Title
Closed Loop Investigation of Hippocampal Replay

Author
Croughan, William D

Publication Date
2023

Peer reviewed|Thesis/dissertation

Closed Loop Investigation of Hippocampal Replay

By

William D Croughan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:
Professor David Foster, Chair
Professor Dan Feldman
Professor Joni Wallis
Professor Frederic Theunissen

Spring 2023

Abstract

Closed Loop Investigation of Hippocampal Replay

by

William David Croughan

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor David Foster, Chair

How people absorb and recall information about the world around them remains a central mystery of neuroscience. One popular framework for studying this process, systems consolidation, describes how episodic memories are initially encoded during an experience, consolidated into long term memory afterward, and recalled later. Each of these processes has been mapped to specific types of neural activity in specific brain regions. This mapping is incomplete. New progress each year reveals more details about how our memory system functions and new insights into the neural activity of our brains. One such type of neural activity is replay events. These events, typically investigated in rodent brains, allow an animal to simulate paths through an environment. A large body of research has shown clear connections between replay events and the consolidation phase of systems consolidation.

This dissertation expands on this research in two ways. In my behavioral results, I describe a series of experiments which suggest that replay may be critical for all three parts of systems consolidation, including the encoding, consolidation, and recall phases. In these experiments, we interfere with the rats' ability to generate replay events and observe a corresponding loss of memory for salient locations in an environment. These experiments employ closed-loop methods to interact with ongoing neural activity in real time, by delivering feedback whenever a replay event occurs. In my other chapter, I present a tool designed to greatly expand what experiments are possible in such closed-loop interactions. While most interactions involve either detection of all replay events or some coarse categorization, this tool allows experimenters to read out the simulated path through the environment directly as the replay event is occurring.

Chapter 1: Introduction

1a) Memory and the hippocampus

When patient Henry Molaisson underwent surgery to cure his epilepsy, he inadvertently became a touchstone representing one of the most striking documented connections between a specific brain region and cognition (Scoville and Milner, 1957). Henry Molaisson, known commonly as HM, had large portions of his temporal lobes removed to treat severe epilepsy. After recovering from the surgery, he displayed a curiously specific set of symptoms. He was able to interact largely normally with those around him, but had no ability to form new long term memories. Shortly after learning a person's name or a new fact, he would completely forget not only the name or fact, but the fact that he had ever heard it at all. Additionally, he exhibited what is known as temporally-graded retrograde amnesia - the extent to which he could recall his memories from before his surgery increased with how long before surgery they had occurred. Thus, memories from childhood were largely intact, while memories from shortly before his surgery were highly degraded.

While HM's case is particularly famous, many other patients with unique neurological injuries allow scientists to investigate how these injuries correlate with many cognitive functions. HM himself was a part of a small cohort of patients who underwent similar surgeries, and displayed similar, but not identical, cognitive effects. Patients with localized damage due to medical intervention, as well as specific injuries or diseases, each contribute unique insights. Years of this research has allowed scientists to paint a wider picture of cognitive functions' mappings onto brain structures. The hippocampus, in particular, has been implicated in working memory (Reed and Squire, 1998) and navigation (Teng and Squire, 1999). More recently, new technology has allowed scientists to temporarily interfere with hippocampal function in healthy adults to causally test these hypotheses, confirming the hippocampus' critical role in navigation and working memory (Goyal et al, 2018). The fact that the hippocampus - one centralized and anatomically separate region of the brain - is critical to such specific cognitive functions gives hope for rigorously probing how the brain handles these specific cognitive functions. At the same time, the hippocampus' location at the center of the brain, with strong connections with a large number of brain regions (Amaral and Witter, 1989), offers hope that progress on these questions may open doors to investigating many other cognitive functions and brain regions.

One model in particular has proven useful in describing how the hippocampus interacts with memory. Systems consolidation breaks our interactions with episodic memory down into three phases. The first phase, encoding, refers to the formation of a memory during experience. The next phase, consolidation, refers to the continuing process of translating the newly encoded experience into long term memory. Finally, recall refers to the process of accessing the stored memory and using it to guide

decision making. Viewed through the perspective of this framework, HM's symptoms indicate that his surgery interfered with his ability to consolidate memories, and perhaps interfered with his encoding ability as well. His inability to incorporate new information into his long term memory indicates either that the information could not be encoded in a way that was amenable to consolidation, or that the consolidation process itself could not take place. Additionally, the temporally graded retrograde amnesia (i.e. his inability to recall memories from just prior to surgery), indicates that those memories had initially been encoded but had not had time to consolidate into long term memory. Memories from long before surgery were already consolidated, and were thus unaffected.

This model, originally put forward in 1900 (Müller and Pilzecker 1900; Squire et al, 2015), has been the subject of extensive research. Today we can note neuroanatomical regions that contribute to the encoding and consolidation functions. Memories are initially encoded in the cortex and hippocampus. Through systems consolidation, the memories are transferred into the cortex, eventually being encoded in synaptic changes outside of the hippocampus. A causal parallel to HM demonstrates this cleanly in an animal model in Anagnostaras et al, 1999. The authors trained mice to fear two different environments before performing a hippocampal lesion. One association was trained 50 days before surgery and one was trained the day before. The first training, far from surgery, was unaffected by the lesion. Rats with lesions and sham surgeries both expressed a robust fear response. However, the learned association from the day before the surgery was highly affected, with control mice again showing a clear fear response, and lesioned mice showing little to no fear response.

1b) A cellular understanding of the hippocampus

With this anatomical connection established, a natural next question to ask is how the neural activity in the hippocampus enables the brain region's role in memory processes. Animal models allow us to probe this question more directly using sophisticated recording techniques and manipulations. In order to arrive at our answer, it is worth considering the properties of hippocampal cells without consideration of memory processes immediately, as a remarkable amount is known about their unique neural activity. We will then return to questions of memory to explore how the cellular and cognitive phenomena are connected.

In 2014, John O'Keefe received a joint Nobel prize for his discovery of place cells in the rodent hippocampus (O'Keefe and Dostrovsky, 1971). O'Keefe found that as rodents explore an environment, individual place cells will fire action potentials only in specific regions. Since this discovery, a wealth of information has been gathered about how these cells behave. Each place cell has a "place field" that governs its activity. When the animal is in this place field, the cell will fire action potentials at a consistent rate. When the animal leaves the place field, this firing stops. This firing reflects the belief by the animal of where it is in the environment, rather than directly responding to

specific stimuli. In experiments in which scientists manipulated environmental cues in various ways, place cells tended to track where an animal believed itself to be, as indicated by the animal's choices (Moser et al, 2008).

Two caveats in particular illustrate the wealth of information we have not yet discovered about place cells. The first is that this canonical activity is only present when a rat is actively engaged in an exploratory task. When a rat is at rest, these cells no longer exhibit the same pattern of activity. In fact, the bulk of this thesis is dedicated to understanding this resting neural activity. The second is that observations in primates do not show this activity with the same prevalence. In addition to cells that act analogously to rodent place cells (Ono et al, 1991), many experiments point to activity in the primate hippocampus that does not directly track physical location, but instead tracks more abstract notions of location. One paper showed that primate hippocampal cells tracked the animal's location in an "abstract value space", which the animals were motivated to keep track of internally as they performed a task (Knudsen and Wallis, 2021). Similarly, human studies show increased activity in the hippocampus when participants are asked to navigate an abstract space in their minds to complete a task (Tavares et al, 2016). These findings suggest that as the primate hippocampus developed, the hippocampal machinery was employed to solve more abstract problems. Thus, research into how rodents use their hippocampi to navigate may illuminate a wide array of cognitive processes in the human brain.

1c) Understanding the activity of many place cells

Considering place cells individually can yield insight into the hippocampus' functions. However, a much more holistic interpretation of their activity can be inferred when recording the activity of many place cells simultaneously. Recording from this many neurons is not trivial. In modern recording setups, hundreds of electrodes are packed together on a device that can be attached to the skull of a rat. These electrodes are wrapped tightly in groups of four - known as tetrodes - and then lowered slowly over the course of weeks to the correct depth in the brain. Continuous voltage readings from each electrode are streamed to a computer where action potentials from neurons are picked up as brief spikes in voltage. This voltage spike has a characteristic shape that is consistent when a single neuron fires multiple action potentials, and varies from neuron to neuron. Additionally, slight differences in position between each electrode in a tetrode mean that the action potential from neurons in different locations will be seen differently on each of the four electrodes. Using all of this information, an experimenter can parse the incoming voltage from a tetrode into firing times of many different neurons in a small radius around the tetrode. Note that it is also possible signals from multiple neurons will overlap and be indistinguishable. Therefore we refer to one isolated group of action potentials as a unit, since each unit might correspond to multiple neurons.

With this technology in hand, experimenters can probe the simultaneous activity of a population of place cells. They have found that each place cell has its own unique pattern of firing rates as a function of the rat's location in an environment. These patterns seem to be independent from one place cell to the next. That is, while aspects of the environment can affect where a cell has a place field, the location of one cell's place field does not seem to impact that of its neighboring place cells (Leutgeb et al, 2005). Therefore, recording from enough place cells simultaneously ensures that the entire recording environment will be covered by place fields. When this is the case, we have an expected profile of activity for the population for every area in the environment. By comparing recorded activity to these expected profiles, we can infer just based on the population's activity where a rat is in its environment.

The most common quantitative method for decoding spatial information from neural activity is known as Bayesian decoding. This method is used throughout the Results section unless otherwise noted. It is notable in its minimal assumptions, as described in Appendix A. In brief, it takes our knowledge about how each cell's activity is dependent on the rat's location, and allows us to infer a location based on each cell's activity. In the results, I explore the use of modifications to this method by introducing Bayesian priors, as well as so-called clusterless decoding (Deng et al, 2015). Note that these methods yield a probability distribution over possible locations the hippocampus can represent. A coordinated set of place cells that all are representing the same location will yield a punctate distribution with probability amassed around one point, whereas random firing by these neurons will typically yield a widely distributed probability distribution with many local maxima.

1d) Replay

With the decoding methodology in hand, we can revisit our earlier observations about the activity of place cells. As should be expected, decoding the spatial information represented when a rodent explores its environment, while the place cells are representing the animal's current location, yields one point of high probability that tracks the animal's location. When we look at cellular activity during rest, we do not see the same reflection of the rat's current location. However, we can see that the probability distributions again become punctate during brief bursts of neural activity, and trace coherent spatial trajectories throughout the environment. These trajectories are different in that they no longer represent the current location of the rat, and the represented location moves through the environment many times faster than a rat would move. This phenomenon was initially dubbed "replay," due to the recreation of patterns of activity seen during behavior. However, more recent results have painted a much more interesting picture of these events, showing that they can trace out paths that the rat has never taken (Gupta 2010), or indeed even trace a path through an environment before the rat visits that environment (Olafsdottir et al, 2015). Additionally, they can

represent paths moving forward or backwards in time (Foster and Wilson 2006), exhibit complex spatial dynamics (Pfeiffer and Foster 2015; Berners-Lee et al, 2022), and modulate their content according to the particular motivations of the rat (Pfeiffer and Foster, 2013; Shin et al, 2019; Olafsdottir et al, 2017).

Note that some papers investigating this phenomenon refer to sharp wave ripples as opposed to replays. A sharp wave ripple is the phenomenon observed in the LFP surrounding the hippocampal cell layer, and tends to co-occur with replay events. Therefore, we can often infer replay's role in cognition by indirectly observing its occurrence via sharp wave ripples.

1e) Replay and consolidation

Now let us return to our question from above: what neural activity might underlie the hippocampus' role in memory consolidation? Replays offer an obvious candidate, being related to previous experiences but occurring at separate times. Indeed, causal experiments in rodents have largely established that replay events are necessary for the consolidation of new information into long term memory, and correlational studies have shown replay's importance in generally processing new information. Replays have been shown to increase in frequency after new knowledge about rewarding areas of an environment has been discovered (Cheng and Frank, 2008; Ambrose et al, 2016). Importantly, this upregulation of replays of an environment persists in rests periods after the rat has been removed from that environment. Replays are also driven to occur during this period by new learning, such that when a rat has new information to consolidate, interfering with its ability to generate replay events will cause it to generate them at a higher rate to compensate (Girardeau et al, 2014). This effect is not seen when the rat has not learned new information, implicating this drive for replays in a specific role in memory consolidation. This is in line with many correlative papers in humans and rodents showing that more ripples correlate with better memory retention (Wagner et al, 2004; Durpet et al, 2010).

A large portion of experiments in this area employ tasks with a learn-wait-test paradigm, wherein knowledge about a task must be retained after a subject has been removed from that task through an intervening wait period. The wait period offers a period in which manipulations to the consolidation process can be easily tested. Two notable recent papers of this sort include work from the Kloosterman (Michon et al, 2019) and Csicsvari labs (Gridchyn et al, 2020). In the former, rats were trained on a radial arm choice task, in which the difficulty of the choice and size of the reward could be varied. Ripple interruption during a wait period showed differential effects according to task difficulty and reward size, wherein more difficult decisions with larger rewards were more heavily affected. In the latter, rats were trained on a simple reward finding task in two different environments. During a rest period, replays of one environment were interrupted, while replays of the other environment were not. They found that

behavior of the environment for which replays were interrupted indicated reduced memory for the exact reward location. These papers illustrate not only a clear role for replay events in memory consolidation, but the importance of the replay content itself in this role.

One notable implementation of the learn-wait-test paradigm is the Morris Water Maze (Morris et al, 1982). In this task, rats are placed in a tank of water, and have to swim until they find a platform on which to stand submerged just below the surface of the water. Since they cannot see the platform, they must first swim until they find the platform by chance. They are then removed from the platform, and, after some intervening wait period, are place back into the water in a different starting location. Rats do not like to swim, and so will generally swim straight toward the place where they earlier found the platform. Morris found that mice with lesioned hippocampi could not perform this task, thus indicating a deficit in memory of the goal location or the ability to navigate to that location.

Further developments on this paradigm have established that portions of hippocampal activity are specifically necessary for longer term memory retention. Steele and Morris (1999) found that rats with NMDA receptors blocked could perform the task as well as control rats if they were placed immediately back in the environment after removal. Interestingly, Silva et al (2015) saw a similar effect of NMDA blockade on replay events, wherein the blockade did not interfere with the production of replays of familiar environments, but did interfere with the ability of rats to integrate their current environment into future replays. Thus, the specific behavioral deficit in memory formation may be due to an inability to form replay events.

Together, these findings suggest that replay events are a critical part of memory retention, and act offline to reinforce knowledge about the location of rewarding places in an environment. This conception of replay is supported by many computational models in which offline reactivation of experiences can help improve performance. Johnson and Redish (2005) modeled the introduction of replay-like simulations into a model based on the basal ganglia, and saw improvements in the speed of learning a ruleset for a new task. Chersi and Pezzulo (2012) explore the hippocampal replay in combination with striatal activity as a method for the brain to implement model-based reinforcement learning. Schaul et al (2016) found that machine learning algorithms which leverage simulations analogous to replay events in their training perform better when these replay events are biased in their content, similar to what has been observed in real replay events (Pfeiffer and Foster, 2013; Gillespie et al, 2021). Mattar and Daw (2018) build a model in which replay events serve the purpose of spreading value information about a specific location across an internal cognitive map. Their model recreates some of the changes seen in the content replay events over the course of learning. See also Findlay et al, 2020 for a larger review in which the connections between replay and reinforcement learning is explored.

Additionally, some molecular investigations of memory support a role for replay events in memory consolidation. Morris established (Morris, 2006) that NMDA receptors are necessary for memory encoding and retention, but not retrieval, in agreement with the above work from Silva et al, and later experiments on the Morris Water Maze. Norimoto (Norimoto et al, 2018), discovered that downregulation of synaptic strength during sleep is dependent on NMDA receptors and replay events. These replay events during sleep have been shown to be important for memory consolidation (Wagner et al, 2004).

1f) Replay and other cognitive functions

This body of work draws a clear connection between replay and memory consolidation, but replay is likely involved in other cognitive functions as well. One such function that has been widely hypothesized to involve replay is that of memory retrieval. This is supported by some correlative findings showing that replays are biased to represent upcoming goals in some tasks (Pfeiffer and Foster 2013; Xu et al, 2018; Olafsdottir et al, 2017). Interruption of replays can actually reduce performance in an alternation task from the previous session (Fernandez-Ruiz et al, 2019, supplemental figures), an effect that cannot easily be explained by consolidation, since a consolidation effect would likely hinder future but not current performance.

However, other papers using different tasks show contradictory results. These experiments show variable or no predictive power of replay in predicting behavior (Gillespie et al, 2021; Shin et al, 2019). In some cases, replay content can actually be anti-correlated with behavior (Carey et al, 2019). While this finding could indicate replay is informing the immediate decision by confirming a lack of reward, it fits more parsimoniously, along with the tasks showing no predictive power, with the hypothesis that replays help to maintain the topology and information about an internal cognitive map. This conception of replay is supported by computational models that assign the hippocampus the role of maintaining a model of the environment (Chersi and Pezzulo, 2012; Momennejad et al, 2017). This picture is complicated, though, by the fact that replay events rapidly respond to topological changes in the environment, indicating that replays are not needed to establish this topology (Widloski and Foster, 2022).

Finally, it is possible that replay activity is simply the product of the hippocampus' natural inclination to produce sequences of activity when highly active, and that the spatial content of replay is not as relevant as the fact that cells are firing in bursts, triggering changes at the molecular level that enable memory retention and other cognitive functions. However, recent results suggest this is unlikely. Gridchyn et al (2020) showed a memory deficit specifically for an environment for which replay events were interrupted, compared to a separate environment for which replay events were not interrupted.

1g) Closed-loop experiments

As technology progresses, manipulations that interact with the brain's ongoing activity provide new ways to ask previously difficult questions. Many of the papers cited above relied on real-time analysis of LFP (Girardeau et al, 2009; Ego-Stengel and Wilson, 2010; Jadhav et al, 2012; Fernandez-Ruiz et al, 2019; Gridchyn et al, 2019; Gillespie et al, 2021). For additional examples of closed-loop feedback, see Nokia et al, 2012 and Knudsen and Wallis, 2020. Some work has harnessed basic content-specificity based on the activity of certain place cells (Lavilléon et al, 2015, Gridchyn et al, 2019) to separate activity related to certain areas. However, many open questions remain that could be probed using more advanced versions of these methods. For instance, questions about the possible different roles of forward and reverse replay in informing behavior (Foster et al 2006; Ambrose et al, 2016; Mattar and Daw 2018) are not accessible because forward and reverse replays are not accurately differentiable based simply on which cells are active. While Girardeau et al (2014) showed a generalized drive for increased ripples after learning, and Ambrose et al (2016) showed an increase in replay for recently experienced environments, the extent to which this drive can specifically motivate certain replays is an open question. For instance, it is unclear whether interrupting only ripples in one part of an environment lead to compensatory replays specifically in that area. Or if interrupting only forward replays would lead to a specific compensation of forward replays.

Progression in both hardware and software is enabling new methods. Vastly more electrodes can be packed into an implantable drive (see Foster et al 2006 vs Widloski and Foster, 2022, for instance) enabling a proportionally larger number of place cells to be accessed simultaneously. This explosion in data volume is accompanied by software that reduces the need to process this data by hand (Chung et al, 2017). Finally, as we learn more about how the hippocampus encodes information, we can leverage this knowledge to glean knowledge from activity that would have been uninterpretable previously (Gillespie et al, 2021; Widloski and Foster 2022; Cao et al, 2021).

1h) Outline of results

Below, I present two sections of novel results. The first section describes a software package that enables an experimenter to decode the spatial information of replay events in real time and specify arbitrary feedback contingencies based on that spatial information. I show that the software performs this task in a time frame that is amenable to real-time interaction with replays. I then describe how we quantified its accuracy on a linear track and in a two dimensional environment. Efforts to maximize this accuracy without sacrificing computational latency are described in detail.

The following section describes an exploration of tasks which we hoped might be amenable to more interesting investigation if closed-loop methodologies were applied.

After some exploration, we decided to record a large amount of behavioral data from a home-away task which has been shown to elicit interesting replay events that are not easily described by the consolidation model, and suggest a role for replay in memory recall or decision making. The task and our results are described in detail, along with various methodologies to rigorously extract maximal information from our dataset.

Chapter 2: CHOIR

2a) Algorithm Design

CHOIR (Crazyfast Hippocampal Online Investigation of Replay) is a software package consisting of an algorithm and a basic command line interface that enables experimenters to interact with replay events in real time. It analyzes incoming neural data to decode its spatial content, detects replay events, and analyzes the decoded content of those events. The algorithm is designed for experiments divided into three phases. In the first “training” phase, neural data is recorded along with position data of a freely moving rat. In the second “processing” phase, individual units are isolated from the recorded neural data, and their place fields are constructed according to the recorded position data. Finally, in the third “online” phase, the results of the processing phase are used to cluster incoming neural data and decode its spatial information. The algorithm is optimized to enable fast decoding of place cell activity, so that feedback can be delivered during the replay event being analyzed. Arbitrary contingencies can be specified that will determine what feedback is delivered in response to various types of replay and behavior.

A key design consideration of this algorithm is to minimize the time between when the brain expresses a replay with interesting spatial content, and when feedback is delivered in reaction to that content. We quantify this more rigorously by noting that in any incoming spike train, we can label one action potential as the “critical” action potential if it is the first spike for which the decoded spatial information including that spike passes some decision threshold and indicates that a replay of interest is in progress. There are two major contributing components to this latency (Fig 1). The first is the delay between the start of a replay event, and the occurrence of the critical action potential. This delay is affected by the choice of algorithm that determines when a replay event is in progress. A permissive algorithm that is quick to declare a potential replay event may enable a short latency, but will incur a high rate of false positives. A stringent algorithm will require more data to be confident in a replay event, and therefore will cause this latency to be higher. A user that intends to disrupt ongoing replay activity should consider this tradeoff, as they are unlikely to achieve both high specificity and low detection latency.

The second aspect contributing to overall detection latency is computational latency. This can be divided into the time it takes to transmit the data, assign a spike to a cluster, integrate the spatial information from that spike, and analyze all of the recorded spikes’ spatial information to determine the appropriate feedback to deliver (Fig 1). CHOIR minimizes the latency between receiving the critical action potential and delivering feedback by keeping a continuously updated record of the spatial information being represented, and updating this record immediately as each action potential contributes additional spatial information. Thus when the critical action potential arrives,

all previous spikes have already been processed and will not contribute to latency. Storing and updating a mathematically correct decoded probability distribution directly - that is, one that is non-negative and sums to 1 - would require multiplication of an incoming spike's place field by the current probability distribution, and subsequent normalization. This includes at a minimum one multiplication, one addition, and one division operation for each place bin. Instead, a proxy for this distribution is kept and updated with each incoming spike. For each place cell spike, the logarithm of its place field is added to the ongoing spatial information. This means each spike requires only one addition per place bin per spike to update the spatial information. When a spike burst is evaluated, these values can be exponentiated and normalized to recover the mathematically correct probability distribution. Note also though, that there is a monotonic relationship between these stored values and the corrected probability. Thus, if only the maximum-probability place bin is required for evaluation, this can be done directly on the stored log-probabilities, and no further processing is required.

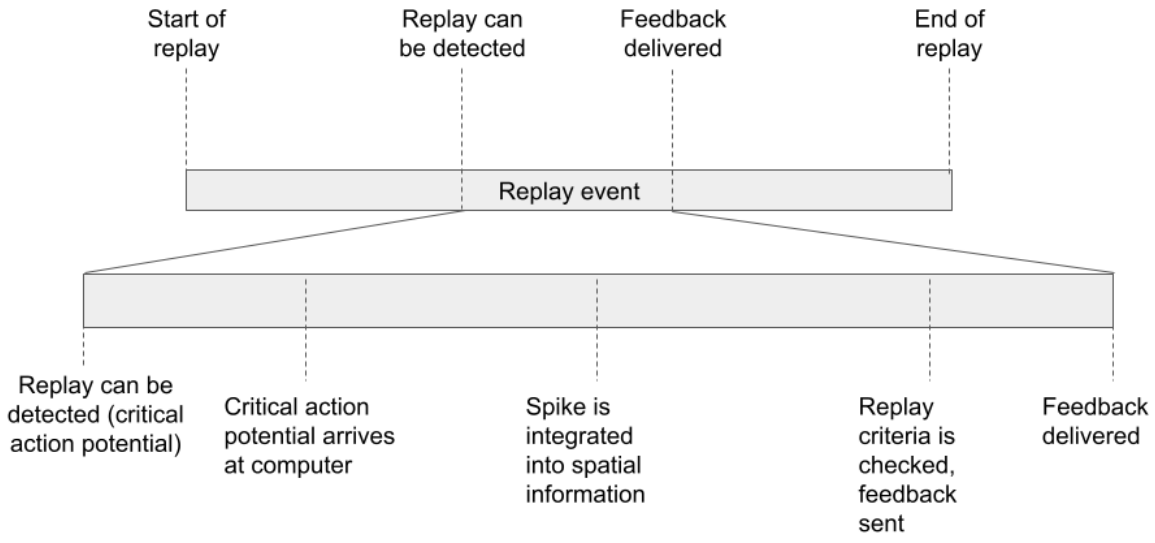

Fig 1. Components of feedback latency *For a single replay event, there is some delay between the start of the event and when enough neural activity has occured to classify it. The bottom row zooms in on the computation at this point, and shows the order of computations that must occur after this point before feedback is delivered.*

2b) Latency measurements

CHOIR divides its computations among a number of threads (Fig 2). The communication thread is in charge of receiving data from the recording software and passing it to the rest of the threads in CHOIR. It also monitors the overall number of spikes and updates a running average and standard deviation. The “worker” threads receive incoming spikes, cluster them, and incorporate their spatial information into the maintained buffer. Another thread monitors incoming video data and tracks whether the rat is moving or in a zone of interest. The main thread monitors the number of spikes in each decoding window. When the spike count exceeds a threshold z-score, it analyzes the spatial information and determines whether to deliver feedback. Since the video thread updates only rarely compared to the other threads and performs negligible computations, we can consider the total computational latency to be the time it takes the worker thread to integrate an incoming spike into the spatial information plus the time it takes the main thread to detect a spike burst and analyze the spatial content. Since these threads are not synchronized with each other, I measured these two durations separately, and we can infer a total latency by summing the two values. Figure 3 shows a histogram of the measured latency of these two threads over the course of a twenty minute session. On average, we expect a computational latency of less than 1ms. In the worst case in which both threads run at their absolute slowest, we would see a computational latency of 3ms. Note any detection algorithm must look back tens of milliseconds to make an accurate determination of replay content, and therefore we can conclude that the computational latency is small enough for the experiments for which this algorithm was designed.

