWR rate during immobile periods (velocity less than 4 cm/sec) by behavioral phase, pooled across all epochs, no ripple size threshold, otherwise as described previously[38], for correct (blue) and error (yellow) trials, error bars = standard deviation, all trials. Behavioral phases defined by reward start and end times only exist for correct trials, and so report no SWR rate for error trials. The two rightmost phases correspond to the pre-outcome and post-outcome periods for error trials.

Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

DocuSigned by:
Hannah Joo
0FD3CB9FB1D34E4...
Author Signature

5/24/2020
Date