

Neurofeedback training can modulate task-relevant memory replay in rats

Anna K. Gillespie , Daniela A. Astudillo Maya, Eric L. Denovellis, Sachi Desse, Loren M. Frank

Departments of Biological Structure and Lab Medicine & Pathology, University of Washington, Seattle, WA, USA • Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA • Howard Hughes Medical Institute, Chevy Chase, MD, USA

 https://en.wikipedia.org/wiki/Open_access

 <https://creativecommons.org/licenses/by/4.0/>

Reviewed Preprint

Published from the original preprint after peer review and assessment by eLife.

[About eLife's process](#)

Reviewed preprint posted

September 12, 2023 (this version)

Sent for peer review

July 19, 2023

Posted to bioRxiv

July 15, 2023

Abstract

Summary

Hippocampal replay – the time-compressed, sequential reactivation of ensembles of neurons related to past experience – is a key neural mechanism of memory consolidation. Replay typically coincides with a characteristic pattern of local field potential activity, the sharp-wave ripple (SWR). Reduced SWR rates are associated with cognitive impairment in multiple models of neurodegenerative disease, suggesting that a clinically viable intervention to promote SWRs and replay would prove beneficial. We therefore developed a neurofeedback paradigm for rat subjects in which SWR detection triggered rapid positive feedback in the context of a memory-dependent task. This training protocol increased the prevalence of task-relevant replay during the targeted neurofeedback period by changing the temporal dynamics of SWR occurrence. This increase was also associated with neural and behavioral forms of compensation after the targeted period. These findings reveal short-timescale regulation of SWR generation and demonstrate that neurofeedback is an effective strategy for modulating hippocampal replay.

eLife assessment

This is an **important** study that tests the effects of using neurofeedback, in the form of reward delivery when large sharp wave-ripples (SWRs) are detected, on neurophysiological and behavioral measures. The authors provide **convincing** evidence that the rate of SWRs increased prior to reward delivery and decreased in the period after reward delivery, with no significant effect on memory performance. The ability to manipulate SWR rate in a naturalistic way is a potentially exciting new tool for studies that seek to understand the function of SWRs.

Introduction

Across species, memory is central to adaptive behavior. Key memory processes are coordinated by the hippocampus, where highly plastic neural ensembles are thought to rapidly encode ongoing experience. Subsequently, patterns of neural activity associated with past experience are

“replayed” during offline states, such as sleep and pauses in ongoing behavior (Findlay et al., 2020). Replay events tend to coincide with sharp wave ripples (SWRs), characteristic bursts of high-frequency oscillatory activity in the CA1 region of the hippocampus that coordinate spiking activity in a wide range of cortical and subcortical brain regions (Buzsaki, 2015; Joo and Frank, 2018; Skelin et al., 2019; Todorova and Zugaro, 2020). The brain-wide coordination of activity during SWRs is thought to enable their proposed role in storing and maintaining memory traces in distributed hippocampal-cortical circuits (Frankland and Bontempi, 2005; Kaefer et al., 2022; Klinzing et al., 2019).

Disrupting SWRs impairs learning and alters memory performance, establishing the necessity of SWRs for intact memory-guided behavior (Aleman-Zapata et al., 2021; Ego-Stengel and Wilson, 2010; Girardeau et al., 2009; Jadhav et al., 2012; Nokia et al., 2012; Todorova and Zugaro, 2020). The importance of SWRs to memory function is further underscored by observations of reductions and abnormalities in SWRs concurrent with memory impairment in several rodent models of neurological disease (Gillespie et al., 2016; Iaccarino et al., 2016; Jones et al., 2019; Prince et al., 2021; Sanchez-Aguilera and Quintanilla, 2021; Stojilkovic et al., 2018). These findings raise the compelling possibility that driving or enhancing SWRs in cases where they are compromised could confer therapeutic benefit.

Indeed, a small number of previous studies have demonstrated positive manipulations of SWRs. Electrical stimulation and optogenetics have been used to evoke exogenous SWRs (Ishikawa et al., 2014; Jiang et al., 2018; Stark et al., 2015; Stark et al., 2014) or to enhance existing SWRs (Fernandez-Ruiz et al., 2019; Maingret et al., 2016). While compelling, some limitations remain. In some cases, studies have demonstrated the selective engagement of subsets of hippocampal neurons (Fernandez-Ruiz et al., 2019; Oliva et al., 2020), but whether these manipulations preserve the full diversity or accuracy of naturally occurring replay events remains unknown. Further, endogenous SWRs tend to occur within permissive windows influenced by cortical activity states (Battaglia et al., 2004; Isomura et al., 2006) and neuromodulatory tone (Novitskaya et al., 2016; Vandecasteele et al., 2014); it is unclear whether exogenous events evoked outside of such states recruit downstream targets or engage brain-wide coordination as endogenous SWRs do. Finally, both electrical and optogenetic interventions are suboptimal for translation to human use. Both approaches would require invasive procedures, and clinical use of optogenetics in the brain faces even more fundamental barriers (White et al., 2020).