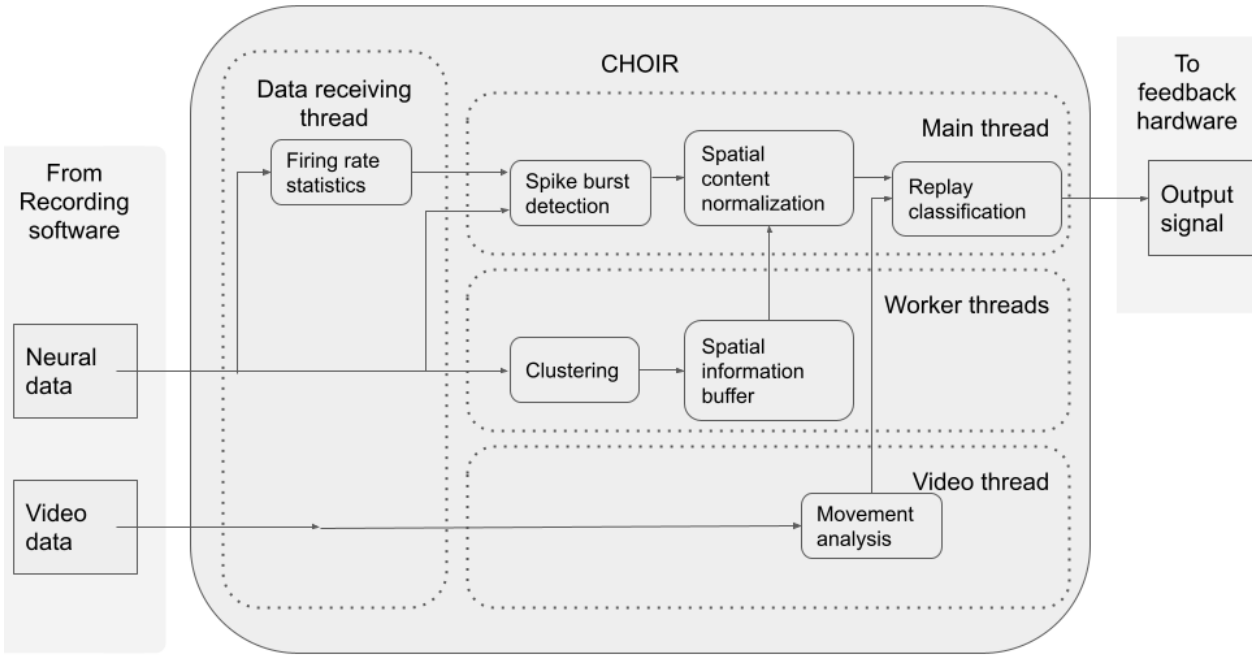

Fig 2. CHOIR Schematic. *Schematic of the information flow through CHOIR, and the separation of tasks into threads.*

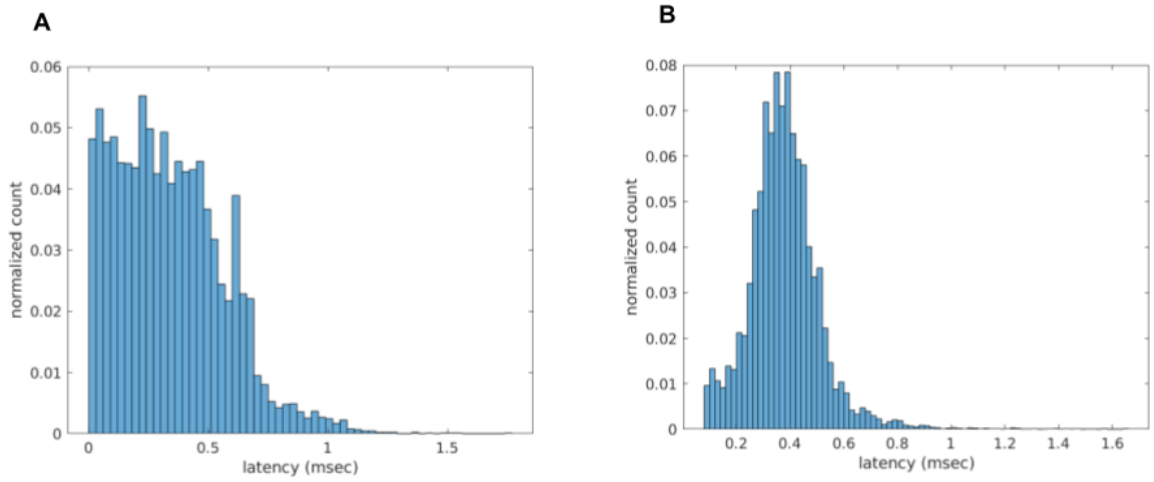

Fig 3. Measured CHOIR latencies. *Measured CHOIR latencies for A) clustering and integrating spatial information from each spike and B) normalizing the spatial information probabilities and classifying the replay.*

2c) Accuracy measurements

The assumed experimental design of this algorithm presents significant challenges to accurately decoding spatial information. Tiny movements of the tetrodes relative to the brain can change the shape of a single units' observed waveform. This means that a decoding model built on previously recorded data will lose decoding accuracy over time. While much work has gone into mitigating this problem across large swaths of neuroscience, it still presents a significant challenge, and most analyses of tetrode recordings never assume that two units from separate recordings can be assigned to the same neurons. Even if the same neurons could be tracked perfectly across recording sessions, this still would not guarantee decoding accuracy, as place cells are known to "remap" between exposures to the same environment. This can be seen as a simple "rate remapping" in which their place fields remain largely unchanged, but their overall firing rate changes. In some cases place fields change locations entirely. In particular, this is prone to happen when the shape or connectivity of the environment changes, or if other aspects of the environment such as reward contingencies change.

We sought to measure the degree to which this instability would impact the accuracy of any categorization of replay events. We used one session separated into two halves to simulate a multi-session experiment. This enabled us to have a confident ground truth of single units tracked over one session, while measuring the impact of tetrode drift over time. The first half was used to cluster place cells and build place fields, and the other half was used to analyze the accuracy of replay decoding. We label the results of this test as online decoding, as it modeled the constraints inherent in our real online decoding algorithm. Ground truth decoding was considered to be the decoded spatial content when clustering and building place fields from the entire session, which we label as offline decoding. Online decoding generally preserved much of the spatial content, though in a somewhat degraded state (Fig 4a). We quantified the extent to which this would impact an online decoding algorithm by simulating a task wherein replays that entered a certain goal zone should be detected (Fig 4b,c). This binary decision gave us a false positive, false negative, true positive, and true negative rate of detection (Fig 4d).

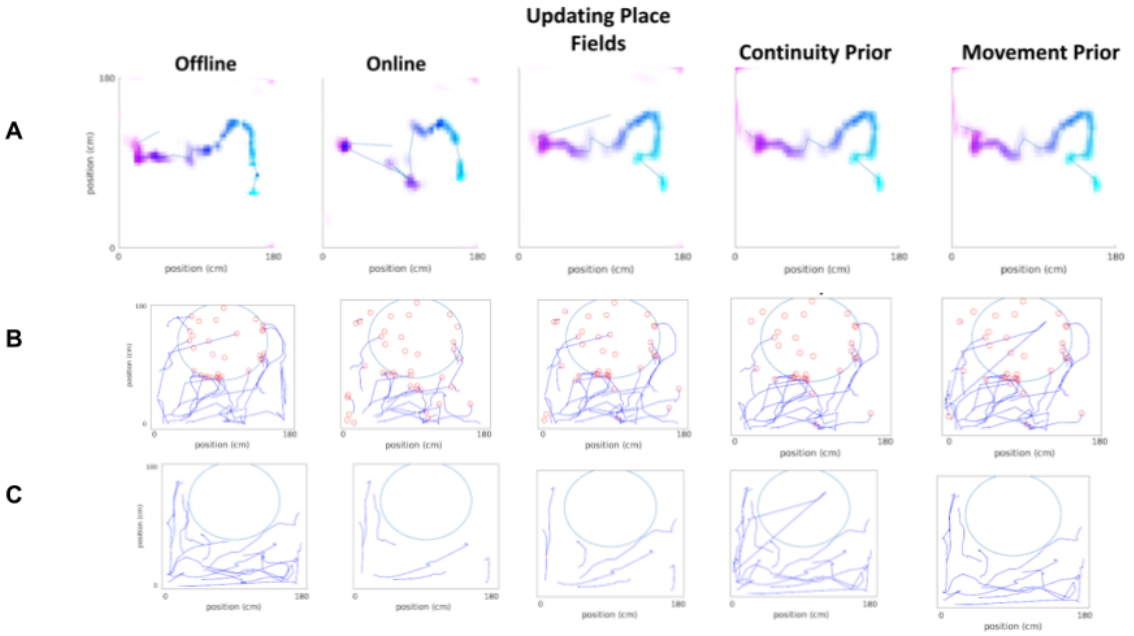

D

	True Positives	True Negatives	False Positives	False Negatives
Online Decoder	131	60	19	19
Updating Place Fields	129	61	18	18
Continuity Prior	126	72	7	7
Movement Prior	129	71	8	8

Fig 4. Replay accuracy classification in 2d. *A) One example replay event decoded using each method. B) Replays identified by each method has having entered the goal zone, marked with a blue circle. The blue lines show the path of the replay before detected entry into the goal zone, with the path defined by the offline-decoded content. Red circles show the location of the replay at the time of detection. C) Same as B but for replays that were classified as not entering the goal zone. D) Classification results for each decoding method.*

The level of classification agreement was well above chance. It also gave us a baseline against which we could test modifications to our decoding method. We first tried continually updating the place fields as incoming spikes were clustered in the online decoding phase. With this method, we could compensate for remapping by incorporating more information about each neuron's firing rate throughout the run. We

then tried modifying our decoding formula to use a bayesian prior that incorporates the knowledge that replays of interest trace out continuous paths through space. We first identified the decoding frame with the most spiking activity, as a proxy for choosing the frame for which we are most confident in the decoding result. Then we iteratively decoded adjacent frames in time, but replaced the usual uniform bayesian prior with a normal distribution centered at the center of mass of the adjacent decoded frame. In this way we signified that we expect the replay content of adjacent frames to be near each other.

Finally, we sought to capture an intuition that replays not only are continuous, but tend to have some momentum in one direction. Thus we used a similar technique but incorporated a prior that assumes a consistent speed and direction of motion. We again started with the frame with the most spiking activity and a uniform prior. We then decoded adjacent frames with a prior that reflected a normal distribution over the distance from the adjacent frame's center of mass. Then for all other frames, we iteratively decoded using a prior that had the same radial gaussian multiplied by a factor incorporating the angle from the adjacent frames' direction of motion. The exact formulas are expanded in Appendix A. Overall, we saw that updating the place fields had a minimal effect on accuracy, whereas incorporating either prior reduced both the false positive and false negative rates to a similar degree. Some concern should be taken to ensure that the process of jumping to the frame with the highest spike rate does not invalidate these results in the context of online decoding, in which decoding must happen during the spiking activity in question. Note though that with the computational efficiency noted above, if a wave of spiking activity occurs which establishes a new frame with highest spiking rate, decoding could begin anew for all previous decoding frames using the iterative gaussian method without incurring a large cost on latency.

Another unique application of this software we predict is the distinction between forward and reverse replays. We thus sought to find the best algorithm with which we could classify these two types of replays. Again, we hoped to determine this in the context of online decoding, in which clustering must be done automatically based on previously established clusters, and decoding must be done using previously established place fields. Furthermore, a classification must be made during the replay event itself, and thus could not be based on, for instance, the start and end points of a replay. To determine the best decision criteria, existing data was analyzed and forward and reverse replays were extracted using normal best practices for offline analysis. Then, various combinations of clustering, event detection, and classification algorithms were assessed based on their agreement level with the offline analysis. A list of all methods tested can be found in appendix B. Among all combinations, using BIRCH and agglomerative clustering, thresholding based on both ripple power and decoded

content, and classifying based on a simple sum of posteriors was found to be the most similar to offline classification.

Chapter 3. Ripples and Behavior

We sought to explore the range of hypotheses that could be tested using closed loop stimulation. We aimed to find a task which seemed likely to involve replay events in decision making, and thus may lead to an observable behavior difference when the subjects were subjected to replay manipulations. We embarked on an exploration of tasks. At the end of this exploration, we observed trends of behavioral differences when ripples were interrupted in the home-away task. This task also displays interesting correlations between replay content and navigation, making it potentially interesting fodder for future content-specific interruption experiments. Data is presented from six rats: Martin, B13, B14, B16, B17, and B18. Data for the initial task explorations was recorded from Martin. All statistical tests were performed using a custom built plotting and statistics library, described at the end of this chapter. All statistical shuffles presented used 5000 shuffles, unless otherwise noted.

3a) Ripple interruption

Recording replay requires a large number of simultaneously recorded neurons. To more quickly explore which tasks may be dependent on replay, we chose to use sharp wave ripples as a proxy for replay activity. Ripples are the LFP signature that coincides with replay events, and only require one high quality signal near the cell layer of the hippocampus to detect, in contrast to the dozens of tetrodes that must be precisely in the cell layer to decode replay. Since the two events co-occur, a ripple-interruption experiment that shows that ripples are necessary for a certain task is suggestive that replay events are necessary for that task. Furthermore, behavioral differences that occur when interference of hippocampal activity is timed to sharp wave ripples implies that the content of replay events informs the observed behavior.

We designed a system which enabled online detection of sharp wave ripples, and delivery of a brief electrical pulse to interrupt ripples. We used Spikegadgets hardware embedded in custom-built micro-drives, and Spikegadgets software to record neural signals from the vicinity of dorsal CA1 and video footage from overhead cameras. The Spikegadgets recording software, Trodes, includes an API that enables users to write their own code that receives a stream of the incoming neural and video data. We used this API to stream LFP data, which was then filtered to extract the power between 150-250Hz. To filter out electrical artifacts, which typically appeared as a burst of power across many frequencies and all tetrodes, we subtracted the filtered power from a baseline tetrode which was left far away from the hippocampal cell layer. The ripple power was then z-scored, and any reading above a set threshold was considered to be a ripple. Upon ripple detection, our software sent a serial message to an arduino, which then sent a brief TTL pulse to a stimulus isolator. This stimulus isolator generated a 200μs pulse between 40-60μA, which was sent to a bipolar stimulating electrode

implanted in the rat's ventral hippocampal commissure. To minimize electrical interference, all cables were shielded with aluminum foil, which was connected to the ground of the acquisition system.

Given that our goal was to draw a connection between ripples and behavior, recording from more animals more easily became more important than recording from many cells in one animal. We redesigned the recording drive to be lighter and easier to build, with only eight tetrodes (Fig 5), as opposed to the usual 40 or 64. This "ripple drive" also includes one concentric bipolar stimulating electrode targeted to the ventral hippocampal commissure. All of the tetrodes and the stimulating electrode could be independently moved vertically. After implanting, the tetrodes were adjusted until ripples could be found, and at least one clear hippocampal cell was present. We then could use this cell to calibrate the stimulating electrode. We adjusted the strength of the stimulation and height of the stimulating electrode until a stimulating pulse had a clear silencing effect on the recorded hippocampal unit. We measured this by recording for 20-30 minutes in the sleep box while stimulation was timed to interrupt ripples, followed by a break with no stimulation for at least 30 minutes, followed by 20-30 minutes of stimulation that was also timed to the detection of ripples, but delayed by 200ms. Functional stimulation showed an immediate reduction in spiking in the interruption condition relative to the delayed condition, and a reduction 200ms later in the delayed condition (Fig 5). Note that no calibration data was collected for Martin in the delayed condition, but he displays a clear immediate reduction in spiking upon stimulation. This, in combination with his behavioral results on the W-maze being in line with published ripple-interruption results (described below), offers clear evidence that interruption was working as intended in this subject.

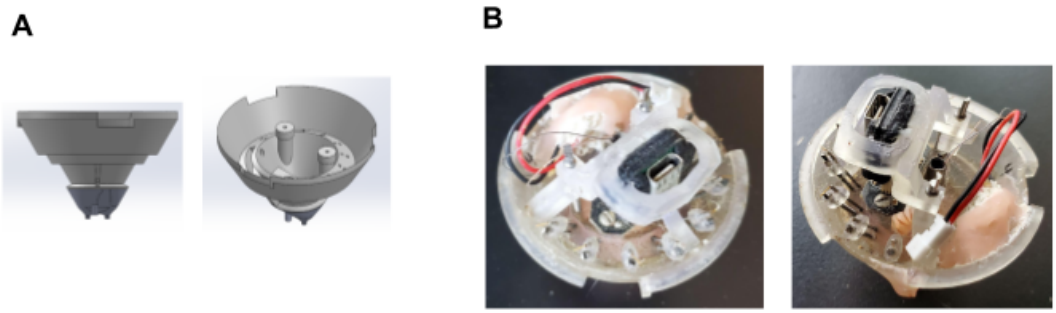

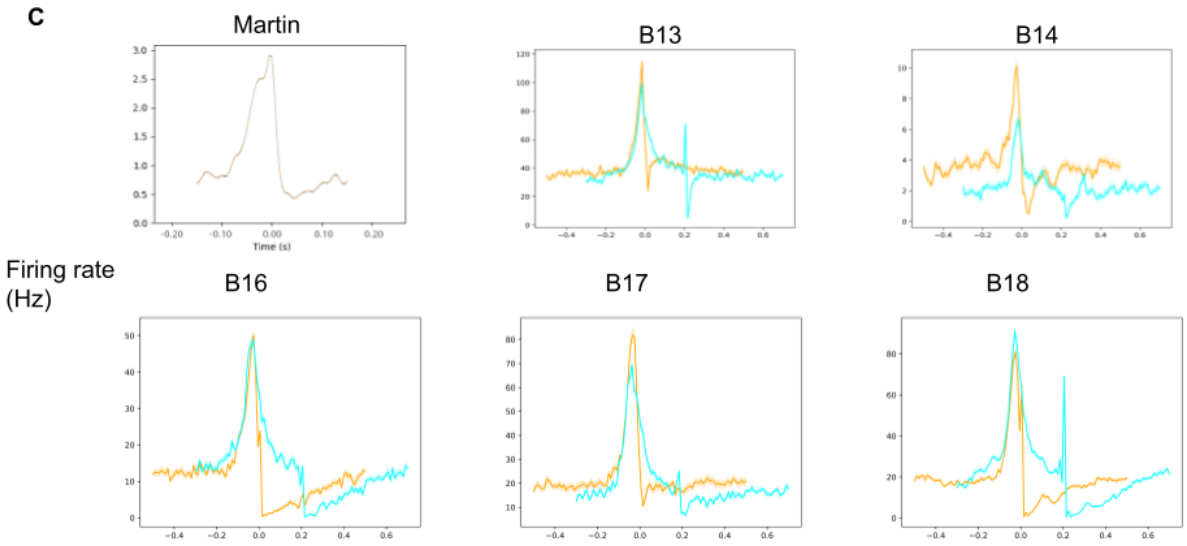

Fig 5. Ripple Interruption. *A) Rendering of ripple drive design. B) Image of ripple drive including tetrodes, stimulating electrode, and the supporting electronics. C) PSTHs showing effective interruption of place cell activity in response to delivered stimulation in each rat.*

3b) Initial task exploration

Our initial inclination was to take tasks which have been shown to be hippocampus-dependent and test whether they depended on replay activity. We thus chose two tasks to replicate. The first task is the Morris Water Maze. In this task, rats or mice are placed in a pool and must find a hidden platform on which to stand before they are removed. Multiple trials are then repeated from different starting locations, and since rats dislike the water environment, they will typically go directly to the hidden platform from the previous trial in order to escape. Rats with hippocampal lesions show an inability to navigate to the platform from the previous trial.

We adapted this protocol to be amenable to electrical recordings by motivating the rats to find reward in a dry environment instead of escaping a wet environment. This environment is the same described below for the home-away task, and consists of a 6ft by 6ft arena with 36 reward wells spread evenly throughout in a square grid. In each block, one well was filled, and the rat was placed in a corner and would then search for this random well. Three more trials were performed in which the same reward well was filled, analogous to the escape platform being made available in the same place in the water maze each trial (Fig 6a). The starting corner for each trial was chosen pseudo-randomly, so that a different corner was used each trial. Similar to the results

from the Morris Water Maze, we observed that our test subject consistently performed better on trials 2-4 than on trial 1, indicating a memory of the reward location from trial 1 (Fig 6b). Contrary to our hypothesis, interrupting ripples did not impair the rat's ability to complete this task. Trial 2 latencies were not significantly different between interruption sessions and control sessions (Fig 6c).

Since our dataset is limited in size, we cannot reliably test for subtle behavioral effects with small effect sizes. However, one interesting trend was observed during a "pseudo-probe" period, which we defined to be the exploratory period of the first trial of a block before the reward had been found. If the previous set of trials was a control block, then the rat tended to spend more time at that block's reward location than if it had been an interruption block. Probe trials are often interpreted as an expression of the strength of a memory or level of confidence. This trend, then, would imply that ripple interruption reduces the confidence of the animal in the location of the reward, and therefore leads to less perseveration around that location at the start of the following block (Fig 6d). Any further interpretation should warrant gathering more data to confirm or deny these preliminary findings. We did not gather any more data on this task; however, these results align interestingly with our results from the home-away task, described in the section below.

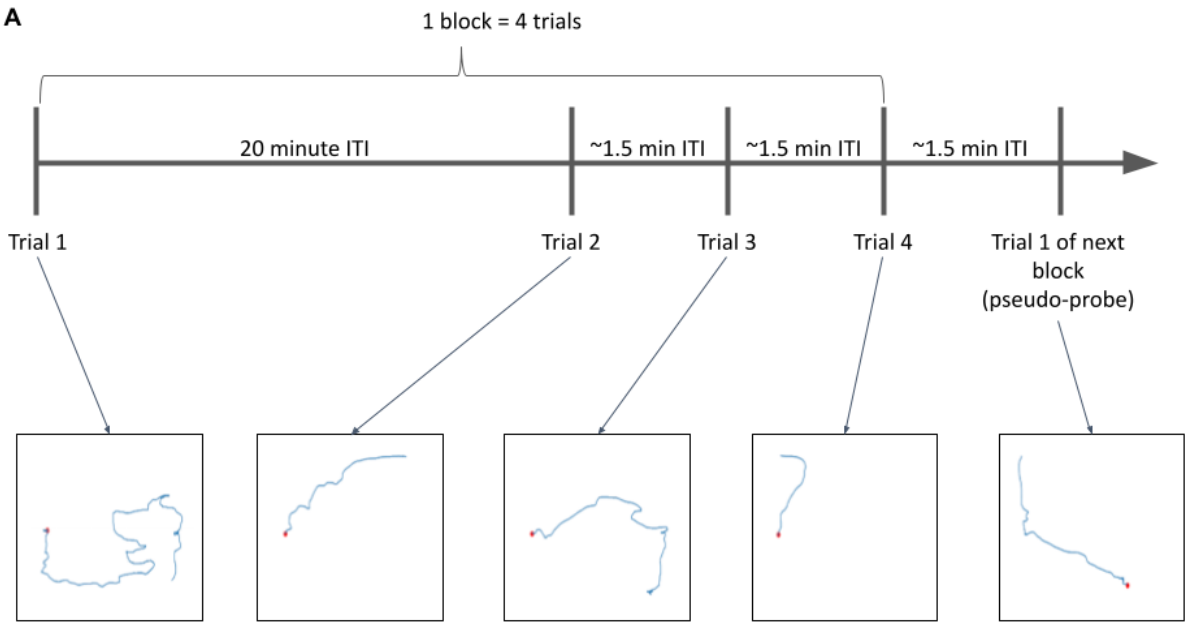
A
1 block = 4 trials
20 minute ITI
~1.5 min ITI
~1.5 min ITI
~1.5 min ITI
Trial 1
Trial 2
Trial 3
Trial 4
Trial 1 of next block (pseudo-probe)

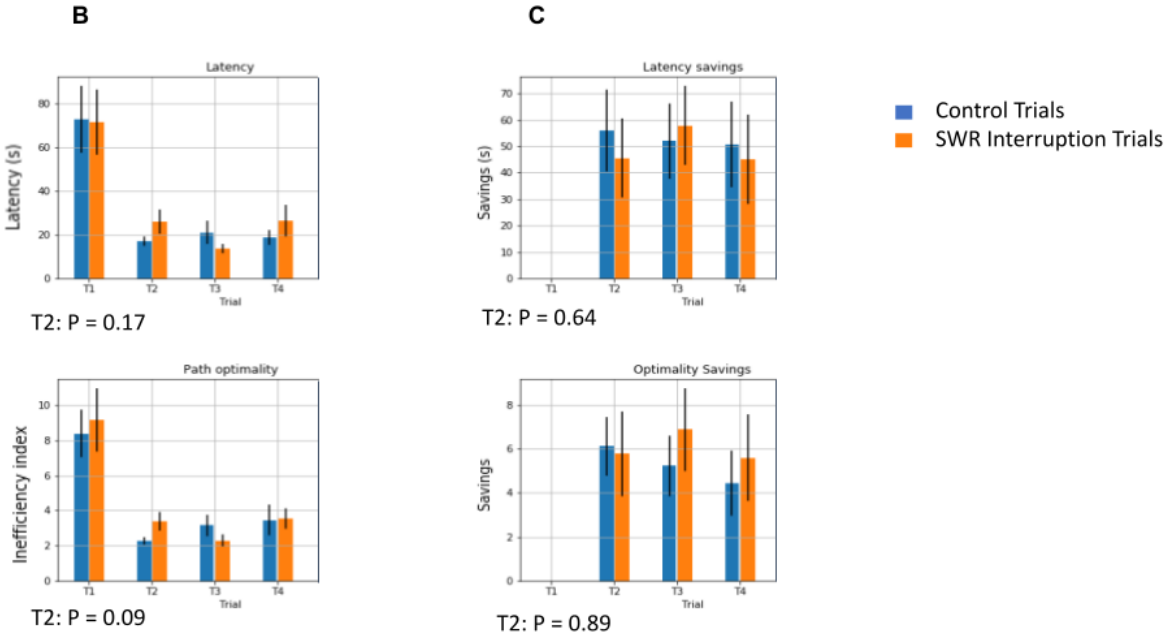
B
C
Latency
Latency (s)
Trial
T2: P = 0.17
Latency savings
Savings (s)
T2: P = 0.64
Control Trials
SWR Interruption Trials
Path optimality
Inefficiency index
T2: P = 0.09
Optimality Savings
Savings
T2: P = 0.89

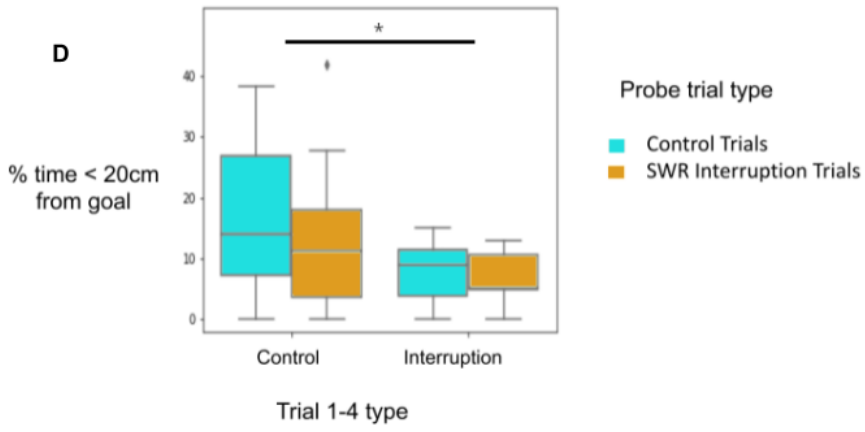

Fig 6. Dry Morris Water Maze. *A) Schematic of the dry morris water maze with example behavior from one block of four trials, and the first block of the following trial, which we also refer to as the pseudo-probe. B) Latency and normalized path length separated by trial and condition in the dry morris water maze. Error bars show the standard error of the mean. P-values shown for t-test of trial 2 data, n=44. C) Latency and normalized path length savings, defined as the difference from the trial 1 value from that block of four trials. P-values shown for t-test of trial 2 data, n=44. D) Percentage of time near the previous block's rewarded well, separated by the condition of the previous and current block. Horizontal lines represent quartiles with one outlier marked as a diamond. 2-way ANOVA, p=0.039, n=29.*

The second preliminary task we explored was the W-maze. In this task, rats are placed on a maze with one central arm and two side arms. Each of the three arms has a reward port at its end. The rat must learn an alternation rule to receive reward (Fig 7a). Previous studies (Jadhav et al, 2012, Fernandez-Ruis et al, 2019) have shown that interrupting ripples impairs performance at this task. Specifically, performance on the outbound choice, in which rats have to remember which side they visited previously and go to the opposite side, was reduced when ripples were interrupted throughout the session. Fernandez-Ruis et al (2019) measured this by performing two sessions each day, one of which was a control session, and subtracting each day's control session performance from its ripple-interruption session performance. We followed the same protocol. Our implementation showed an effect of interrupting ripples that visually

matched up with these previous results, showing a trend that is in line with their observed effects (Fig 7b). At this point, we were focused on qualitative results, and did not attempt to statistically quantify the extent to which our results aligned with the reported data. We ultimately decided not to record more animals on this task for a number of reasons, noted in the discussion section. Nonetheless, this is a useful confirmation that our interruption protocol appears to reproduce reported effects of ripple interruption on behavior.

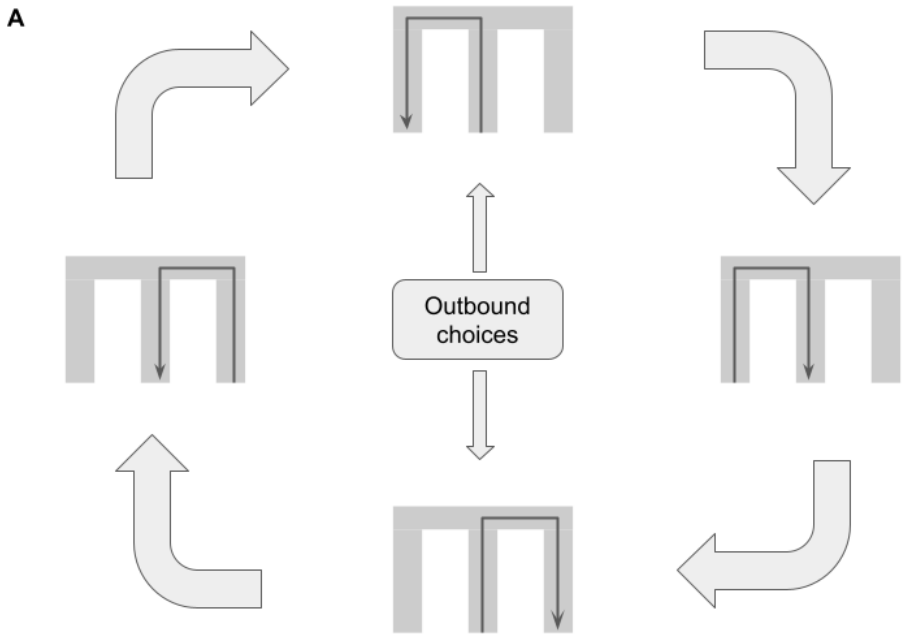

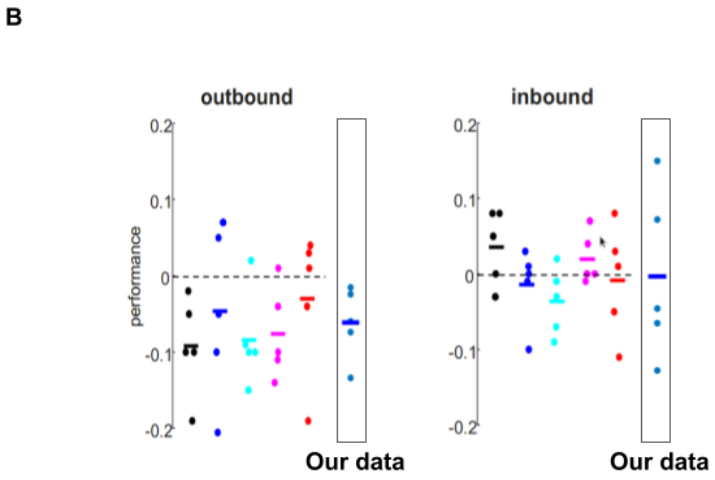

Fig 7. W-maze. *A) Schematic of correct performance on the W-maze. Outbound choices involve starting in the center arm and choosing either the right or left arm. B) Figure modified from Fernandez-Ruis, et al, 2019, to compare our observed data with their presented data on the same task.*

3c) Home-away task

We chose to focus on recording data from a home-away task, in which rats have been shown to exhibit interesting replay content (Pfeiffer and Foster, 2013). This task takes place in a square environment measuring six feet across. At the center of each square foot, a reward well is embedded into the environment floor. These reward wells can be filled and drained remotely and silently by an experimenter. Before the task begins, the experimenter chooses one well to be the home well for that session, and a series of other wells to be away wells. When the rat is first placed in the environment, they must search for and find the home well to receive reward. Then the experimenter fills the first away well, and the rat must find that well. Whenever the rat finds reward at an away well, they then must return to the home well to receive reward again, at which another randomly chosen away well is filled. The task proceeds in this way, alternating between foraging for a random away well and returning to the consistent home well, until either twenty minutes have passed or a preset number of wells have been found. Away wells are never repeated, so at the end of this task the rat has received reward multiple times at a single home well and once at each of the away wells. Each session, a new home well and series of away wells is chosen. The task structure is illustrated in Fig 8a.

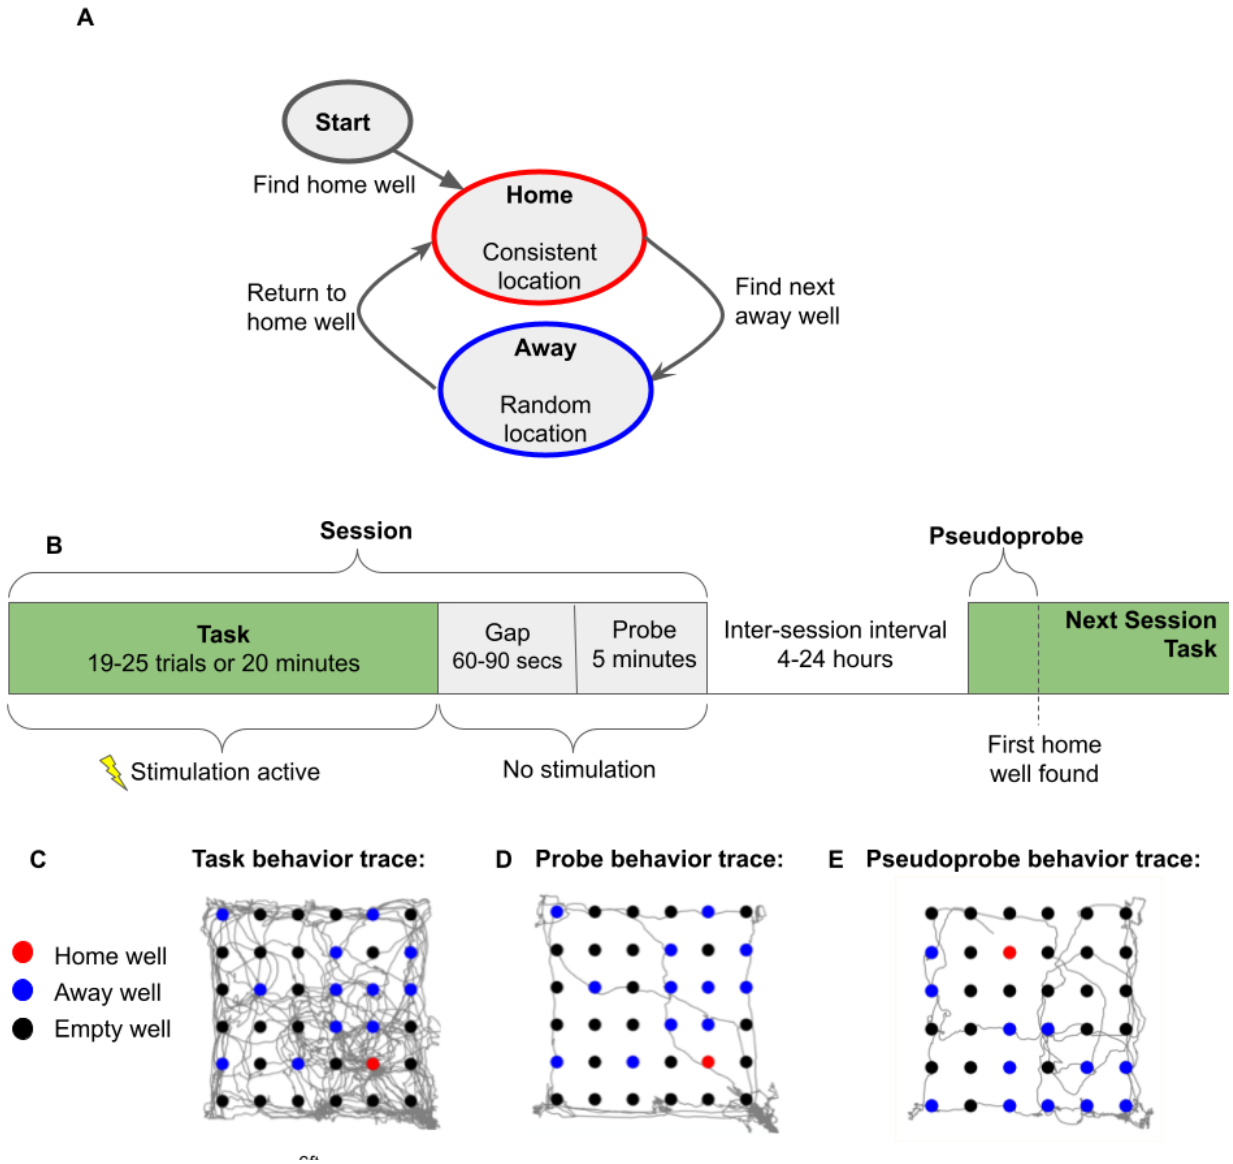

Fig 8. Home-away task. *A) Schematic of the home-away task. B) Schematic of task timeline. C) The full path of the rat during one example session. D) The full path of the rat on the immediately following probe. E) The full path of the rat on the pseudoprobe, which is the task phase of the following session until the first home well is found.*

After recording data from one rat on this task, we made two modifications for other rats. The first is that we restricted the home well to be one of the 16 wells that were not along the wall on the environment. This is because we wanted a clear behavioral indication that the rats learned the location of the home well. Typically, they would spend a lot of time circling the exterior of the maze. If the home well was along the exterior, it would be difficult to establish when they chose to approach the home well

and when they had simply stumbled on it. The other modification was to fill two away wells at a time during away trials. Whenever one away well was found, the other was drained, and then refilled in addition to a new away well on the following away trial. In this way, the random foraging part of the task was made easier, while still preserving the overall structure of the task. For our first rat, this was used as the last step of training before moving to the task with only one away well filled at once. However, the fact that this made the task easier seemed to lead to more consistent behavior, as rats would typically stop searching if they could not find an away well on one particular trial after a few minutes. Therefore we kept this task structure for recording for the other five rats. If the rat did not find all of the away wells, then there could be a confounding factor here where rats were more likely to find away wells in places that they looked more often. However, in the majority of sessions the rats found all of the away wells, and therefore the most important aspects of the task structure, that each away well was rewarded exactly once and that the location of away wells was randomly distributed, were preserved.

The condition of each session was assigned pseudorandomly to avoid any confounding incidental correlations between the condition of consecutive sessions or the time of day of each session. Two sessions were run each day, with condition assignments determined four days at a time. The first of the four days consisted of one control and one interruption session. The order was chosen at random. The following day was the same but the order was the opposite of the first day. The third day was either two control or two interruption sessions, determined at random. And the fourth day was the inverse of the third day.

Before recording, rats are well trained on the task. Therefore the learned information each session is restricted to the locations of the reward wells. Rats found all of the wells on most sessions (Fig 9a). Their understanding of the task structure, and learning of the home location each session, can be confirmed by looking at how quickly they completed each home trial compared to away trials. Each rat consistently decreased its latency to the home well over trials, indicating a memory being formed for that location (Fig 9b). The same pattern is not observed in away trials, as is expected for random foraging. Note that for all rats except Martin, the latencies to away wells is typically lower on the first few trials. This is the expected result of the modifications noted above, including restricting the home well to be away from the border and filling two away wells each away trial. The more important observation for our interpretation of these results is the pattern of latencies over time, which shows clear learning of the home well without a similar pattern for away wells.

Our primary question we wanted to ask was whether ripples were necessary for good performance at this task. The replay phenomena noted earlier in this task led us to hypothesize that replays were used to recall the location of the home, help the rat navigate to that location, or both. We looked for evidence of this by comparing the

home-trial latencies in control and interruption sessions. No consistent difference was observed between conditions, tested by shuffling the condition label of each session and comparing the mean values for each trial (Fig 9c). We also quantified performance on each trial with the normalized path length, which is the total distance traveled by a rat during a trial divided by the straight-line distance between its starting and ending location. Again we saw a clear difference between home and away trials, but no difference between the home trials of the different conditions (Fig 9d,e).

Earlier this year, a group found that rats can complete this task without any hippocampal activity (Duszkiewicz et al, 2023). This is in line with our results. They did, however, see a small effect of their lesion on the home trials in early learning. Specifically, in the first three returns to the home well, they found that rats without hippocampal activity had longer path lengths on their home trials. We did not see this effect in our data when measured either by path length, normalized path length, or latency (Fig 9f-h). Results are displayed for individual rats, and an across rat shuffle, as described in the statistics section 3f, is displayed to the right. Briefly, each rat's data was shuffled by randomly reassigning condition labels to each data point. One across-rat statistic was then generated by taking the mean of the summary statistic for each rat, where the summary statistic is the difference between the means of the two conditions, weighted by the number of datapoints collected for that rat. The true value is shown as a red line above the shuffled distribution with 5000 shuffles.

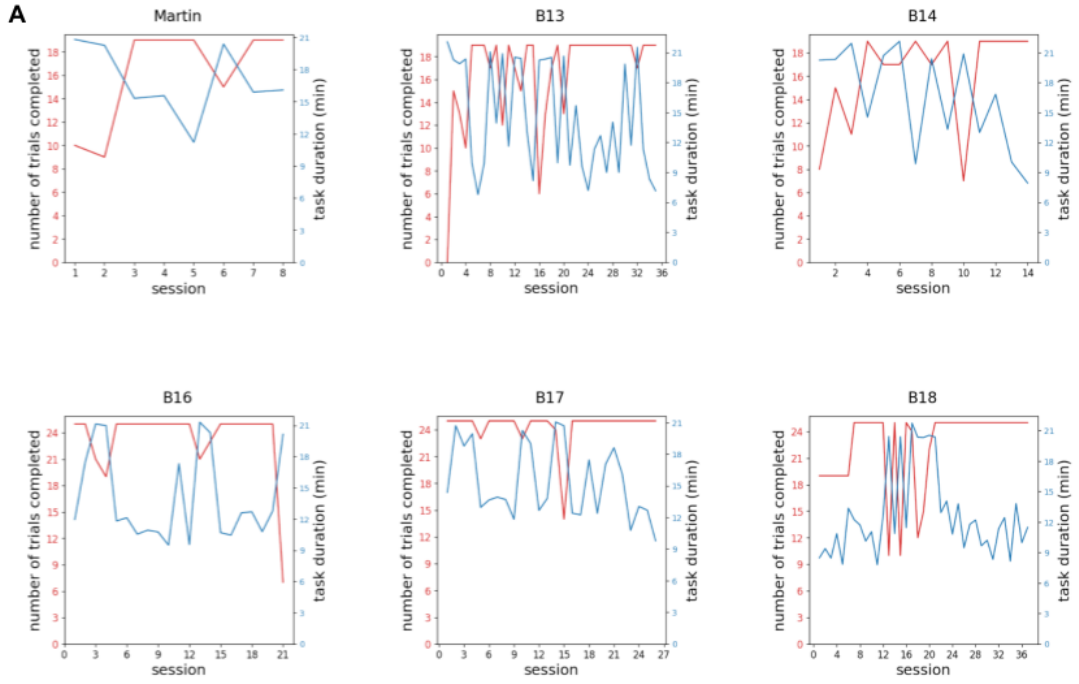

B

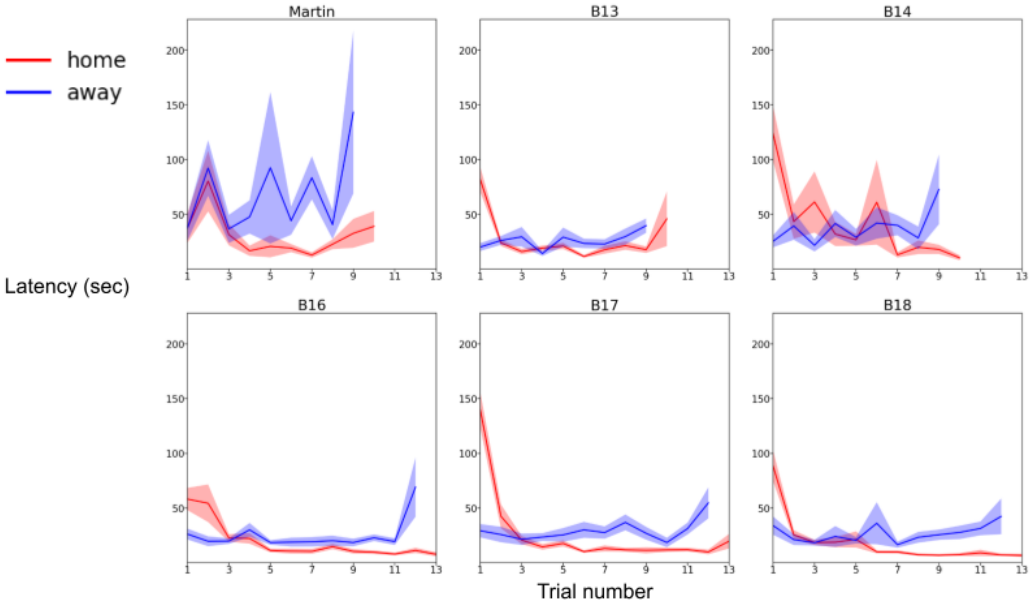
home
away
Martin
B13
B14
B16
B17
B18
Latency (sec)
Trial number

C

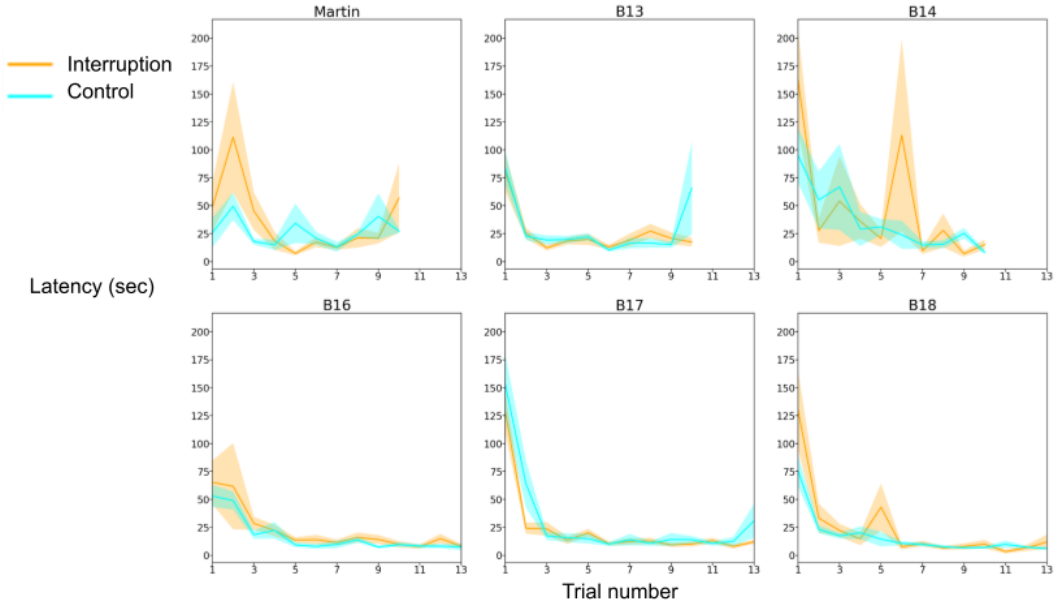
Interruption
Control
Martin
B13
B14
B16
B17
B18
Latency (sec)
Trial number

D

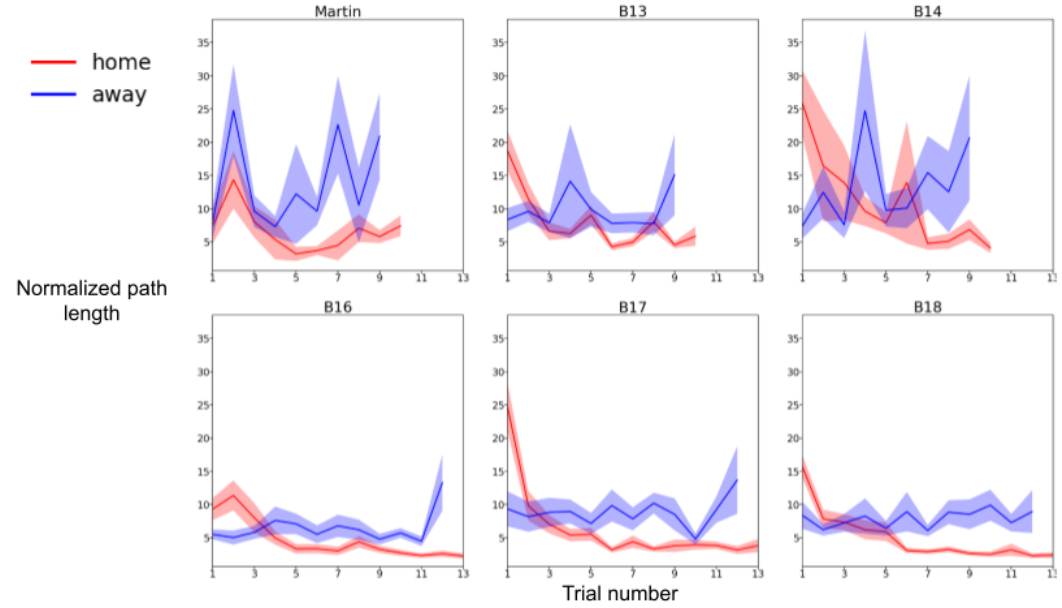
home
away
Martin
B13
B14
B16
B17
B18
Normalized path length
Trial number

E

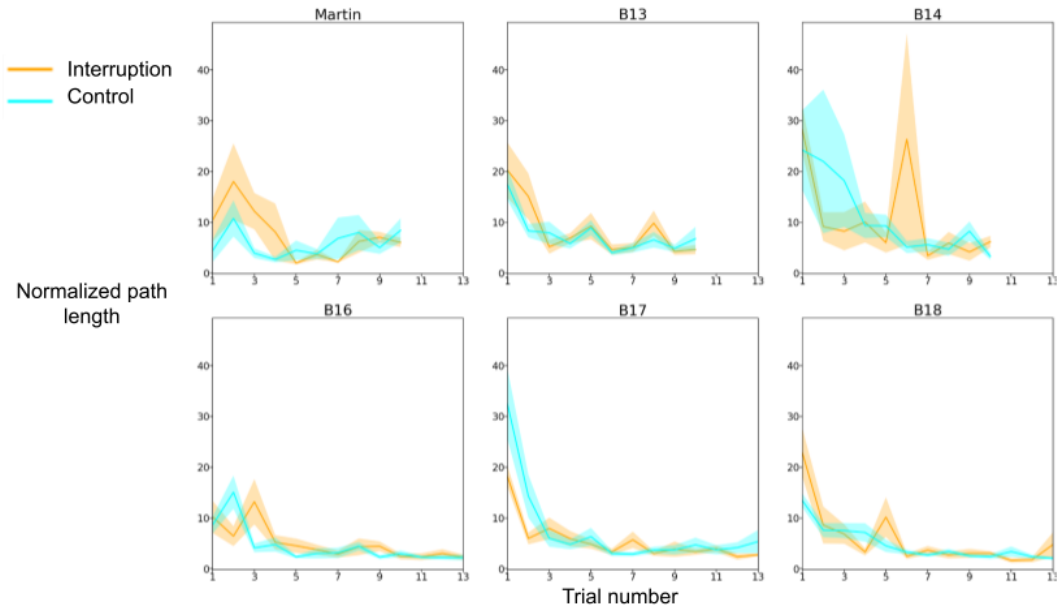
Interruption
Control
Martin
B13
B14
B16
B17
B18
Normalized path length
Trial number

F

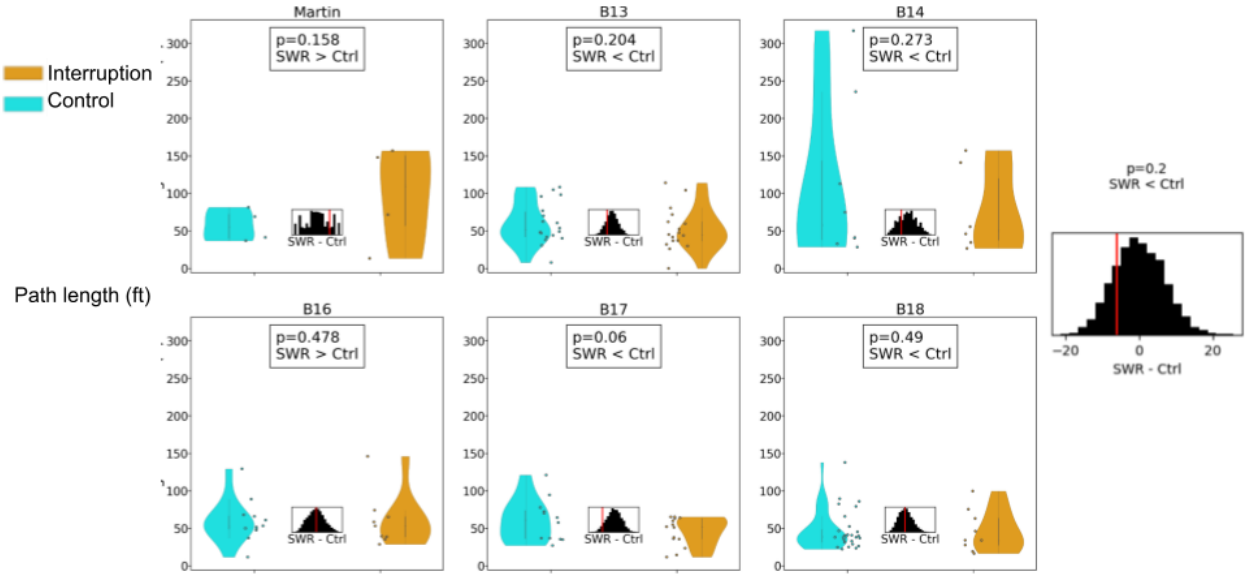
Interruption
Control
Path length (ft)
Martin
p=0.158
SWR > Ctrl
B13
p=0.204
SWR < Ctrl
B14
p=0.273
SWR < Ctrl
B16
p=0.478
SWR > Ctrl
B17
p=0.06
SWR < Ctrl
B18
p=0.49
SWR < Ctrl
SWR - Ctrl
p=0.2
SWR < Ctrl
SWR - Ctrl

G

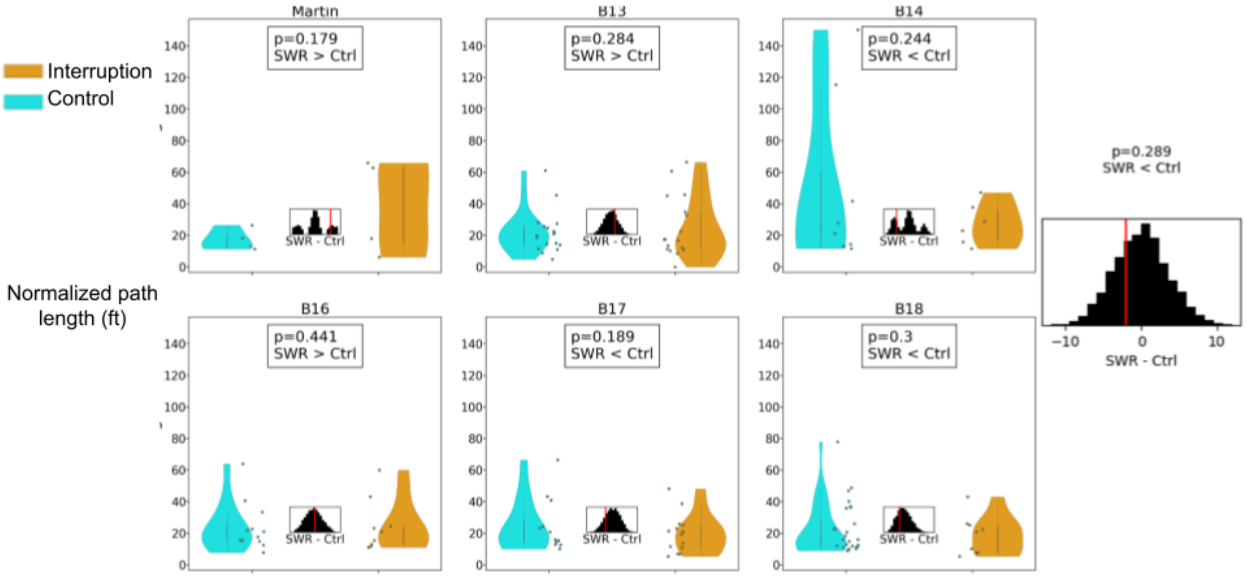
Interruption
Control
Normalized path length (ft)
Martin
p=0.179
SWR > Ctrl
B13
p=0.284
SWR > Ctrl
B14
p=0.244
SWR < Ctrl
B16
p=0.441
SWR > Ctrl
B17
p=0.189
SWR < Ctrl
B18
p=0.3
SWR < Ctrl
SWR - Ctrl
p=0.289
SWR < Ctrl
SWR - Ctrl

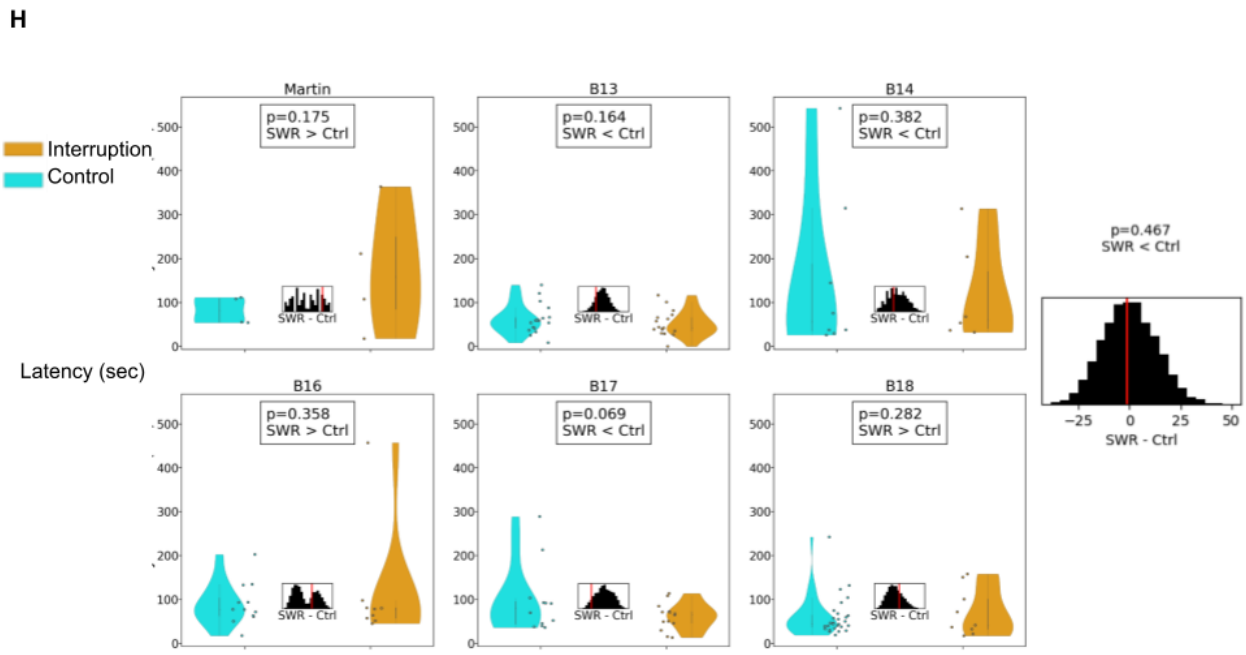

Fig 9. Home-away task behavior. *A) For each rat, the number of wells found (red) and total duration (blue) in each session, plotted chronologically. B) Latency of each home trial (red) and away trial (blue). C) Latency of each home trial in interruption sessions (orange) and control sessions (cyan). D) Normalized path length of each home trial (red) and away trial (blue). E) Normalized path length of each home trial in interruption sessions (orange) and control sessions (cyan). F) Average path length of home trials 2, 3, and 4. Left: data for each rat. Each dot represents one session, while violin plots show overall distribution of the dots. Inset histograms show distribution of the difference of means between interruption and control sessions, generated by shuffling the category of each session. Actual difference of means in data is shown as a red vertical line. RIght: An "across" rat shuffle was performed as described in the statistic library section. A red line showing the value for the data is shown with the histogram of shuffled values. G) Same as F, but for normalized path length. H) Same as F, but for latency.*

<u>3d) Probe trial</u>

To more deeply probe the effect we may be having on the rats' memory, we included a probe after completion of the task. In this probe, the rats are placed back into the same environment in which they performed the task, and are allowed to explore with either no reward given, or sparse reward delivered after some exploration time as described below (Fig 8b). Probes are employed in the water maze experiment, and many other spatial memory tasks (i.e. Gridchyn et al, 2020), for a number of reasons. In the task itself, it is difficult to gauge the confidence of the animal's knowledge of the home

location. Once the rat approaches and finds reward at the home well, we cannot gather more data from that trial. In the exploration phase of a probe trial we gain greater insight into the level of confidence because they may either spend longer near the home well expecting reward, or revisit that location more often, either of which indicates a degree of confidence in their memory. One other important difference between behavior during the probe and during the task is that stimulation is delivered only during the task portion of a session. Thus, behavior differences during the probe due to ripple interruption must reflect a lasting effect of the earlier stimulation. Finally, one widely studied navigational strategy by which rats could complete this task without use of their hippocampus is path integration. This strategy involves the rat tracking its total movement away from the home well, and then simply reversing that total movement to return to the home well. Removing them from the environment interferes with this process, and should bias the animals to use other methods of navigation that may be more hippocampus-dependent.

We first quantified the extent to which the rats' behavior reflected their memory of the home well over other locations of the environment. The rats displayed a clear preference for revisiting the home well, measured both as the number of visits and the total time spent near that well (Fig 10a,b). These quantities were compared to the same quantities measured for other away wells which were not on the border to establish a baseline. Five of the six rats displayed a highly significant preference for the home well over away wells, and this difference remained highly significant when tested across all rats.

Note that Martin, B13, and B14 experienced a probe trial with no reward. However, we noted that the rats tended to not explore on many probe trials (Fig 10c), presumably having learned the task structure and that probe trials were unrewarded. Therefore when we recorded data from B16, B17, and B18, we instituted a pseudo random reward schedule. Once during each probe, the home well from that session's task phase would be filled according to the following pseudo random schedule. Among a group of four days, the four control sessions would be assigned a fill time of one, two, three, and four minutes in a random order. The same was done for the interruption sessions. When that time was reached, the home well would be filled.

We hypothesized that interrupting ripples would interfere with the formation of the memory of the home well location, and would lead to reduced preference for the home well during the following probe. We considered the same measures of total time and number of entries to the home well, this time separated by the interruption condition in the preceding task phase. We found that the data trended in the hypothesized direction, but we did not observe a significant difference at a threshold of 0.05 (Fig 10d,e).

Note that most of Martin's control sessions, and B18's later control sessions were performed with no stimulation, instead of delayed stimulation. This was due to technical difficulties that arose during recording. As we did not see a difference between these no-stimulation sessions and delayed-stimulation sessions, we included them all as

control data points. The same results with these sessions excluded are shown in Appendix C.

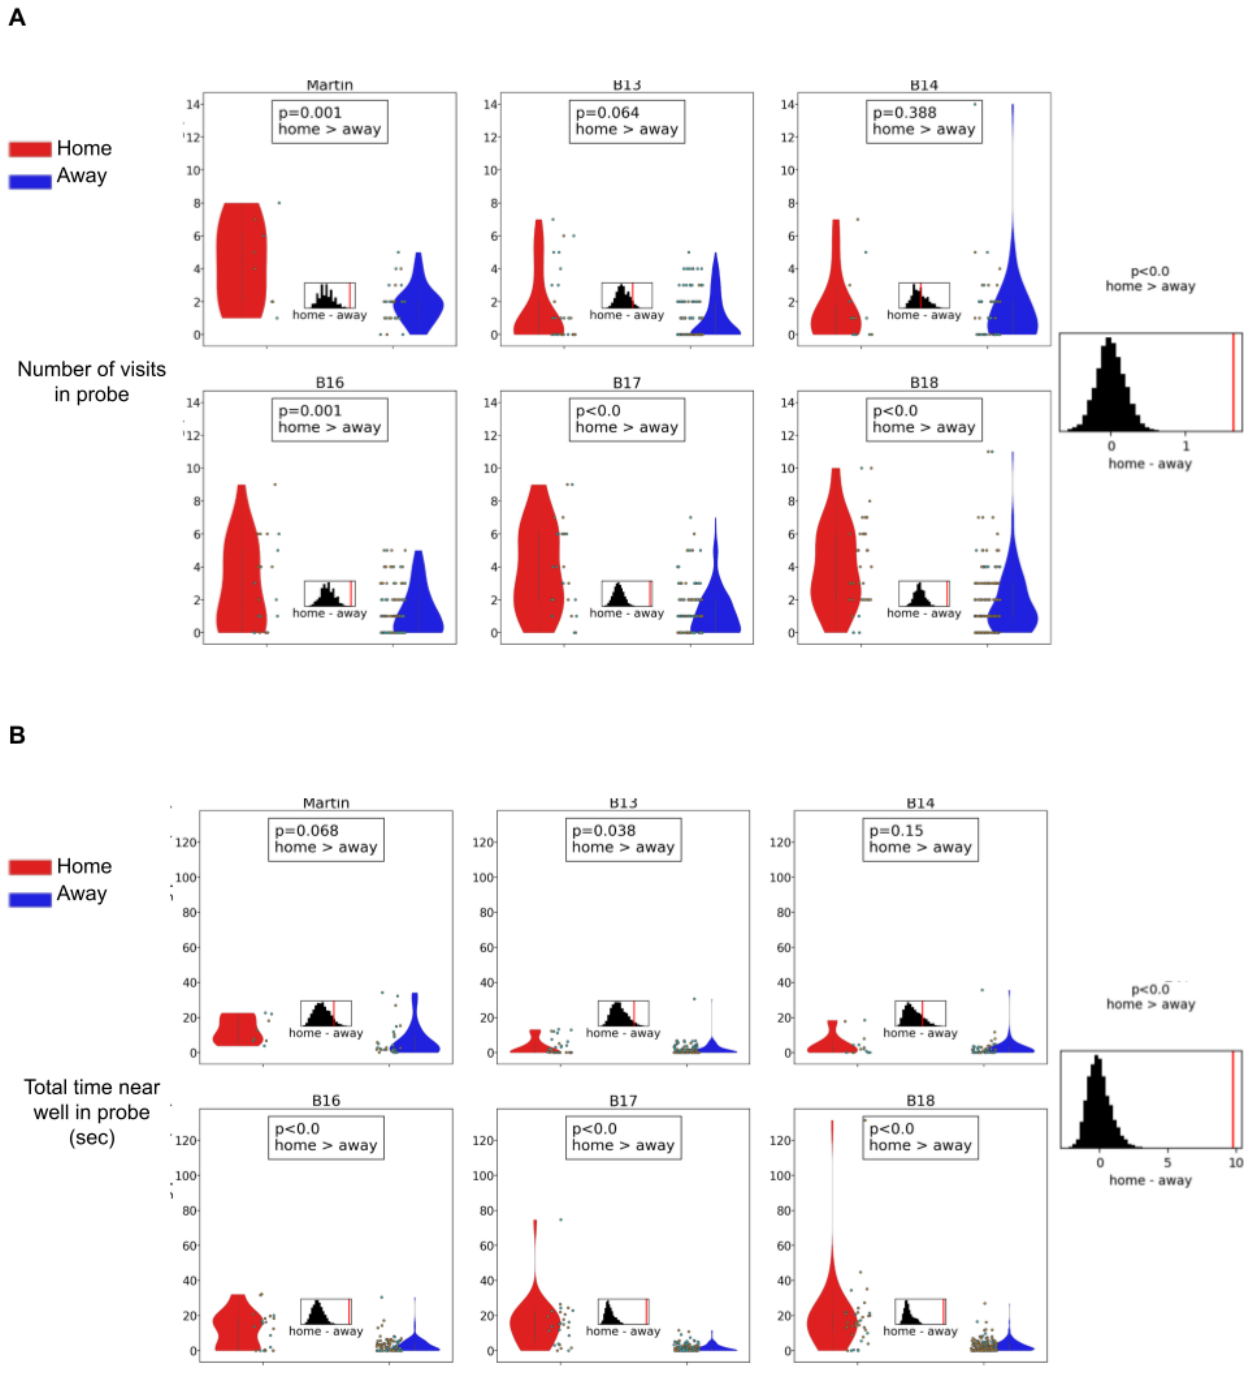

C

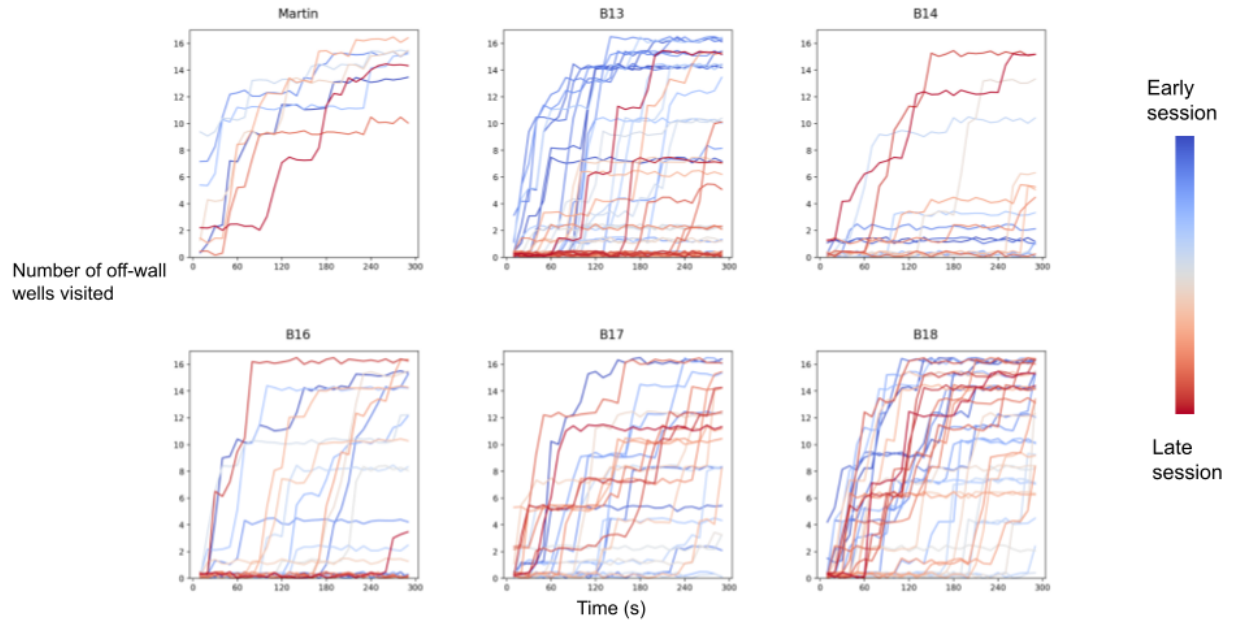
Martin
B13
B14
B16
B17
B18
Number of off-wall wells visited
Time (s)
Early session
Late session

D

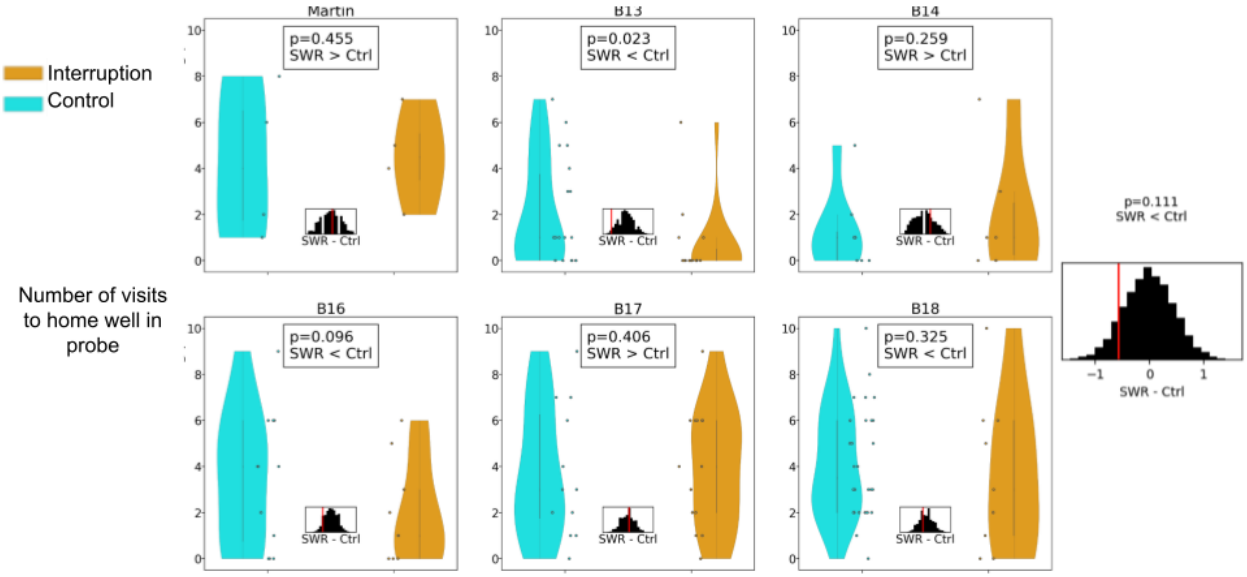
Interruption
Control
Martin
p=0.455
SWR > Ctrl
B13
p=0.023
SWR < Ctrl
B14
p=0.259
SWR > Ctrl
B16
p=0.096
SWR < Ctrl
B17
p=0.406
SWR > Ctrl
B18
p=0.325
SWR < Ctrl
SWR - Ctrl
Number of visits to home well in probe
p=0.111
SWR < Ctrl

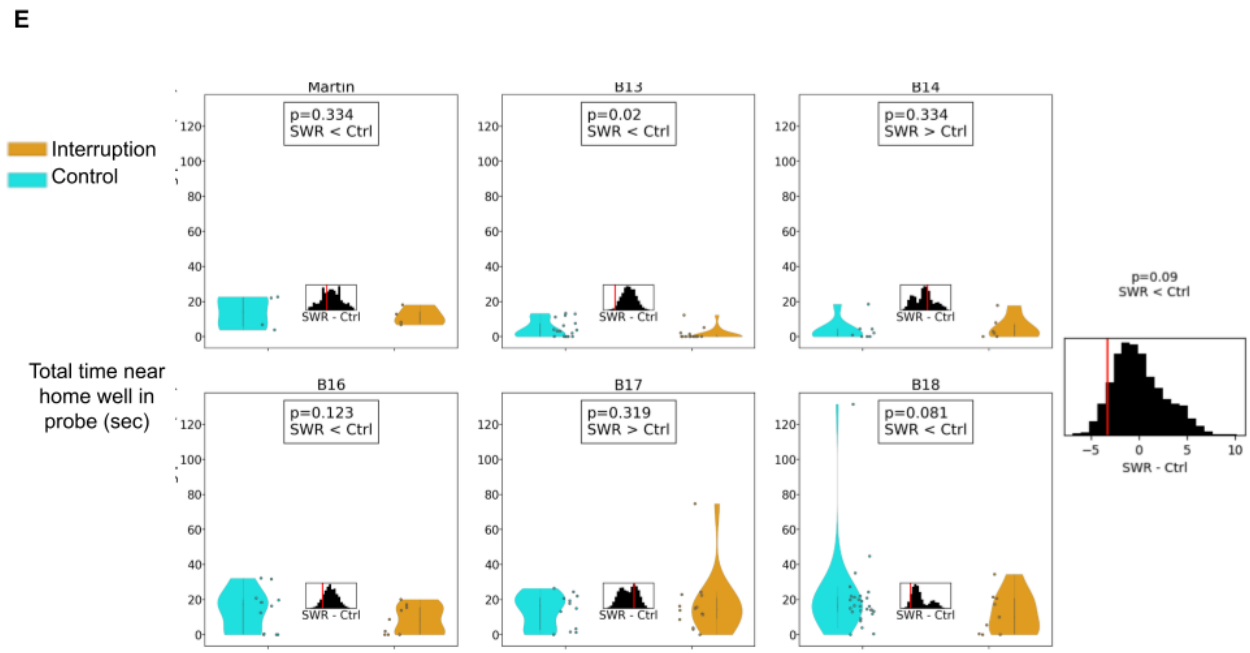

Fig 10. Home-away probe behavior. *A) Number of visits within 6 inches of the home and away wells during the probe. For a detailed description, see figure 9F. B) Total time within six inches of the home and away wells during the probe. C) Cumulative number of reward wells in the center of the environment checked during the probe. D) Number of visits to the home well during the probe E) Total time near home well during the probe*

3e) Data Mining and testing for behavioral differences

The measures introduced above offer a coarse way to quantify patterns in behavior that we hypothesized might differ between sessions with and without ripple interruption. As we gathered and analyzed the data, we began to make more granular hypotheses about potential behavioral differences. The rigor of testing these predictions begins to come into question though as we run into possible issues of overfitting to our data and running multiple comparisons. Therefore, we chose to employ a test-validation strategy in which we chose a subset of the collected data to explore and search for possible patterns, which we could then test rigorously on the rest of the data.

We chose to use B17's data to data mine for possible differences between conditions, as he had a decent number of trials from both control and interruptions sessions, and displayed consistent exploration during probes. A complete accounting of all of the methods for quantifying behavioral differences can be found in Appendix D, and their implementations can be found in the Datamining.py file in the associated codebase. We ran each of these measures with the specified range of parameters, controls, and statistical comparisons. 100 shuffles were performed for each test,

enabling a rough idea of significance without incurring an overwhelming computational time. We then manually observed the resulting p values, in order to separate measures that were significant across a range of parameters, indicating a clear difference, from measures that were not robust to parameter changes, indicating likely spurious significance. Interestingly, no robust differentiators were observed in task behavior. The previously observed differences on probe behavior were correctly identified.

We also observed interesting differences relating the interruption condition of a session to the initial behavior in the following session. We found a strong trend for rats to revisit the previous home well more times when that previous session was a control session (Fig 11a). Furthermore, rats tended to revisit the previous home well more often in this same period if they were currently in an interruption session (Fig 11b). The implications of these observations are noted in the discussion section.

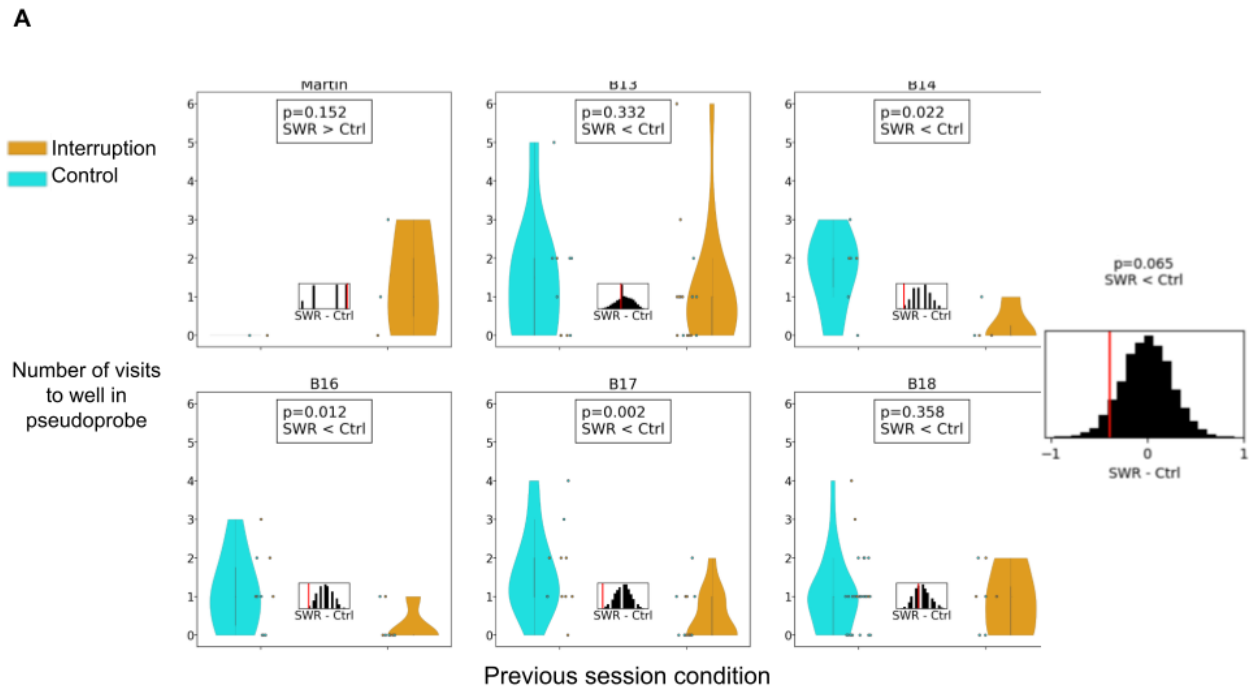

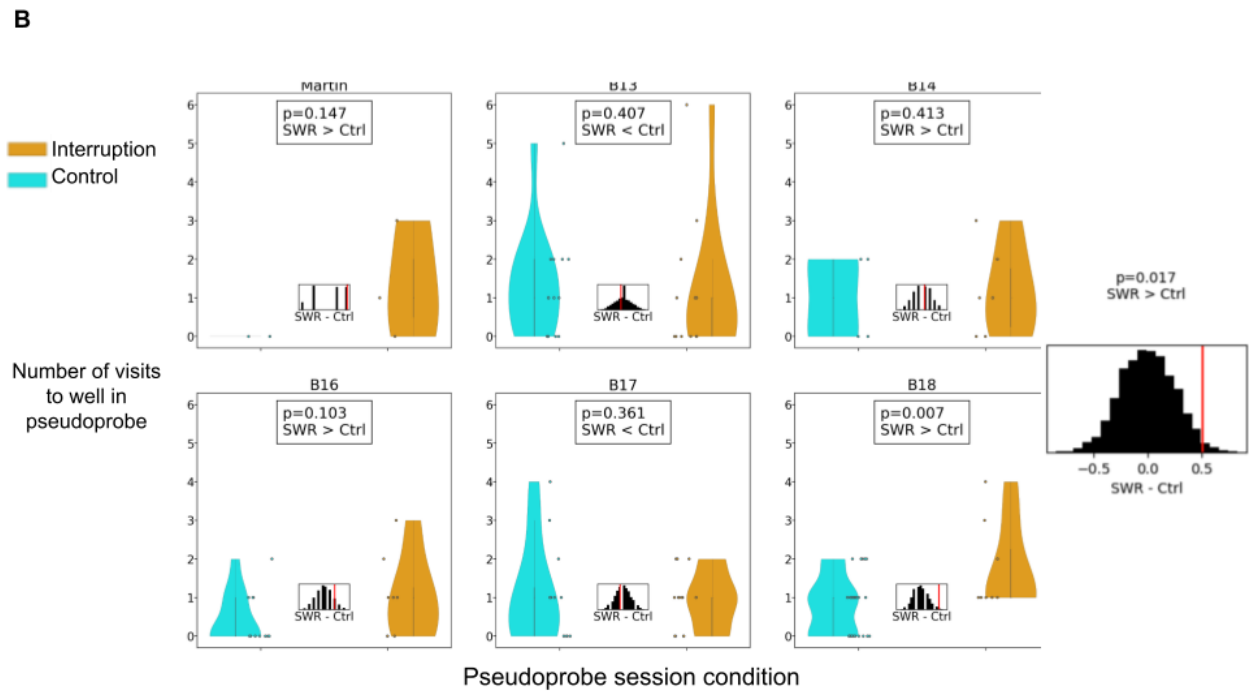

Fig 11. Home-away pseudo-probe behavior. *A) Number of visits to the home well during the pseudo-probe, divided by the session's condition. B) Number of visits to the home well during the pseudo-probe, divided by the condition during the pseudo-probe.*

3f) Plotting and statistics library

Developing the pipeline to data mine B17's data inherently incentivized me to write flexible plotting and statistics code that could be widely applied to various quantitative measures. I therefore separated the code handling the plotting and statistics into its own package which may be of wider use when others are performing similar comparisons of quantitative measures across groups. The logic behind the statistics is described briefly below. A full illustrative set of examples is also available at https://github.com/fosterlabberkeley/BillCroughanThesis .

The shuffling library is organized to test for a range of relationships among different categories of data. The tests are specified using four different types of effects. A main effect of a category refers to the difference between the mean value of datapoints within that category compared to the mean value of datapoints outside of that category. Note in the case that there are only two categories, this is equivalent to comparing the means of the two groups directly. The other three effect types are defined recursively: within, across, and interaction effects. Any length of recursion is possible, and the final effect type is always a main effect, as that is the only non-recursive effect considered. The recursive effects considered are within-group effects, in which only data from the specified group is considered; across-group effects, in which the observed

effects of the recursive effects are calculated within each group and then summed; and interaction effects, in which the difference between the recursively defined effect is compared within the specified group compared to outside of the specified group. Illustrated examples of these effects are shown in figure 12, and exact formulas are given in Appendix E.

The statistics are designed to integrate easily with the plotting functions of the library. Whenever a user creates a plot, they may choose to attach any categorical or continuous data to that plot to be analyzed. For instance, when creating the plot for each animal in figure 11, I attached the y values of each point, and their corresponding category in either the control or interruption group. They may also specify a single shuffle that should be immediately run and represented visually on the plot itself. In figure 11, I specified that the session conditions should be shuffled. Before saving the figure, the library will run this shuffle and indicate the results as a p-value, effect direction, and histogram of shuffled differences of means. After making any number of plots the user can choose to run a series of statistical tests on the data that was saved. This can include any combination of recursive effects. The examples shown in Fig 12 were automatically calculated in this way. Note that the recursive effects displayed are chosen to illustrate the patterns in the data and are not a comprehensive illustration of every effect that could be tested.

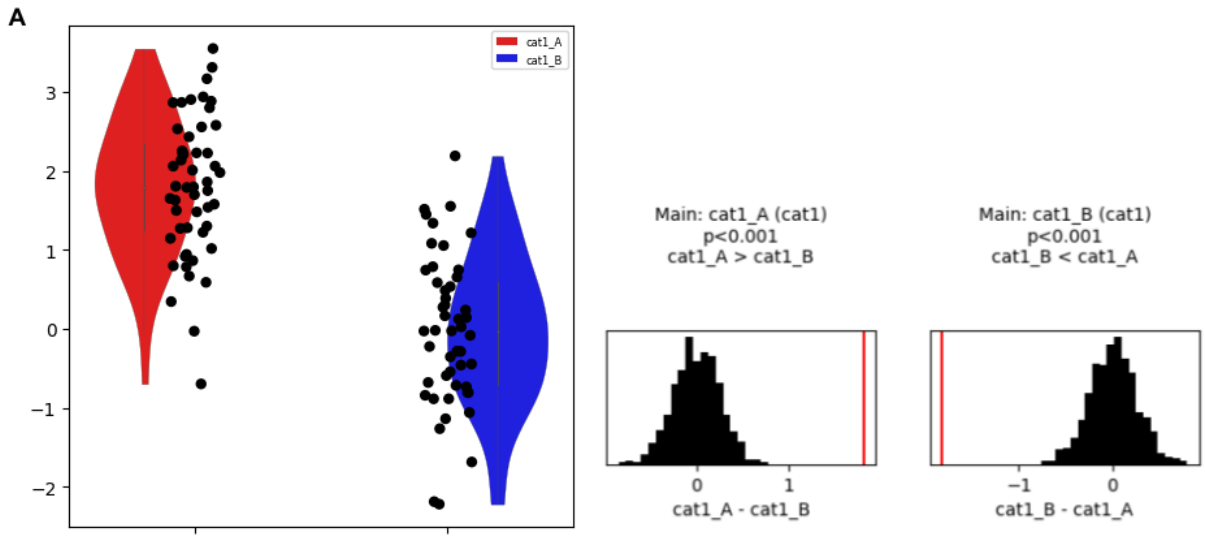

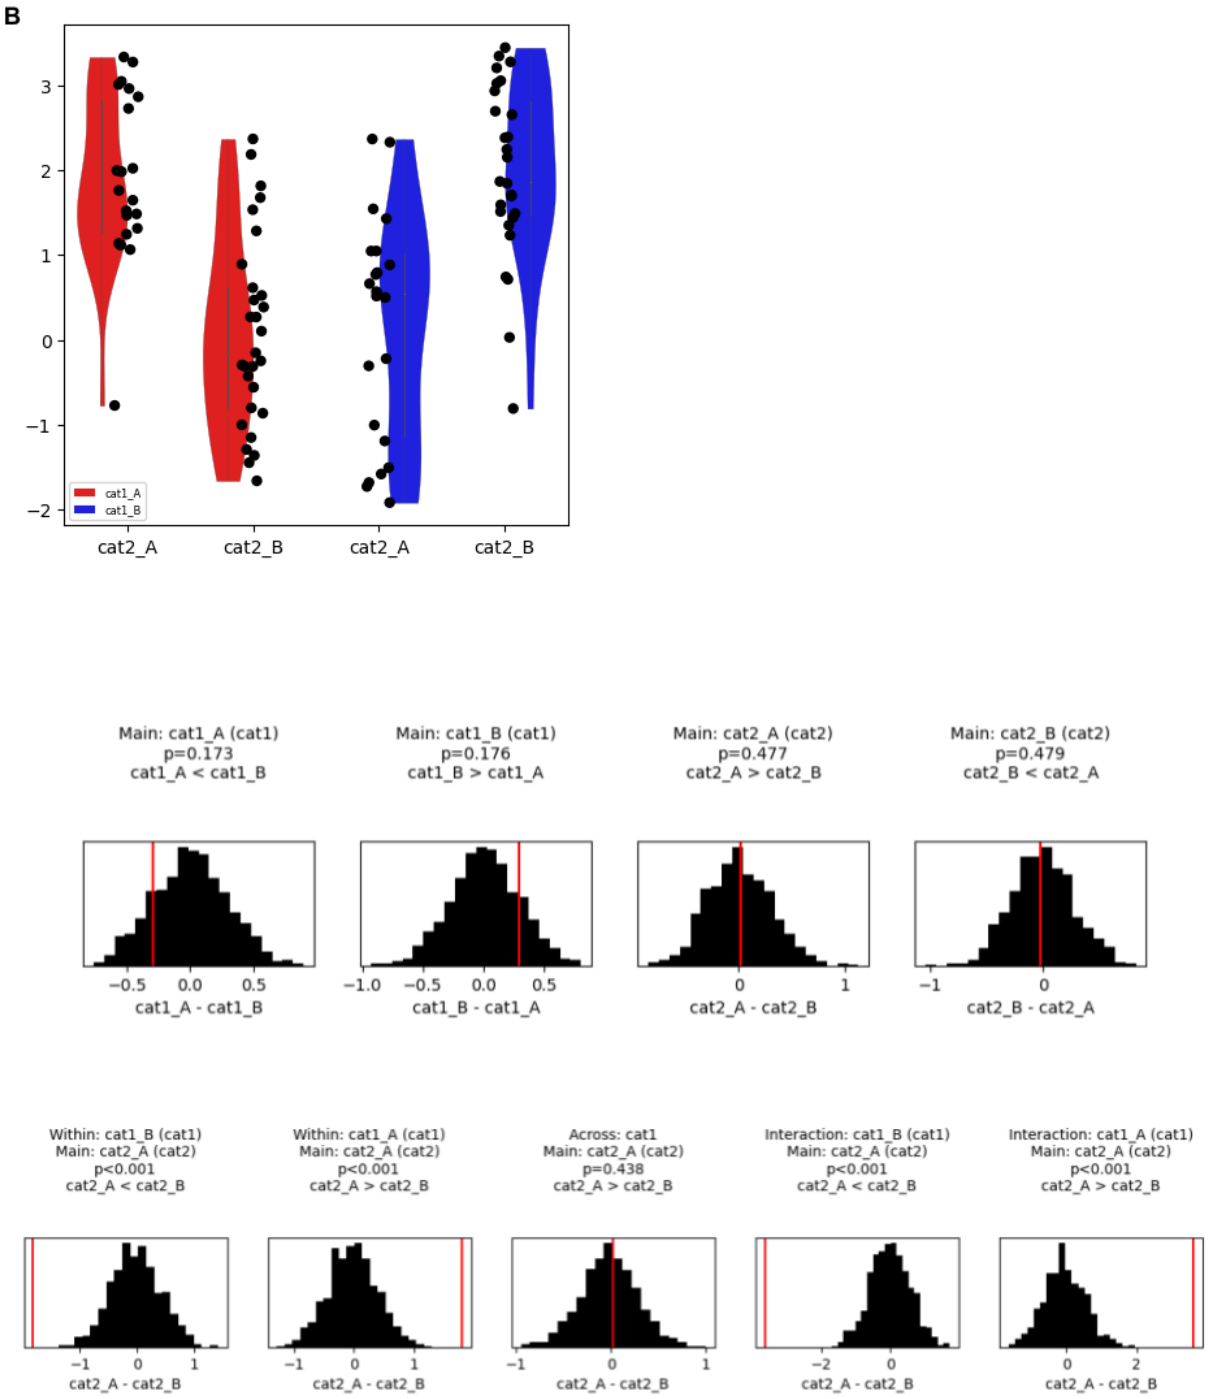
B
cat1_A
cat1_B
cat2_A
cat2_B
cat2_A
cat2_B
Main: cat1_A (cat1)
p=0.173
cat1_A < cat1_B
Main: cat1_B (cat1)
p=0.176
cat1_B > cat1_A
Main: cat2_A (cat2)
p=0.477
cat2_A > cat2_B
Main: cat2_B (cat2)
p=0.479
cat2_B < cat2_A
cat1_A - cat1_B
cat1_B - cat1_A
cat2_A - cat2_B
cat2_B - cat2_A
Within: cat1_B (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A < cat2_B
Within: cat1_A (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
Across: cat1
Main: cat2_A (cat2)
p=0.438
cat2_A > cat2_B
Interaction: cat1_B (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A < cat2_B
Interaction: cat1_A (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
cat2_A - cat2_B
cat2_A - cat2_B
cat2_A - cat2_B
cat2_A - cat2_B
cat2_A - cat2_B

C

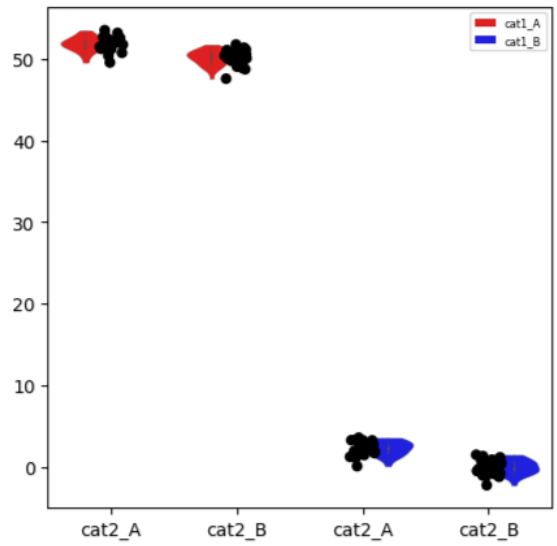
cat1_A
cat1_B
50
40
30
20
10
0
cat2_A
cat2_B
cat2_A
cat2_B

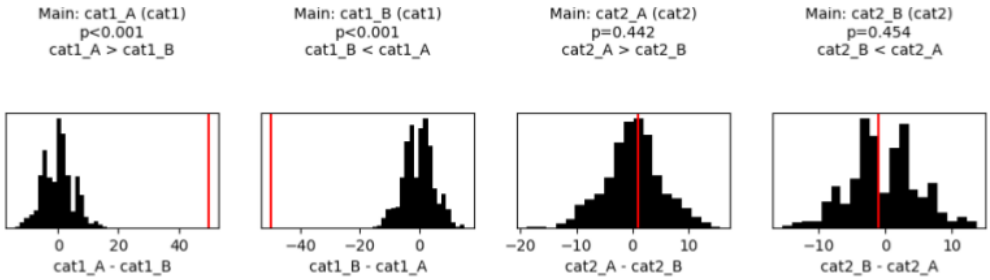
Main: cat1_A (cat1)
p<0.001
cat1_A > cat1_B
Main: cat1_B (cat1)
p<0.001
cat1_B < cat1_A
Main: cat2_A (cat2)
p=0.442
cat2_A > cat2_B
Main: cat2_B (cat2)
p=0.454
cat2_B < cat2_A
0 20 40
cat1_A - cat1_B
−40 −20 0
cat1_B - cat1_A
−20 −10 0 10
cat2_A - cat2_B
−10 0 10
cat2_B - cat2_A

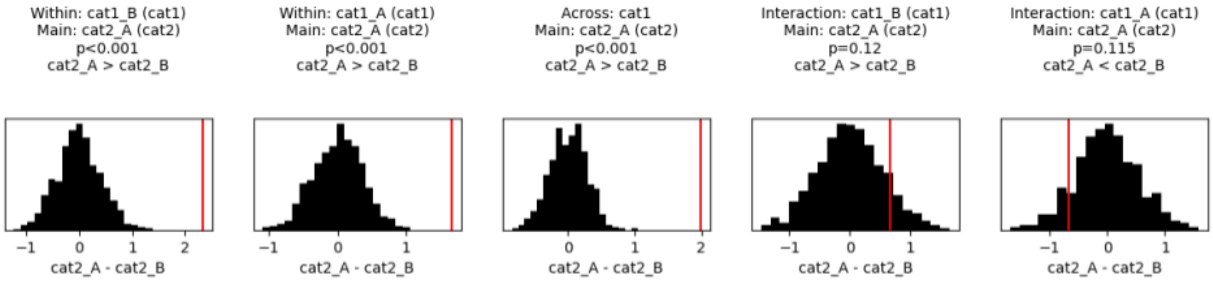
Within: cat1_B (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
Within: cat1_A (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
Across: cat1
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
Interaction: cat1_B (cat1)
Main: cat2_A (cat2)
p=0.12
cat2_A > cat2_B
Interaction: cat1_A (cat1)
Main: cat2_A (cat2)
p=0.115
cat2_A < cat2_B
−1 0 1 2
cat2_A - cat2_B
−1 0 1
cat2_A - cat2_B
0 1 2
cat2_A - cat2_B
−1 0 1
cat2_A - cat2_B
−1 0 1
cat2_A - cat2_B

D

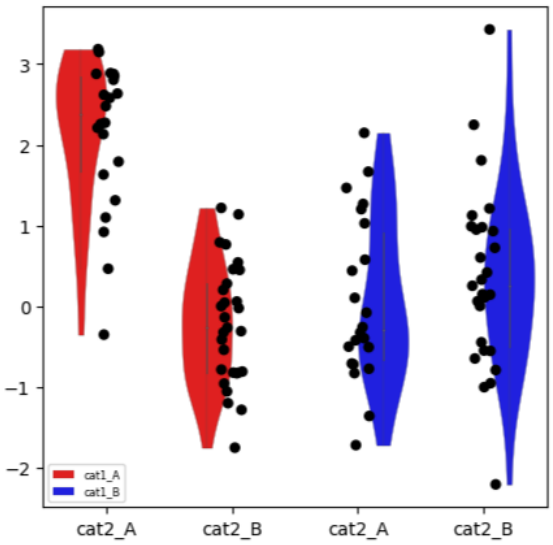
cat1_A
cat1_B
cat2_A
cat2_B
cat2_A
cat2_B

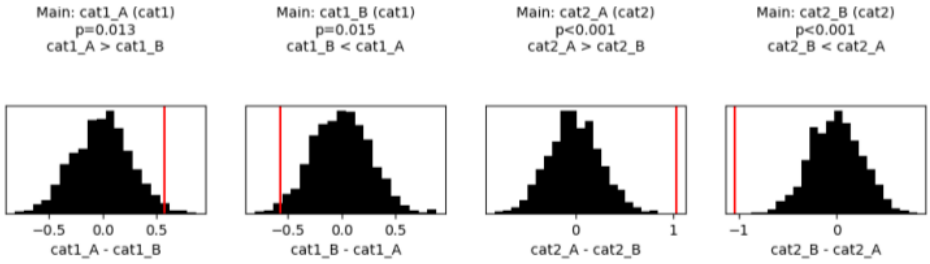
Main: cat1_A (cat1)
p=0.013
cat1_A > cat1_B
cat1_A - cat1_B
Main: cat1_B (cat1)
p=0.015
cat1_B < cat1_A
cat1_B - cat1_A
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
cat2_A - cat2_B
Main: cat2_B (cat2)
p<0.001
cat2_B < cat2_A
cat2_B - cat2_A

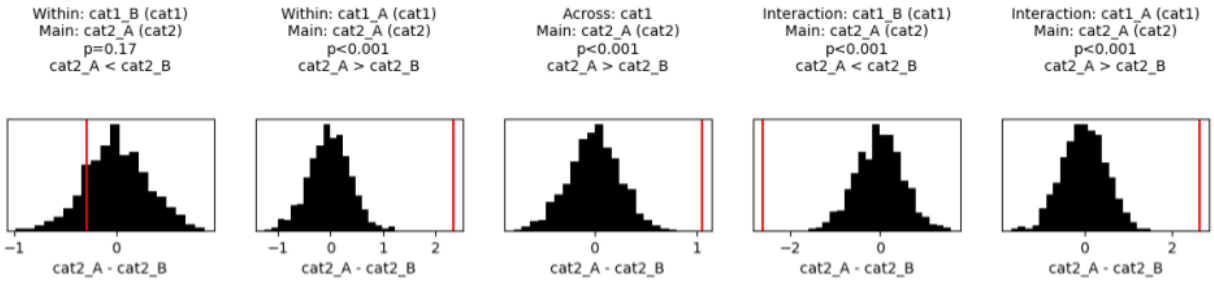
Within: cat1_B (cat1)
Main: cat2_A (cat2)
p=0.17
cat2_A < cat2_B
cat2_A - cat2_B
Within: cat1_A (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
cat2_A - cat2_B
Across: cat1
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
cat2_A - cat2_B
Interaction: cat1_B (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A < cat2_B
cat2_A - cat2_B
Interaction: cat1_A (cat1)
Main: cat2_A (cat2)
p<0.001
cat2_A > cat2_B
cat2_A - cat2_B

E

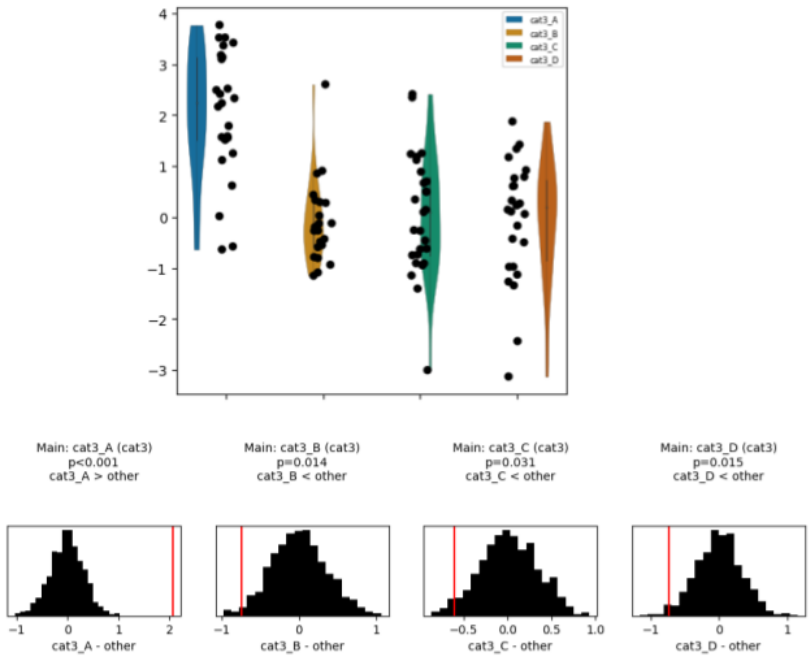

Fig 12. Example Statistical Analyses. *A) Example with one categorical variable with two values, and different means based on that categorical value. B) Example with two categories each with two values, and an interaction effect between the two categories. C) Example with two categories each with two values, and a separate effect of each category. D) Example with two categories each with two values, and a different effect of category 2 based on category 1. E) An example with one categorical variable with four values, where one category has a different mean than the rest.*

Chapter 4: Discussion

4a) CHOIR

Above, I describe a software package capable of decoding replay events and classifying them in real time. Closed-loop techniques have been used increasingly in the past decade to investigate the function and mechanism of replay (Zapata et al, 2022). This technique offers the ability to perform causal experiments, in many cases yielding much more useful data than correlational studies. These techniques are still in development, and many questions cannot be investigated using currently published methods. In particular, current closed loop methods can only effectively test hypotheses that involve all sharp wave ripples, or simple content distinctions that can be inferred from small groups of neural activity. Given that replays represent coherent trajectories through space, incorporating information about this trajectory into feedback mechanisms would allow us to probe their function and mechanism more deeply. For instance, replays can use largely overlapping groups of place cells to represent multiple directions of movement over the same area of an environment. These representations can also represent that movement proceeding forward or backward in time; these types of replays, forward and reverse, have been hypothesized to serve complimentary cognitive functions (Ambrose et al, 2016). However, they cannot reliably be differentiated by simple detection of which cells are active, and instead necessitate a full spatial decoding analysis to differentiate. CHOIR represents an important addition to the toolbox of neuroscientists hoping to investigate such phenomena.

Some limitations of our methods should be noted. In particular, our method for measuring the effect of online constraints on decoding uses two halves of a session, which are closer in time than would be the case in a similar experiment that employs CHOIR. This likely underestimates the effect of tetrode drift, which would increase with the delay between the two sessions. It may also underestimate the effectiveness of our method for updating place fields during the second session, as this was designed to combat place field remapping, which is not typically observed between two halves of a single session in a familiar environment. This could be improved in the future by using a dataset in which a rat is removed from its environment for a period of time before being placed back, while activity is recorded throughout both sessions and the intervening rest session. This would allow for tracking the same neurons between the two sessions, while also introducing a longer delay and the opportunity for remapping to play a bigger role. Another possibly more informative test scenario would be to run the same analysis in a session where some environmental or reward contingency changes. For instance, a modification of the home-away task in which the home well is changed halfway through the task. This may trigger some level of place cell remapping and illuminate the effect of this remapping on decoding accuracy.

Developing this software illustrates a more general lesson related to building tools for experiments: thinking through the complete application of the tool is often more important than optimizing the initial design. While initial results look promising for CHOIR, much emphasis during development was placed on the latency, and measurements of accuracy were not as immediately addressed. Once the benefit of modifications to improve classification accuracy was realized, these were incorporated, but the development process would have been faster had these been explored and incorporated from the beginning.

4b) Ripples and Behavior

Our work connecting sharp wave ripples to behavior offers new insights into how ripples may contribute to cognition. Published work to date has focused largely on replay's role in consolidation (Gridchyn et al, 2020; Michon et al, 2019; Girardeau et al, 2009; Ego-Stengel and Wilson, 2010), or has specifically revealed an effect during task learning, such as in the W-maze (Jadhav et al, 2012; Fernandez-Ruis et al, 2019). Our results presented above implicate replay in separate memory processes, while also demonstrating this effect in a well-trained task. Although we saw promising preliminary results on the W-maze task, we ultimately decided to use a different task for a number of reasons. As noted above, this effect is seen during task learning. Practically, this limits the amount of data that can be gathered from one subject. One of our intentions in this ripple investigation was to lay the groundwork for a follow up investigation that interacts replay content using CHOIR. Given the amount of time and effort that goes into recording from enough neurons in a single animal to effectively decode replay, we wanted to focus on tasks for which an arbitrary number of sessions could be administered. Effects seen only during learning also complicate the interpretation of these results, because the memory being encoded, consolidated, and recalled cannot be as cleanly isolated when the subject is also learning about the task itself. Additionally, the information associated with a replay through 2d space can be conjectured to relate specifically to information around the places it represents. However, the possible uses for replays in linearized environments with discrete choices is more complicated. For instance, if replays are hypothesized to be part of a simulation function in which a certain future path is tested to evaluate its outcome, a replay through a 2d environment would reveal information about the specific areas it visited. However, in a binary choice task, a replay down the unrewarding arm could be just as informative as a replay down the rewarding arm.

Our observed trend of reduced memory expression during the probe suggests a role for replay in encoding. Interruption is only present during the task itself, not the probe, meaning that this effect cannot be due to any impact on recall. The fact that our probe takes place immediately following the task without an extended rest period suggests that consolidation does not play a large role in this memory expression. Our

effect seen on pseudo-probe behavior similarly suggests an effect of interruption on encoding. These delays between sessions ranged from four hours to overnight. This is well within the range of time in which the hippocampus is known to consolidate memories with ripples. Therefore, the persistent effect seen during the pseudo-probe indicates that the memory was not as well formed at the start of this rest period, and therefore could not be recalled later after a complete opportunity for consolidation.

Finally, the effect on the pseudo-probe of ripple interruption during that session itself implicates replay in memory recall, since these interruptions occur well after any encoding and consolidation, and impact ongoing behavior at the time of decision making. Interestingly, the memory of the previous home well is higher on interruption pseudo-probes, whereas one might have hypothesized the opposite effect. One might explain this by noting that our ripple interruption cannot be considered to prevent or stop all ripples. Instead, stimulating the commissure during a ripple prevents replays from reaching their endpoint while also adding in random neural signaling which may be interpreted as aberrant information in downstream circuits. Different models of replay may suggest different outcomes in this scenario when the content of replay is modified. For instance, consider a model of replay as facilitating memory recall by allowing the rat to simulate the result of taking various paths through the environment. If the rat is relying on replay to predict the outcome of taking a path, the aberrant signaling associated with interruption may be interpreted as information about the end point of the replay. Viewed through this lens, normal replay of a previous session's home well may actually help the rat to understand the context of that memory and learn not to continue to check that older well, whereas partial or modified replay of the previous home well may hinder this process.

Although our hypothesis of replays being important for encoding during the task would fit with the results seen by Duszkievicz (Duszkievicz et al, 2023), we did not see the same effect in our data. Given the small size of this effect, it is possible we would see this same effect if we gathered data from more rats; here we present data from 6 rats compared to their 15. However, we have nearly as many sessions since we recorded from more sessions per rat on average. Another possible explanation is that their lesions encompassed more hippocampal activity than just replays, and another aspect of hippocampus activity, such as place cell activity during movement, underlies this difference. Finally, this could reflect the fact that our replay interruptions could not be comprehensive, because we had to detect sufficient ripple activity before delivering stimulation. This means that some replay activity must have been unaffected. Their lesions stopped hippocampal activity uniformly, and so would have had a more complete effect on shutting down ripples entirely. Thus, it is possible that replays are specifically necessary for this task, but the amount of ripples that were allowed to occur in our detection-interruption setup was enough to enable full performance at the task.

Of course, memory consolidation is only one framework through which to view replay. Much of the work in the W-maze has instead considered the role of ripples in the context of working memory, arguing that they are necessary for the maintenance of the short-term working memory allowing the rat to remember where it came from and integrate that into its next choice. Viewed from this perspective, the effect we observe during the probe could be interpreted as an inability to maintain such a working memory. However, our results on the pseudo-probe cannot be explained this way, and fit much more naturally into the memory consolidation framework.

Two changes to the experimental protocol could be instrumental in expanding on our experiments. The first would be to switch to using silicone probes over tetrodes. Cell yield is less of a concern than maximizing behavioral data from each animal and ensuring sufficient amount of animals. The time and effort spent adjusting tetrode depth does not contribute to these findings, whereas having immediate access to the hippocampal layer and a small number of cells that using a silicon probe would enable the same experiments to take place. The second modification would be to add a more controlled section of exploration time at the start of a trial where the pseudo-probe currently is. One such option would be to delay the first home well reward delivery by one minute. This would give a consistent amount of time for the rat to explore before the next session begins, and would enable a more rigorous comparison of behavior between sessions, and between the immediate probes and the pseudo probes at the start of the next session.

Common pitfalls of large analysis screens should be considered when interpreting our results from data mining B17's behavioral data, and the results from the algorithm optimization search for replay direction classification. For B17, care was taken not to interpret these statistical results directly. We did not consider the p-values resulting from these statistical tests to be true indications of statistical significance, but instead used them as a proxy for what may be significant when tested on the large dataset. Furthermore, instead of testing some number of these comparisons directly, we first looked at patterns in the measurements, used those to arrive at specific hypotheses, reformulated the measurements with these hypotheses in mind, and then ran those tests. In this way we essentially recreated the pattern of having preliminary data and testing that on new subjects, but were able to complete this process with only existing data. A similar concern can be expressed for the direction classification algorithm sweep. Note that in this sweep, we were not looking for a particular level of significance, but instead trying to maximize an objective function of classification accuracy, and therefore we would not fall prey to multiple comparison fallacies. Care should still be taken to avoid overfitting to our data. In this case, we tested relatively few measures, so can have reasonable confidence that these findings will generalize to other datasets. Further confirmation of these results should be done before immediately applying them in a closed-loop experiment.

References:

- Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron*, *91*(5), 1124–1136. https://doi.org/10.1016/J.NEURON.2016.07.047
- Bayley, P. J., Gold, J. J., Hopkins, R. O., & Squire, L. R. (2005). The neuroanatomy of remote memory. *Neuron*, *46*(5), 799–810. https://doi.org/10.1016/j.neuron.2005.04.034
- Berger, T. W., Song, D., Chan, R. H. M., Marmarelis, V. Z., Lacoss, J., Wills, J., Hampson, R. E., Deadwyler, S. A., & Granacki, J. J. (2012). A Hippocampal Cognitive Prosthesis: Multi-Input, Multi-Output Nonlinear Modeling and VLSI Implementation. *Ieee Transactions on Neural Systems and Rehabilitation Engineering*, *20*(2), 198. https://doi.org/10.1109/TNSRE.2012.2189133
- Born, J., & Wilhelm, I. (2012). System consolidation of memory during sleep. *Psychological Research*, *76*(2), 192–203. https://doi.org/10.1007/S00426-011-0335-6
- Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, *25*(10), 1073–1188. https://doi.org/10.1002/HIPO.22488
- Bye, C. M., & McDonald, R. J. (2019). A specific role of hippocampal nmda receptors and arc protein in rapid encoding of novel environmental representations and a more general long-term consolidation function. *Frontiers in Behavioral Neuroscience*, *13*, 8. https://doi.org/10.3389/FNBEH.2019.00008/BIBTEX
- Cao, L., Varga, V., & Chen, Z. S. (2021). Uncovering spatial representations from spatiotemporal patterns of rodent hippocampal field potentials. *Cell Reports Methods*, *1*(7), 100101. https://doi.org/10.1016/J.CRMETH.2021.100101
- Carey, A. A., Tanaka, Y., & van der Meer, M. A. A. (2019). Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nature Neuroscience*, *22*(9), 1450–1459. https://doi.org/10.1038/S41593-019-0464-6
- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience 2011 14:2*, *14*(2), 147–153. https://doi.org/10.1038/nn.2732
- Cerf, M., Thiruvengadam, N., Mormann, F., Kraskov, A., Quiroga, R. Q., Koch, C., & Fried, I. (2010). On-line, voluntary control of human temporal lobe neurons. *Nature*, *467*(7319), 1104–1108. https://doi.org/10.1038/NATURE09510
- Cheng, S., & Frank, L. M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron*, *57*(2), 303–313. https://doi.org/10.1016/J.NEURON.2007.11.035

- Chersi, F., & Pezzulo, G. (2012). Using hippocampal-striatal loops for spatial navigation and goal-directed decision-making. *Cognitive Processing*, *13 Suppl 1*(1 SUPPL). https://doi.org/10.1007/S10339-012-0475-7
- Chersi, F., Donnarumma, F., & Pezzulo, G. (2013). Mental imagery in the navigation domain: a computational model of sensory-motor simulation mechanisms. *Http://Dx.Doi.Org/10.1177/1059712313488789*, *21*(4), 251–262. https://doi.org/10.1177/1059712313488789
- Colgin, L. L. (2011). Oscillations and hippocampal-prefrontal synchrony. *Current Opinion in Neurobiology*, *21*(3), 467–474. https://doi.org/10.1016/J.CONB.2011.04.006
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal Replay of Extended Experience. *Neuron*, *63*(4), 497–507. https://doi.org/10.1016/J.NEURON.2009.07.027
- Davidson, T. L., & Jarrard, L. E. (1993). A role for hippocampus in the utilization of hunger signals. *Behavioral and Neural Biology*, *59*(2), 167–171. https://doi.org/10.1016/0163-1047(93)90925-8
- Deng, X., Liu, D. F., Kay, K., Frank, L. M., & Eden, U. T. (2015). Clusterless Decoding of Position from Multiunit Activity Using a Marked Point Process Filter. *Neural Computation*, *27*(7), 1438–1460. https://doi.org/10.1162/NECO_A_00744
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*(10), 1241–1242. https://doi.org/10.1038/NN1961
- Diekmann, N., & Cheng, S. (2023). A model of hippocampal replay driven by experience and environmental structure facilitates spatial learning. *ELife*, *12*. https://doi.org/10.7554/ELIFE.82301
- Dolan, R. J., & Dayan, P. (2013). Goals and Habits in the Brain. *Neuron*, *80*(2), 312–325. https://doi.org/10.1016/J.NEURON.2013.09.007
- Dragoi, G., & Buzsáki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*, *50*(1), 145–157. https://doi.org/10.1016/J.NEURON.2006.02.023
- Dudai, Y. (2012). The restless engram: consolidations never end. *Annual Review of Neuroscience*, *35*, 227–247. https://doi.org/10.1146/ANNUREV-NEURO-062111-150500
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, *13*(8), 995–1002. https://doi.org/10.1038/NN.2599
- Dvorkin, A., Szechtman, H., & Golani, I. (2010). Knots: Attractive Places with High Path Tortuosity in Mouse Open Field Exploration. *PLOS Computational Biology*, *6*(1), e1000638. https://doi.org/10.1371/JOURNAL.PCBI.1000638

- Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, *20*(1), 1–10. https://doi.org/10.1002/HIPO.20707
- Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: spatial navigation and beyond. *Nature Neuroscience 2017 20:11*, *20*(11), 1504–1513. https://doi.org/10.1038/nn.4656
- Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science (New York, N.Y.)*, *318*(5853), 1147–1150. https://doi.org/10.1126/SCIENCE.1148979
- Fernández-Ruiz, A., Oliva, A., de Oliveira, E. F., Rocha-Almeida, F., Tingley, D., & Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. *Science (New York, N.Y.)*, *364*(6445), 1082–1086. https://doi.org/10.1126/SCIENCE.AAX0758
- Findlay, G., Tononi, G., & Cirelli, C. (2021). The evolving view of replay and its functions in wake and sleep. *Sleep Advances : A Journal of the Sleep Research Society*, *1*(1). https://doi.org/10.1093/SLEEPADVANCES/ZPAB002
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature 2006 440:7084*, *440*(7084), 680–683. https://doi.org/10.1038/nature04587
- Foster, D., Morris, R., & Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule - Foster - 2000 - Hippocampus - Wiley Online Library. *Hippocampus*. https://onlinelibrary.wiley.com/doi/10.1002/(SICI)1098-1063(2000)10:1%3C1::AID-HIPO1%3E3.0.CO;2-1
- Foster, D. J. (2017). Replay Comes of Age. *Https://Doi.Org/10.1146/Annurev-Neuro-072116-031538*, *40*, 581–602. https://doi.org/10.1146/ANNUREV-NEURO-072116-031538
- Gillespie, A. K., Astudillo Maya, D. A., Denovellis, E. L., Liu, D. F., Kastner, D. B., Coulter, M. E., Roumis, D. K., Eden, U. T., & Frank, L. M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron*, *109*(19), 3149-3163.e6. https://doi.org/10.1016/j.neuron.2021.07.029
- Girardeau, G., & Zugaro, M. (2011). Hippocampal ripples and memory consolidation. *Current Opinion in Neurobiology*, *21*(3), 452–459. https://doi.org/10.1016/J.CONB.2011.02.005
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, *12*(10), 1222–1223. https://doi.org/10.1038/NN.2384

- Girardeau, G., Cei, A., & Zugaro, M. (2014). Learning-Induced Plasticity Regulates Hippocampal Sharp Wave-Ripple Drive. *Journal of Neuroscience*, *34*(15), 5176–5183. https://doi.org/10.1523/JNEUROSCI.4288-13.2014
- Gridchyn, I., Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2020). Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. *Neuron*, *106*(2), 291-300.e6. https://doi.org/10.1016/j.neuron.2020.01.021
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, *65*(5), 695–705. https://doi.org/10.1016/J.NEURON.2010.01.034
- Igata, H., Ikegaya, Y., & Sasaki, T. (2021). Prioritized experience replays on a hippocampal predictive map for learning. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(1). https://doi.org/10.1073/PNAS.2011266118/-/DCSUPPLEMENTAL
- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science (New York, N.Y.)*, *336*(6087), 1454–1458. https://doi.org/10.1126/SCIENCE.1217230
- Jadhav, S. P. P., Rothschild, G., Roumis, D. K. K., & Frank, L. M. M. (2016). Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. *Neuron*, *90*(1), 113–127. https://doi.org/10.1016/J.NEURON.2016.02.010
- Ji, D., & Wilson, M. A. (2006). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience 2007 10:1*, *10*(1), 100–107. https://doi.org/10.1038/nn1825
- Johnson, A., & Redish, A. D. (2005). Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. *Neural Networks*, *18*(9), 1163–1171. https://doi.org/10.1016/J.NEUNET.2005.08.009
- Johnson, A., & Redish, A. D. (2007). Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. *Journal of Neuroscience*, *27*(45), 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007
- Joo, H. R., & Frank, L. M. (2018). The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nature Reviews. Neuroscience*, *19*(12), 744–757. https://doi.org/10.1038/S41583-018-0077-1
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, *12*(7), 913–918. https://doi.org/10.1038/NN.2344
- Kennedy, P. J., & Shapiro, M. L. (2004). Retrieving Memories via Internal Context Requires the Hippocampus. *Journal of Neuroscience*, *24*(31), 6979–6985. https://doi.org/10.1523/JNEUROSCI.1388-04.2004

- Knudsen, E. B., & Wallis, J. D. (2021). Hippocampal neurons construct a map of an abstract value space. *Cell*, *184*(18), 4640-4650.e10. https://doi.org/10.1016/J.CELL.2021.07.010
- Kovács, K. A., O'Neill, J., Schoenenberger, P., Penttonen, M., Guerrero, D. K. R., & Csicsvari, J. (2016). 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*(10). https://doi.org/10.1371/JOURNAL.PONE.0164675
- Lansink, C. S., Goltstein, P. M., Lankelma, J. v., McNaughton, B. L., & Pennartz, C. M. A. (2009). Hippocampus Leads Ventral Striatum in Replay of Place-Reward Information. *PLOS Biology*, *7*(8), e1000173. https://doi.org/10.1371/JOURNAL.PBIO.1000173
- de Lavilléon, G., Lacroix, M. M., Rondi-Reig, L., & Benchenane, K. (2015). Explicit memory creation during sleep demonstrates a causal role of place cells in navigation. *Nature Neuroscience*, *18*(4), 493–495. https://doi.org/10.1038/NN.3970
- Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M. B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science (New York, N.Y.)*, *309*(5734), 619–623. https://doi.org/10.1126/SCIENCE.1114037
- Lisman, J., & Redish, A. D. (2009). Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1193. https://doi.org/10.1098/RSTB.2008.0316
- Liu, K., Sibille, J., & Dragoi, G. (2018). Generative Predictive Codes by Multiplexed Hippocampal Neuronal Tuplets. *Neuron*, *99*(6), 1329-1341.e6. https://doi.org/10.1016/J.NEURON.2018.07.047
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience 2018 21:11*, *21*(11), 1609–1617. https://doi.org/10.1038/s41593-018-0232-z
- McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neuroscience 2014 17:12*, *17*(12), 1658–1660. https://doi.org/10.1038/nn.3843
- Michon, F., Sun, J. J., Kim, C. Y., Ciliberti, D., & Kloosterman, F. (2019). Post-learning Hippocampal Replay Selectively Reinforces Spatial Memory for Highly Rewarded Locations. *Current Biology*, *29*(9), 1436-1444.e5. https://doi.org/10.1016/j.cub.2019.03.048
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement

learning. *Nature Human Behaviour 2017 1:9*, *1*(9), 680–692. https://doi.org/10.1038/s41562-017-0180-8

- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature 1982 297:5868*, *297*(5868), 681–683. https://doi.org/10.1038/297681a0
- Morris, R. G. M. (2006). Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. *The European Journal of Neuroscience*, *23*(11), 2829–2846. https://doi.org/10.1111/J.1460-9568.2006.04888.X
- Moser, E. I., Kropff, E., & Moser, M. B. (2008). Place Cells, Grid Cells, and the Brain's Spatial Representation System. *Https://Doi.Org/10.1146/Annurev.Neuro.31.061307.090723*, *31*, 69–89. https://doi.org/10.1146/ANNUREV.NEURO.31.061307.090723
- Nicola, W., & Clopath, C. (2019). A diversity of interneurons and Hebbian plasticity facilitate rapid compressible learning in the hippocampus. *Nature Neuroscience 2019 22:7*, *22*(7), 1168–1181. https://doi.org/10.1038/s41593-019-0415-2
- Nokia, M. S., Mikkonen, J. E., Penttonen, M., & Wikgren, J. (2012). Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning. *Frontiers in Behavioral Neuroscience*, *6*(NOVEMBER 2012). https://doi.org/10.3389/FNBEH.2012.00084
- Norimoto, H., Makino, K., Gao, M., Shikano, Y., Okamoto, K., Ishikawa, T., Sasaki, T., Hioki, H., Fujisawa, S., & Ikegaya, Y. (2018). Hippocampal ripples down-regulate synapses. *Science (New York, N.Y.)*, *359*(6383), 1524–1527. https://doi.org/10.1126/SCIENCE.AAO0702
- Nour, M. M., Liu, Y., Higgins, C., Woolrich, M. W., & Dolan, R. J. (2023). Reduced coupling between offline neural replay events and default mode network activation in schizophrenia. *Brain Communications*, *5*(2). https://doi.org/10.1093/BRAINCOMMS/FCAD056
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. https://doi.org/10.1016/0006-8993(71)90358-1
- O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, *11*(2), 209–215. https://doi.org/10.1038/NN2037
- Freyja Ólafsdóttir, H., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *ELife*, *4*(JUNE2015). https://doi.org/10.7554/ELIFE.06063

- Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2017). Task Demands Predict a Dynamic Switch in the Content of Awake Hippocampal Replay. *Neuron*, *96*(4), 925-935.e6. https://doi.org/10.1016/j.neuron.2017.09.035
- Ólafsdóttir, H. F., Bush, D., & Barry, C. (2018). The Role of Hippocampal Replay in Memory and Planning. *Current Biology*, *28*(1), R37–R50. https://doi.org/10.1016/J.CUB.2017.10.073
- Ono, T., Nakamura, K., Fukuda, M., & Tamura, R. (1991). Place recognition responses of neurons in monkey hippocampus. *Neuroscience Letters*, *121*(1–2), 194–198. https://doi.org/10.1016/0304-3940(91)90683-K
- Penagos, H., Varela, C., & Wilson, M. A. (2017). Oscillations, neural computations and learning during wake and sleep. *Current Opinion in Neurobiology*, *44*, 193–201. https://doi.org/10.1016/J.CONB.2017.05.009
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., & Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience*, *12*(7), 919–926. https://doi.org/10.1038/NN.2337
- Pezzulo, G., Rigoli, F., & Chersi, F. (2013). The mixed instrumental controller: Using value of information to combine habitual choice and mental simulation. *Frontiers in Psychology*, *4*(MAR), 92. https://doi.org/10.3389/FPSYG.2013.00092/BIBTEX
- Pezzulo, G., van der Meer, M. A. A., Lansink, C. S., & Pennartz, C. M. A. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences*, *18*(12), 647–657. https://doi.org/10.1016/J.TICS.2014.06.011
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature 2013 497:7447*, *497*(7447), 74–79. https://doi.org/10.1038/nature12112
- Preston, A. R., & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology : CB*, *23*(17), R764. https://doi.org/10.1016/J.CUB.2013.05.041
- YL, Q., BL, M., WE, S., & CA, B. (1997). Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *352*(1360), 1525–1533. https://doi.org/10.1098/RSTB.1997.0139
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience 2016 17:3*, *17*(3), 147–159. https://doi.org/10.1038/nrn.2015.30
- Roux, L., Hu, B., Eichler, R., Stark, E., & Buzsáki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nature Neuroscience*, *20*(6), 845–853. https://doi.org/10.1038/NN.4543

- Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2015). Prioritized Experience Replay. *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*. https://arxiv.org/abs/1511.05952v4
- Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. *Neuron*, *104*(6), 1110-1125.e7. https://doi.org/10.1016/J.NEURON.2019.09.012/ATTACHMENT/9526F7BE-6AD5-4B2B-A5EC-3F175A31762D/MMC1.PDF
- Singer, A. C., & Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, *64*(6), 910–921. https://doi.org/10.1016/J.NEURON.2009.11.016
- Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR Activity Predicts Correct Decisions during the Initial Learning of an Alternation Task. *Neuron*, *77*(6), 1163–1173. https://doi.org/10.1016/j.neuron.2013.01.027
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*(2), 149–172. https://doi.org/10.1002/(SICI)1098-1063(1996)6:2
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience 2017 20:11*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650
- Stella, F., Baracskay, P., O'Neill, J., & Csicsvari, J. (2019). Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. *Neuron*, *102*(2), 450-461.e7. https://doi.org/10.1016/j.neuron.2019.01.052
- Knudsen, E. B., & Wallis, J. D. (2021). Hippocampal neurons construct a map of an abstract value space. *Cell*, *184*(18), 4640-4650.e10. https://doi.org/10.1016/J.CELL.2021.07.010
- van der Meer, M., Kurth-Nelson, Z., & Redish, A. D. (2012). Information Processing in Decision-Making Systems. *Http://Dx.Doi.Org/10.1177/1073858411435128*, *18*(4), 342–359. https://doi.org/10.1177/1073858411435128
- van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, *92*(5), 968–974. https://doi.org/10.1016/j.neuron.2016.10.020
- Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Sleep inspires insight. *Nature*, *427*(6972), 352–355. https://doi.org/10.1038/NATURE02223

- Wikenheiser, A. M., & Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, *18*(2), 289–294. https://doi.org/10.1038/NN.3909
- Wu, C. T., Haggerty, D., Kemere, C., & Ji, D. (2017). Hippocampal awake replay in fear memory retrieval. *Nature Neuroscience*, *20*(4), 571–580. https://doi.org/10.1038/NN.4507
- Wu, X., & Foster, D. J. (2014). Hippocampal Replay Captures the Unique Topological Structure of a Novel Environment. *Journal of Neuroscience*, *34*(19), 6459–6469. https://doi.org/10.1523/JNEUROSCI.3414-13.2014
- Xu, H., Baracskay, P., O'Neill, J., & Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze. *Neuron*, *101*(1), 119-132.e4. https://doi.org/10.1016/j.neuron.2018.11.015
- Yu, J. Y., & Frank, L. M. (2015). Hippocampal-cortical interaction in decision making. *Neurobiology of Learning and Memory*, *117*, 34–41. https://doi.org/10.1016/J.NLM.2014.02.002
- Aleman-Zapata, A., van der Meij, J., & Genzel, L. (2022). Disrupting ripples: Methods, results, and caveats in closed-loop approaches in rodents. *Journal of Sleep Research*, *31*(6). https://doi.org/10.1111/JSR.13532
- Zielinski, M. C., Tang, W., & Jadhav, S. P. (2020). The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus*, *30*(1), 60–72. https://doi.org/10.1002/HIPO.22821

Appendix A: Bayesian decoding

Assumptions:

- Neurons' firing follows a poisson distribution, with a mean rate dependent only on the rat's location.
- Each neuron's firing rate given the rat's location is independent of all other neuron's firing rates

Definitions:

$$
\begin{aligned}
X_t &= \text{position at time } t \\
N_{i,t} &= \text{firing rate of neuron } i \text{ at time } t \\
\lambda_{i,X_t} &= \text{average firing rate of neuron } i \text{ at position } X_t \\
N(a\,;\mu,\sigma) &= \text{Gaussian with parameters } \mu \text{ and } \sigma \text{ evaluated at } a \\
I &= \text{time frame with maximal firing rate within an event} \\
\text{poisson}(a\,;\lambda) &= \text{poisson distribution with mean } \lambda \text{ evaluated at } a
\end{aligned}
$$

Standard bayesian decoding with uniform prior

$$
\begin{aligned}
\mathbb{P}(X_t|\vec{N}) &= \frac{\mathbb{P}(\vec{N}|X_t)\cdot\mathbb{P}(X_t)}{\mathbb{P}(\vec{N})} \\
\mathbb{P}(X_t) &= \text{constant} \\
\\
\implies \mathbb{P}(X_t|\vec{N}) &\propto \mathbb{P}(\vec{N}|X_t) \\
\mathbb{P}(\vec{N}|X_t) &= \prod_{i\in\text{neurons}} \mathbb{P}(N_{i,t}|X_t) \\
&= \prod_{i\in\text{neurons}} \text{poisson}(N_{i,t}|\lambda_{i,X_t})
\end{aligned}
$$

Gaussian continuity prior based on nearby time frames

$$
\begin{aligned}
\mathbb{P}(X_I) &= \text{constant} \\
\mathbb{P}(X_{i>I}) &= N(X_i\,; X_{i-1}, \sigma) \\
\mathbb{P}(X_{i<I}) &= N(X_i\,; X_{i+1}, \sigma)
\end{aligned}
$$

$$
\begin{aligned}
\implies \mathbb{P}(X_I|\vec{N}) &\propto \prod_{i\in\text{neurons}} \text{poisson}(N_{i,t}\,; \lambda_{i,X_t}) \\
\mathbb{P}(X_{i\neq I}|\vec{N}) &\propto \left[\prod_{i\in\text{neurons}} \text{poisson}(N_{i,t}\,; \lambda_{i,X_t})\right] \\
&\quad \cdot N(X_i\,; X_{i\pm1}, \sigma)
\end{aligned}
$$

Movement prior based on velocity between nearby time frames

$$
\begin{aligned}
\mathbb{P}(X_I) &= \text{constant} \\
\mathbb{P}(X_{I+1}) &= N(||X_i - X_{i-1}||\,; \mu, \sigma) \\
\mathbb{P}(X_{I-1}) &= N(||X_i - X_{i+1}||\,; \mu, \sigma) \\
\mathbb{P}(X_{i>I+1}) &= N(||X_i - X_{i-1}||\,; \mu, \sigma) \cdot \left[1 + cos(\angle(\overrightarrow{X_i, X_{i-1}}, \overrightarrow{X_{i-1}, X_{i-2}}))\right] \\
\mathbb{P}(X_{i<I-1}) &= N(||X_i - X_{i+1}||\,; \mu, \sigma) \cdot \left[1 + cos(\angle(\overrightarrow{X_i, X_{i+1}}, \overrightarrow{X_{i+1}, X_{i+2}}))\right]
\end{aligned}
$$

Appendix B: CHOIR replay direction classification

Three components of CHOIR were varied in order to test the accuracy of replay classification: clustering, detection, and decision criteria. Replays were classified as either forward or reverse, and compared against a ground truth established using offline clustering and detection. The offline detection algorithm with which ground truth was established detected replay candidates as peaks in ripple power, and then classified them using weighted correlation. Weighted correlation measures the extent to which place bins with high probability move across the environment smoothly over time, as described in Wu and Foster, 2014. Values range from -1 to 1, with negative values indicating movement in one direction, and positive values the other.

Clustering:

- Winclust: A graph-based algorithm developed by David Foster in Matt Wilson's lab. The algorithm considers each action potential in 4d-space, where each dimension is the amplitude of a separate tetrode channel at the time of the peak of the waveform. Iteratively larger distance thresholds are considered. At each threshold, a graph is constructed where each node is an action potential and an edge exists between two nodes if they are closer than the considered distance threshold. Connected components are found, and if they exceed a minimum number of nodes, they are considered to be a cluster. As the distance threshold expands, new action potentials are assigned to an existing cluster if they are connected to only that cluster. In this way, groups of action potentials that are closer to each other than to surrounding spikes are identified.
- BIRCH and agglomerative clustering: This method harnesses existing algorithms in series. First, principal components of the action potential waveforms are extracted. Then, BIRCH is run on these principal components to assign a hierarchical clustering to the action potentials. Finally, agglomerative clustering combines these clusters up to a set temperature threshold.

Note also that since running this analysis, Mountainsort has established itself as an efficient and powerful method for automatically clustering large amounts of action potentials, and should be considered for future applications of this software.

Event Detection:

- Ripple power threshold: This is the simplest method, wherein replays are considered as soon as the z-score of the power in the ripple band (150-250 Hz) of the LFP signal exceeds a set threshold of 3.
- Ripple power and weighted correlation threshold: In this method, a replay is considered at the point its ripple power exceeds a z-score of 3, and the absolute value of the weighted correlation exceeds a set threshold.

- Ripple power and decoding content threshold: In this method, a replay is considered at the point the ripple power exceeds a z-score of 3, and a majority of the decoded spatial content lies in the center of the environment, indicating a replay - which typically runs from one end to the other - is roughly halfway across the environment.

Decision Criteria:

- Local start, posterior sum: This method assumes that the replay event begins at the current location of the rat, which is the case a majority but not all of the time. It then determines whether a replay is forward or reverse by determining whether the map leading to the rat's current location or away from that location has more total decoded posterior.
- Weighted correlation, posterior sum: This method considered the map that has more total decoded posterior, but then decides which direction the replay is proceeding along that map by looking at the sign of the weighted correlation of that map's posterior distribution.
- Simple posterior sum: Replays are assumed to start at one end of the linear environment. Each map is divided into the two halves of the environment, and the posterior in that half is summed and added to the opposite half of the environment on the opposite map. In this way, any posterior that would indicate forward replay is totalled, and compared to the total posterior that would indicate reverse replay. The directionality is then decided based on which total is greater.

Appendix C: Results with delayed-stimulation controls only

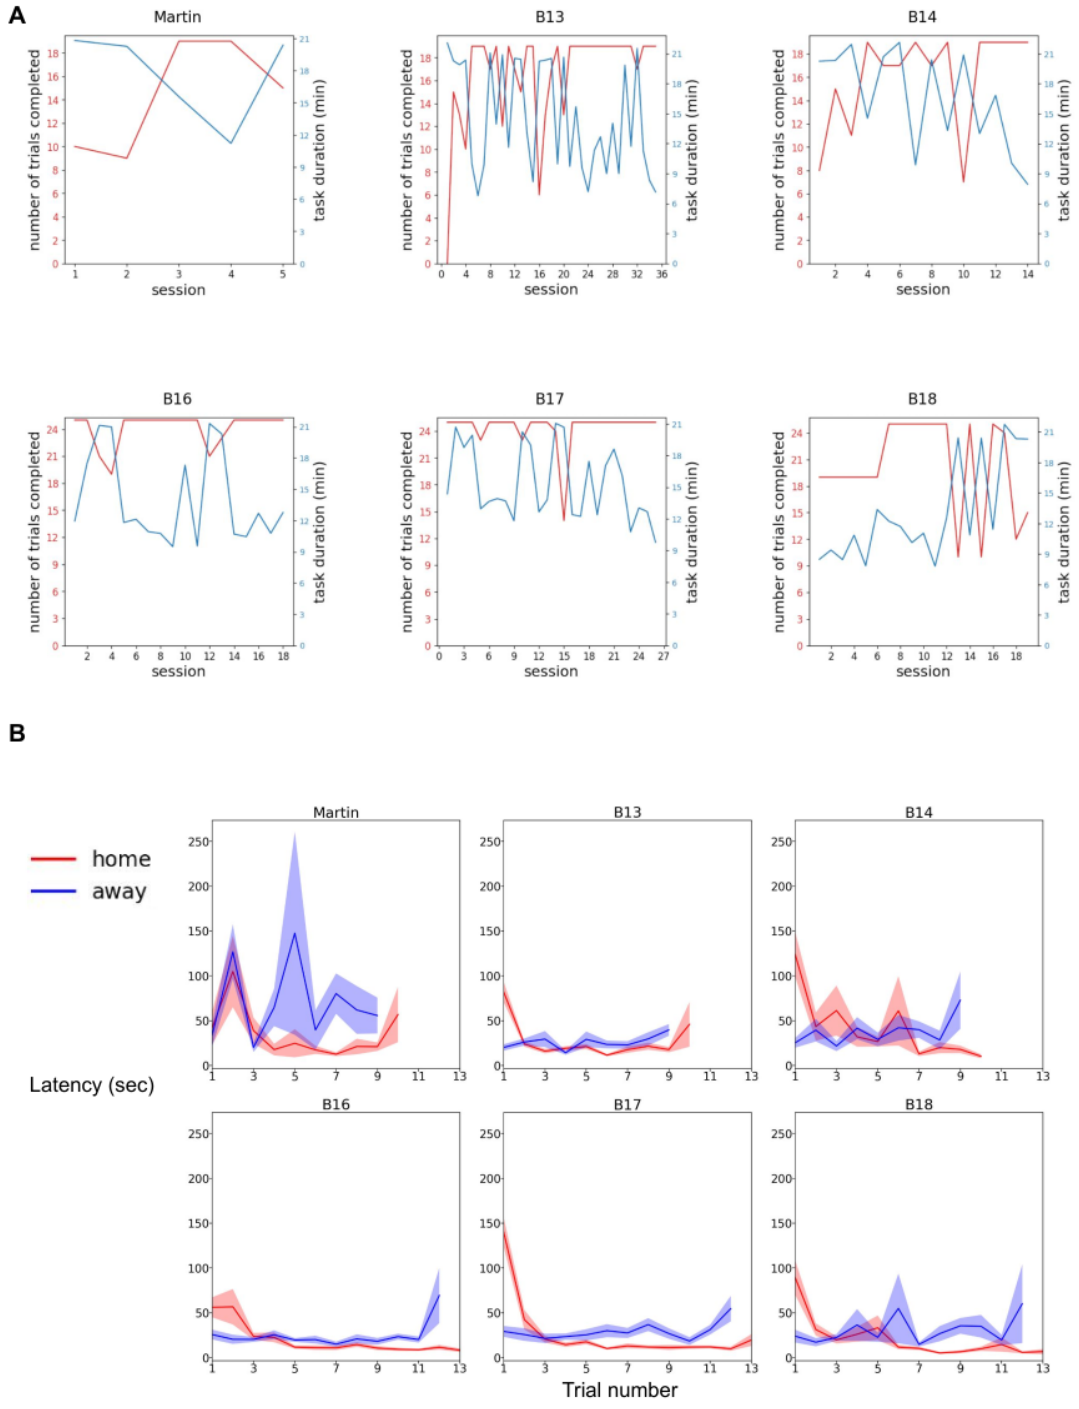

C

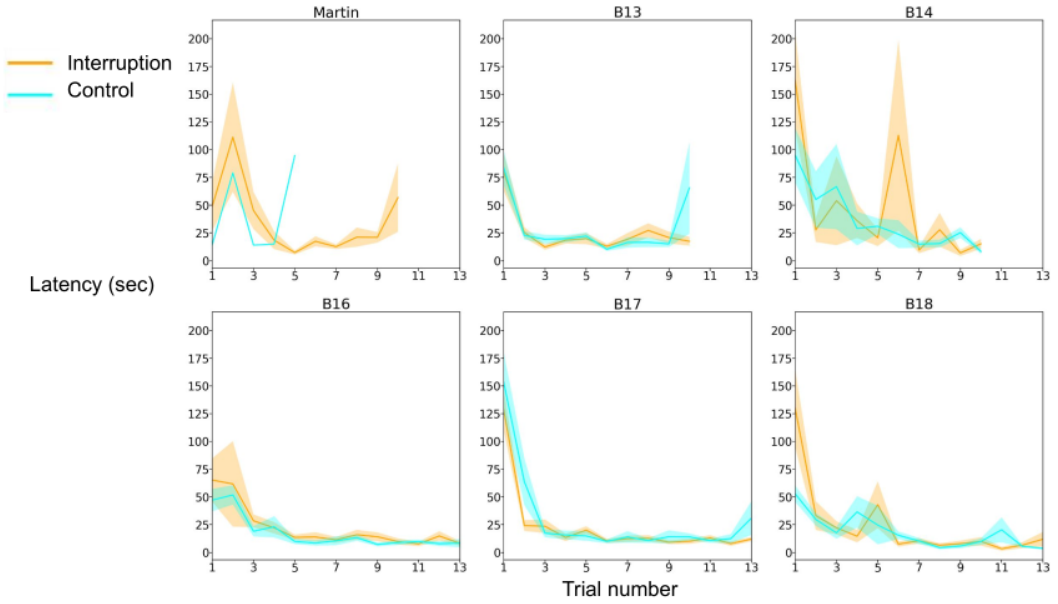
Martin
B13
B14
B16
B17
B18
Interruption
Control
Latency (sec)
Trial number

D

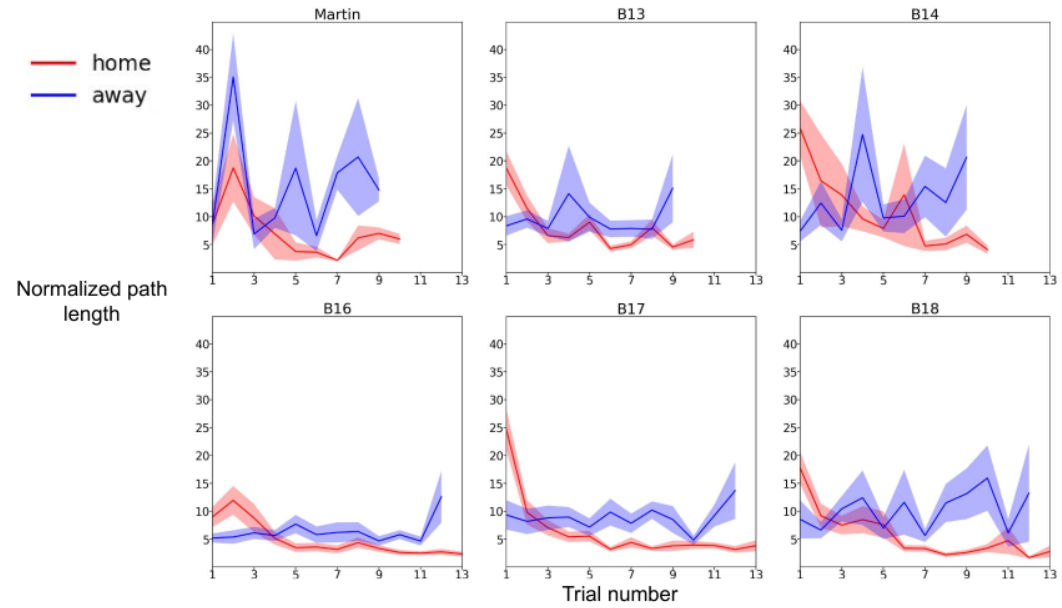
Martin
B13
B14
B16
B17
B18
home
away
Normalized path length
Trial number

E

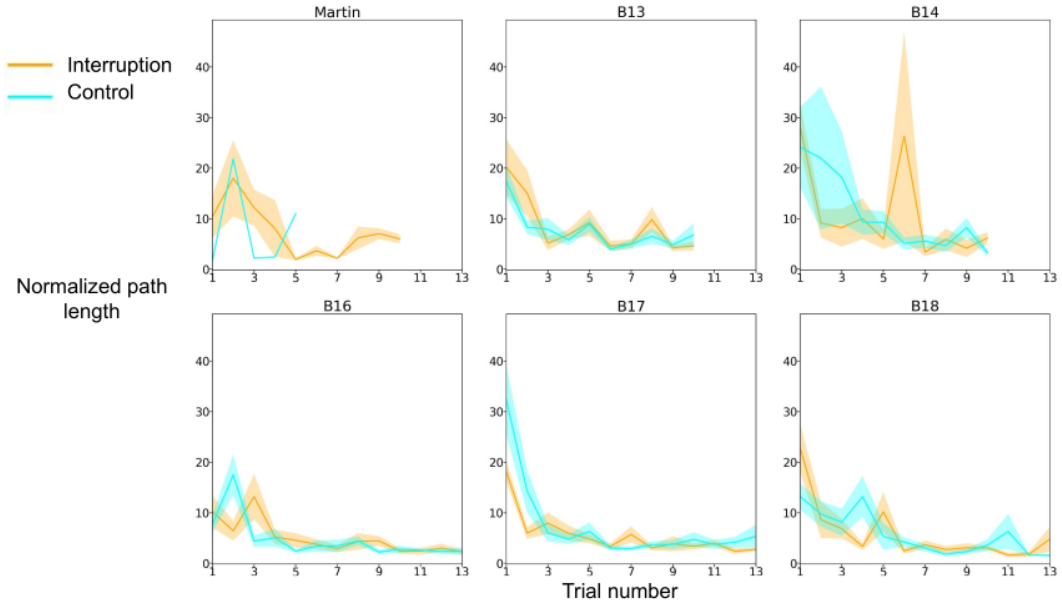
Interruption
Control
Martin
B13
B14
B16
B17
B18
Normalized path length
Trial number

F

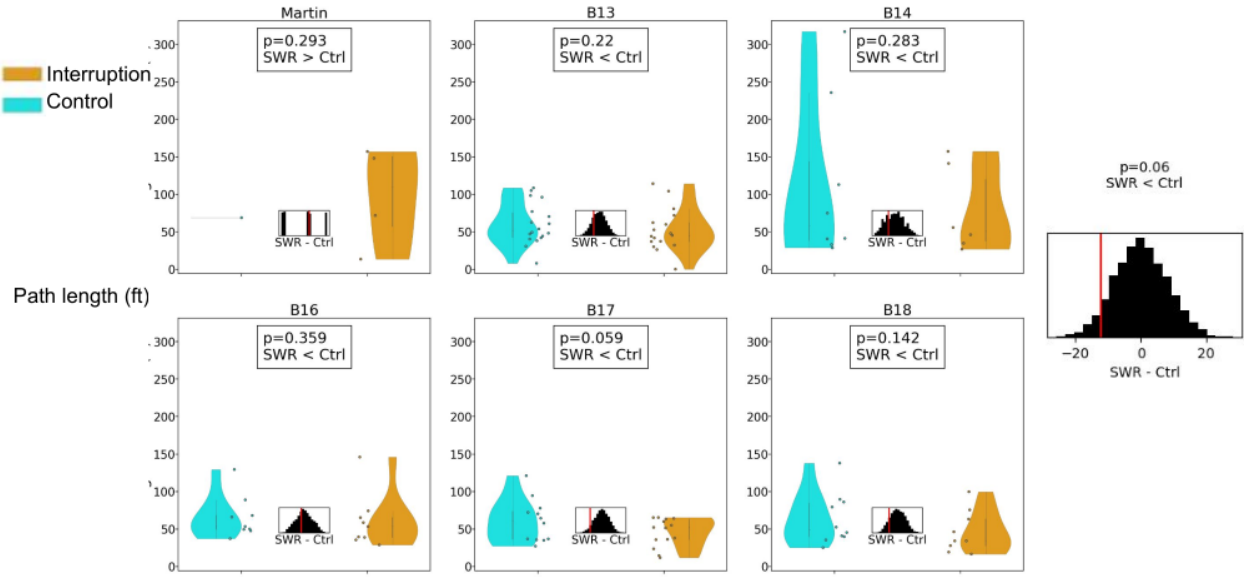
Interruption
Control
Martin
p=0.293
SWR > Ctrl
B13
p=0.22
SWR < Ctrl
B14
p=0.283
SWR < Ctrl
B16
p=0.359
SWR < Ctrl
B17
p=0.059
SWR < Ctrl
B18
p=0.142
SWR < Ctrl
SWR - Ctrl
Path length (ft)
p=0.06
SWR < Ctrl

G

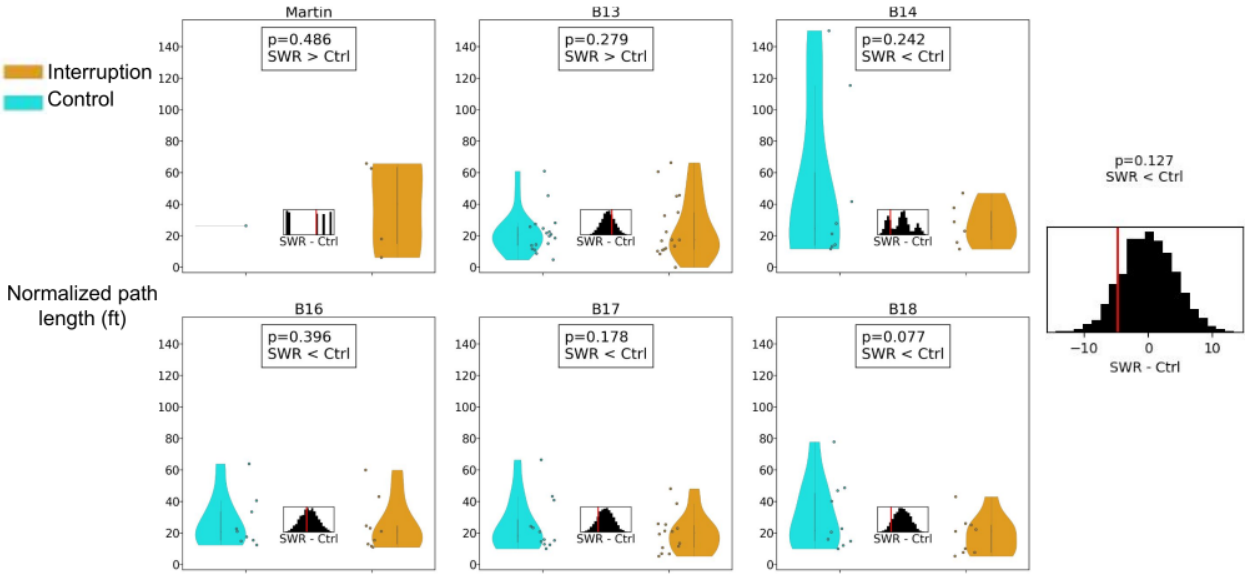
Interruption
Control
Normalized path length (ft)
Martin
p=0.486
SWR > Ctrl
B13
p=0.279
SWR > Ctrl
B14
p=0.242
SWR < Ctrl
B16
p=0.396
SWR < Ctrl
B17
p=0.178
SWR < Ctrl
B18
p=0.077
SWR < Ctrl
p=0.127
SWR < Ctrl
SWR - Ctrl

H

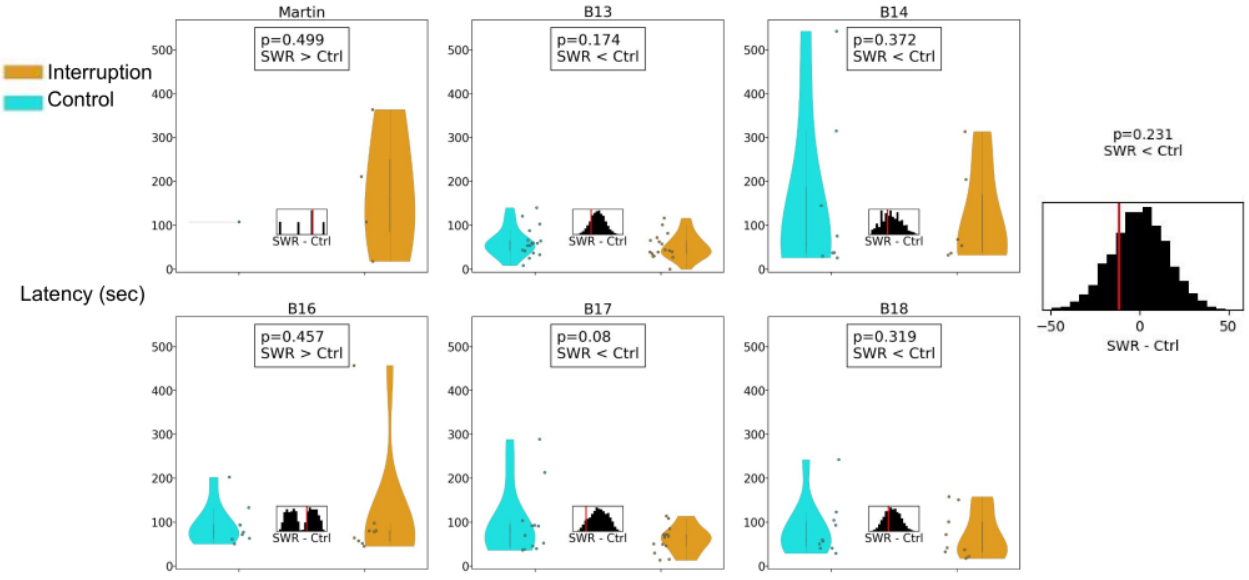
Interruption
Control
Latency (sec)
Martin
p=0.499
SWR > Ctrl
B13
p=0.174
SWR < Ctrl
B14
p=0.372
SWR < Ctrl
B16
p=0.457
SWR > Ctrl
B17
p=0.08
SWR < Ctrl
B18
p=0.319
SWR < Ctrl
p=0.231
SWR < Ctrl
SWR - Ctrl

I

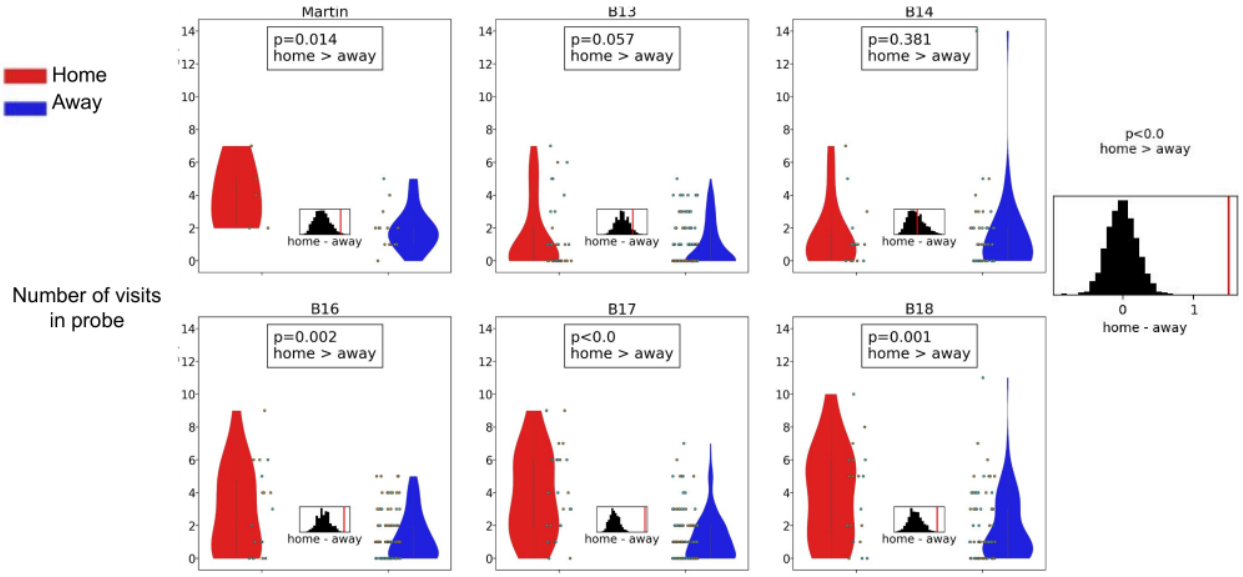
Home
Away
Number of visits in probe
Martin
p=0.014
home > away
B13
p=0.057
home > away
B14
p=0.381
home > away
B16
p=0.002
home > away
B17
p<0.0
home > away
B18
p=0.001
home > away
home - away
p<0.0
home > away
0
1
home - away

J

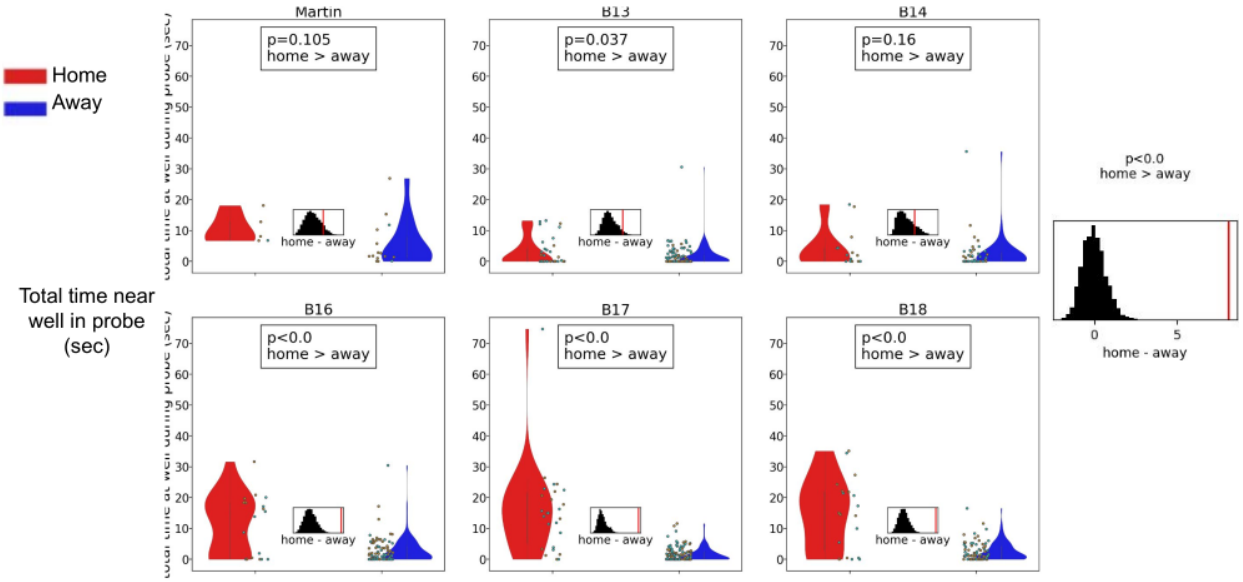
Home
Away
Total time near well in probe (sec)
Martin
p=0.105
home > away
B13
p=0.037
home > away
B14
p=0.16
home > away
B16
p<0.0
home > away
B17
p<0.0
home > away
B18
p<0.0
home > away
home - away
p<0.0
home > away
0
5
home - away

K

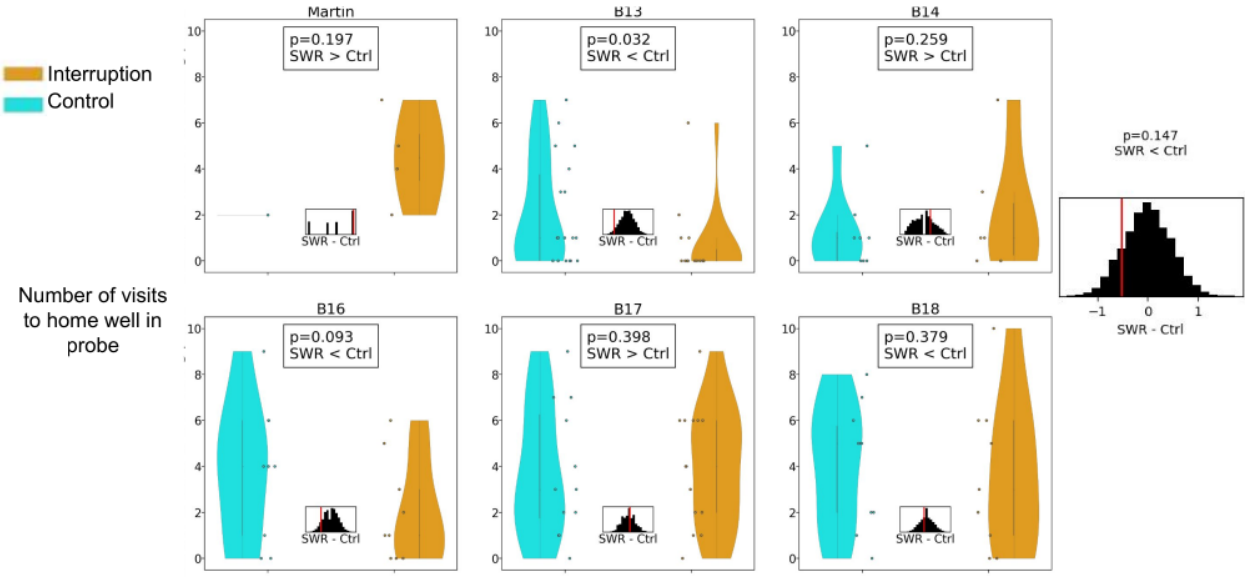
Interruption
Control
Number of visits to home well in probe
Martin
p=0.197
SWR > Ctrl
B13
p=0.032
SWR < Ctrl
B14
p=0.259
SWR > Ctrl
p=0.147
SWR < Ctrl
B16
p=0.093
SWR < Ctrl
B17
p=0.398
SWR > Ctrl
B18
p=0.379
SWR < Ctrl
SWR - Ctrl

L

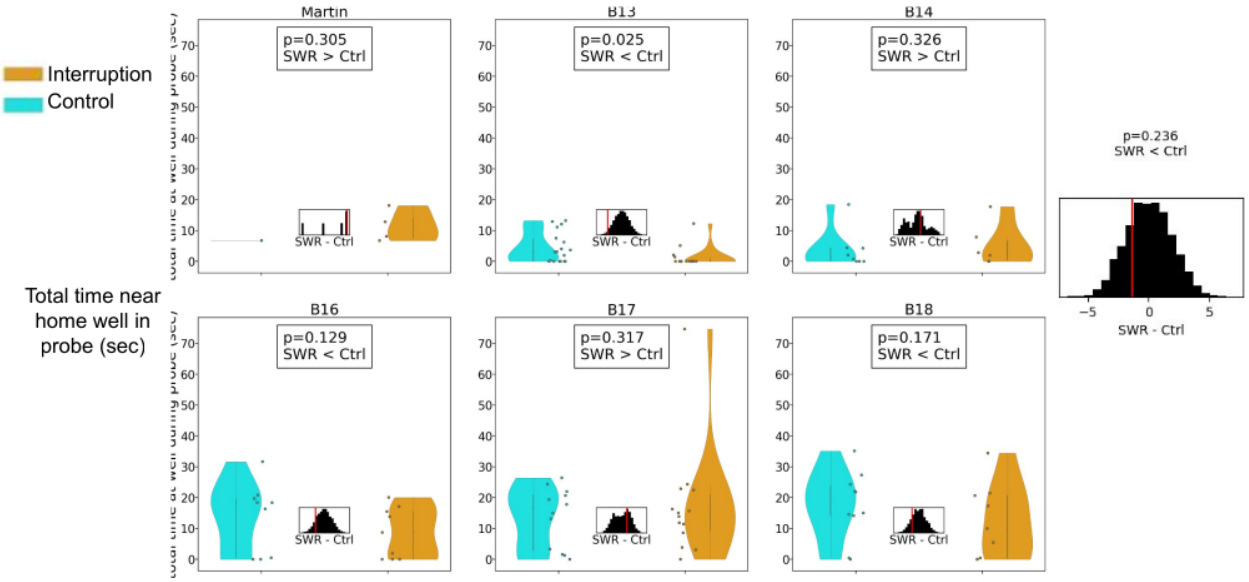
Interruption
Control
Total time near home well in probe (sec)
Martin
p=0.305
SWR > Ctrl
B13
p=0.025
SWR < Ctrl
B14
p=0.326
SWR > Ctrl
p=0.236
SWR < Ctrl
B16
p=0.129
SWR < Ctrl
B17
p=0.317
SWR > Ctrl
B18
p=0.171
SWR < Ctrl
SWR - Ctrl

M

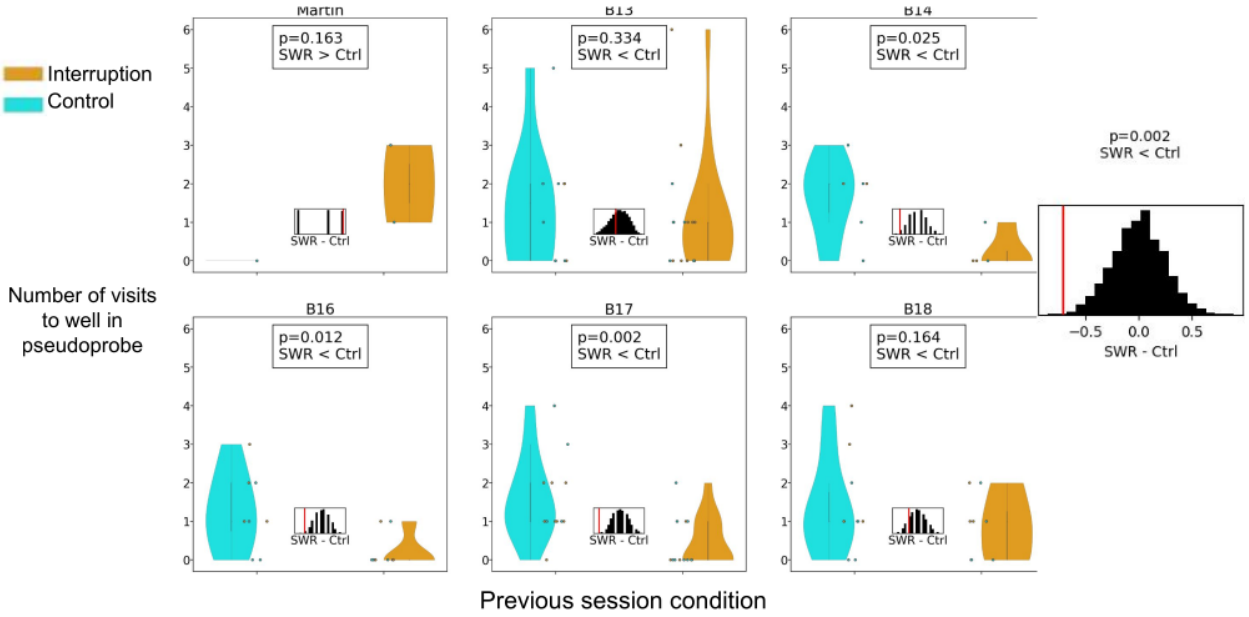

N

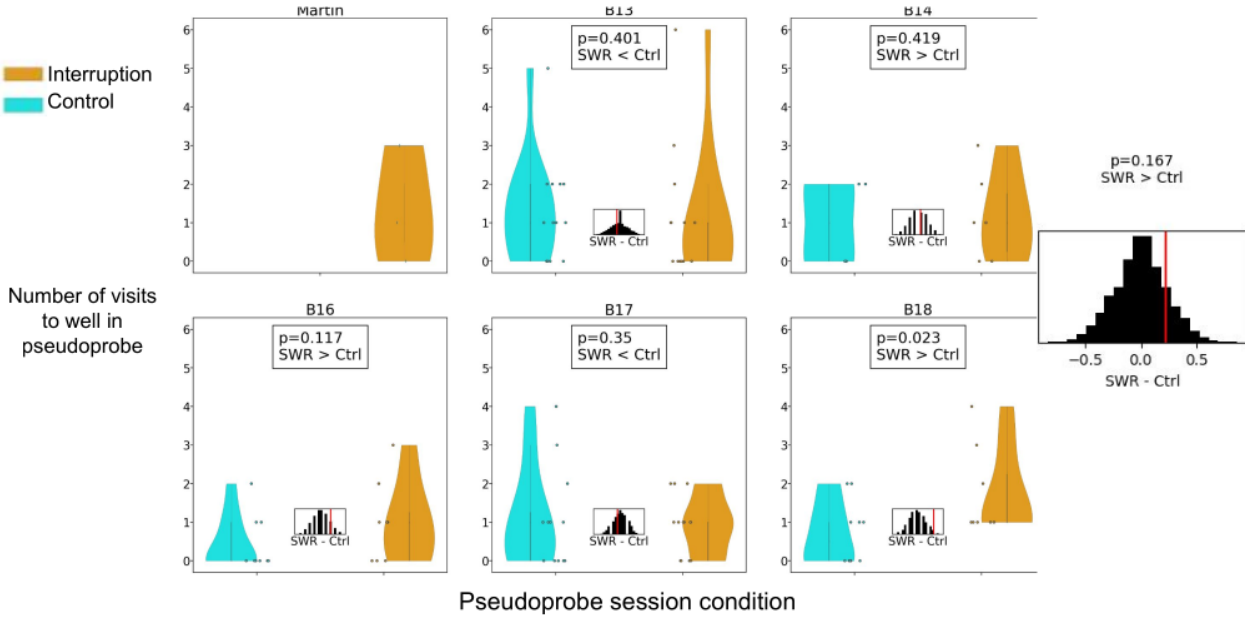

A-H: Fig 9A-H with only delay-stimulation control sessions included
I,J: Fig 10A,B
K,L: Fig 10D,E
M,N: Fig 11A,B

Appendix D: Behavioral Quantification Methods

Each of the measures listed below was computed as described over every combination of parameters listed. Each measure is labeled as a session, trial, or location measure according to how it was computed.

Session measures, such as total task duration, yielded one number per session. Values were compared across session conditions (interruption sessions vs. control sessions). Correlation was also tested for within each condition and across all sessions between this measure and the ripple rate during the task, ripple rate during the probe, stimulation rate during the task, total number of ripples during the task, total number of ripples during the probe, and total number of stimulations during the task.

Trial measures yielded one number per trial. Values were compared across session condition, within home vs away trials across session conditions, and the difference between home and away trials was compared across conditions.

Location measures were computed for every location in the environment, at a resolution of 2.3”2.3”, and then the value at each well was sampled from this map after convolving with a cone-shaped filter to apply smoothing. Values were found at the home well of each session and compared to values at away wells from that session, symmetrically located wells from that session, the same well at the following session, the same well in all later sessions, and the same well in all other sessions. Comparisons were performed across session conditions comparing home well and control location values directly, difference between home and control locations, as well as correlating values across and within conditions with the ripple rate during the task, ripple rate during the probe, stimulation rate during the task, total number of ripples during the task, total number of ripples during the probe, and total number of stimulations during the task.

Where behavior period is listed as a parameter with no values specified, the following behavior periods were considered:

1. Full task
2. Task only during exploration
3. Task only during movement
4. Task only during movement and exploration
5. Task only during home trials
6. Task only during away trials
7. Task only during trials 2-7
8. Task only during home trials 2-4
9. Task only during away trials 2-3
10. Task only during trials 10-15

11. Task only during home trials 6-8
12. Task only during away trials 6-8
13. Task during the first home trial
14. Task during the second home trial
15. Task during the second away trial
16. Periods 5-15 above with the first and last three seconds of each trial removed
17. Full probe
18. The first minute of the probe
19. The first two minutes of the probe
20. The probe until the home well was filled, or the full probe for probes in which the well was not filled
21. Probe only during movement
22. Probe only while the rat was away from the wall
23. Probe only during movement and while the rat was away from the wall
24. 21-23 but restricted to the time before the well was filled, or the full probe for probes in which the well was not filled

Off-wall excursions are defined as time periods in which the rat was at least one foot from the nearest wall.

Exploration bouts are defined as time periods in which the rat moved faster than 8 cm/s, where such time periods separated by less than one second were combined. Bouts defined in this way in which the rat visited fewer than 4 wells were removed from analysis.

Path optimality is defined as the total distance traveled in a given interval divided by the straight-line distance between the start and endpoints of that path.

Gravity is defined as the fraction of entries into a larger radius around a location that the rat visited a given smaller radius around that location. A gravity of 0 thus indicates the rat passed through the larger radius without ever entering the smaller radius, while a gravity of 1 means that every time the rat passed through the larger radius, it also passed through the smaller radius.

Curvature was adapted from Dvorkin et al, 2010. It is defined here as the angle between the entry and exit direction to and from a given location, at a radius of 8cm.

To calculate the spotlight measure, each recorded frame of behavior was analyzed and a value was added to a running total for each location according to the location and heading direction of the animal. As a default, the dot product between the heading

direction and the direction from the rat to the location was added at each time step. In binary mode, this dot product was passed through a binary thresholding function at each time step. A distance weight was specified which weighted the value added at each location exponentially by the distance between the location and the rat. Negative weights led to further distances being weighted more highly, while positive weights led to closer distances being weighted more highly.

Measure	Parameters
Session Measures	
Total task duration	
Number of wells visited per exploration bout	Behavior period: full probe, full task Count repeats to each well: T/F Count off-wall wells only: T/F
Number of wells visited per path away from the wall	Behavior period: full probe, full task Count repeats to each well: T/F
Average exploration bout duration	Behavior period: full probe, full task
Average exploration bout path length	Behavior period: full probe, full task
Average off-wall excursion duration	Behavior period: full probe, full task
Average off-wall excursion path length	Behavior period: full probe, full task
Trial Measures	
Average trial duration	Behavior periods: all trials, trials 2-4, 2-6, 1-2, 1-8
Trial path optimality	Behavior periods: all trials, trials 2-4, 2-6, 1-2, 1-8
Trial path length	Behavior periods: all trials, trials 2-4, 2-6, 1-2, 1-8
Number of wells visited	Behavior periods: all trials, trials 2-4, 2-6, 1-2, 1-8 Count returns to wells: T/F Only considered off-wall wells
Location Measures	
Total time at position	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1, 1.5 ft.
Number of visits to position	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5,

	1, 1.5 ft.
Average dwell time at position	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1, 1.5 ft.
Average curvature at position	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1, 1.5 ft.
Path optimality to position	Behavior period Convolution radius: 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1 ft. Value given to non-visted locations: NaN, 8.5 ft, max of visited values, mean of visited values
Path length to position	Behavior period Convolution radius: 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1 ft. Value given to non-visted locations: NaN, 8.5 ft, max of visited values, mean of visited values
Path length within off-wall excursion	Behavior period Convolution radius: 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1 ft. Value given to non-visted locations: NaN, 8.5 ft, max of visited values, mean of visited values Which excursions to consider: first, last, first where location was visited, last where location was visited, mean value of all excursions, mean value of all excursions in which location was visited
Path optimality within off-wall excursion	Behavior period Convolution radius: 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1 ft. Value given to non-visted locations: NaN, 8.5 ft, max of visited values, mean of

	visited values Which excursions to consider: first, last, first where location was visited, last where location was visited, mean value of all excursions, mean value of all excursions in which location was visited
Latency to location	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1, 1.5 ft. Value given to non-visited locations: NaN, maximum of visited values, mean of visited values, maximum possible duration for the given behavior period, double the maximum possible duration for the given behavior period
Fraction of excursions visited	Behavior period Convolution radius: 0, 0.5, 1 ft. Radius considered at position: 0.25, 0.5, 1, 1.5 ft.
Spotlight measure	Behavior period Radius considered at position: 0.25, 0.5, 1, 1.5 ft. Normalize by time in behavior period: T/F Only sum positive values: T/F Distance weight: -1, -0.5, 0, 0.5, 1
Gravity at location	Behavior period Convolution radius: 0, 0.5 ft Larger pass radius: 0.5, 1, 1.5 ft. Smaller pass radius (as fraction of larger radius): 0.2, 0.4

Appendix E: Statistical comparison definitions

For each effect, a value V is determined based on the datapoints Y, and is compared to the values computed after shuffling the categorical labels, C. Some effects are also defined with respect to another effect, W, or a categorical label, l.

Main effect:

Difference of expected values, separated into sets for which the categorical label does or does not equal the given label l.

$$V_l = E[Y_i \ni C_i = l] - E[Y_i \ni C_i \neq l]$$

Within effect:

The effect W, restricted to data for which the categorical label equals l.

$$V_{W,l} = W(C = l)$$

Across effect:

The summed effect W calculated within each categorical label l. The sum is weighted by the number of datapoints with each label.

$$V_W = \frac{1}{|C|} \sum_{l \in C} count(C = l) \cdot W(C = l)$$

Interaction effect:

The effect W within a categorical label l, minus the the effect excluding datapoints in l.

$$V_{W,l} = W(C = l) - W(C \neq l)$$