These limitations motivated us to take an alternative approach. Neurofeedback, wherein subjects are given real-time external feedback that allows them to alter their own brain states, has been shown to enable modulation of patterns of brain activity across species (Sitaram et al., 2017). The efficacy of neurofeedback-based therapy has been established for some clinical indications, such as ADHD and stroke, but the use of neurofeedback for memory-related symptoms has shown variable efficacy (Staufenbiel et al., 2014; Wang and Hsieh, 2013) and is still in early stages (Jiang et al., 2022; Klink et al., 2021; Trambaiolli et al., 2021). Moreover, most neurofeedback paradigms target the modulation of oscillatory activity in specific frequency bands linked to cognitive functions, such as gamma or theta (Engelhard et al., 2013; Reiner et al., 2014; Wang and Hsieh, 2013) or overall activity levels (Klink et al., 2021), rather than focusing on specific memory-related hippocampal activity patterns (e.g. SWRs).

We therefore asked whether providing external positive neurofeedback coupled to SWRs would effectively drive SWR occurrence, a possibility consistent with prior work using intracranial stimulation of a reward-related brain region to achieve operant conditioning of SWRs (Ishikawa et al., 2014). Specifically, we set out to determine whether subjects could learn to use noninvasive positive neurofeedback to achieve a SWR-conducive state and increase SWR generation. Further, since the replay during SWRs is thought to be central to their function, we also sought to evaluate whether SWR-based neurofeedback could drive replay events with physiologically relevant spatial content.

We were able to develop a neurofeedback training paradigm for rats that increases the rates of SWRs and replay during the targeted neurofeedback period. Importantly, SWRs during the targeted period were reliably associated with replay events which represented a range of spatial content consistent with that seen in our previous study of replay in a very similar task (Gillespie et al., 2021 [↗](#)). Interestingly, we also observed some compensatory changes outside of the targeted neurofeedback period. Our results demonstrate the efficacy of a neurofeedback strategy to modulate hippocampal replay in young, healthy animals. Further, they suggest that SWR generation during task performance is influenced by homeostatic regulation on the timescale of seconds, an important factor to consider when developing strategies to drive SWRs and replay therapeutically.

Results

A neurofeedback paradigm to promote SWRs

To reinforce SWR activity, our neurofeedback paradigm linked the online detection of large SWRs with positive reinforcement: delivery of a tone and food reward (**Figure 1A** [↗](#)). We embedded the neurofeedback requirement into a flexible spatial memory task (**Figure 1B** [↗](#)) that we have studied extensively (Gillespie et al., 2021 [↗](#)), allowing us to compare neural and behavioral data from subjects that perform the task with neurofeedback manipulation to those without the manipulation component. All subjects performed the same overall task, which comprised structured trials consisting of three steps (**Figure 1C** [↗](#)): first, initiating a trial by nosepoking at the home port; second, maintaining a nosepoke at whichever center port illuminates (randomly assigned on each trial) for a duration determined by the port's criteria (see below); and finally, choosing one of eight arms to visit before returning to the home port to initiate the next trial.

Here we present data from two cohorts of animals. In the manipulation cohort, one center port served as the neurofeedback port and the other served as the delay port. At the neurofeedback port, the subject had to maintain its nose in the port until a suprathreshold SWR was detected, which triggered delivery of a tone and food reward (**Figure 1A** [↗](#), Supplementary Figure 1A,B; and see Methods). The pre-reward period at the center port on neurofeedback trials is thus considered the “targeted interval” for neurofeedback. At the delay port, the subject was required to maintain its nose in the port for an unpredictable amount of time. The delay period length was randomly chosen from the pre-reward duration of recent neurofeedback trials, such that the amount of time spent waiting at each of the center ports was approximately matched (Supplementary Figure 1C, and see Methods). Delay trials thus serve as an internal control for each subject, allowing us to assess the specificity of any neurofeedback effect. In the control cohort, whose data have been presented previously (Gillespie et al., 2021 [↗](#)), both center ports required the subject to maintain its nose in the port for an unpredictable delay period (see Methods), after which a tone and food reward was delivered.

The memory aspect of this task is required at the third step of each trial: at any given time, one of the eight arms provides a large food reward; the others provide nothing. Over multiple trials, the subject must sample arms to discover which one provides reward and remember to return to that location on subsequent trials. Occasionally, the rewarded arm changes, prompting the subject to return to sampling different arms to find the new goal location (Gillespie et al., 2021 [↗](#)). Before implant, all subjects were pretrained on a delay-only version of this task, as summarized in the experimental timelines in **Figure 1D** [↗](#). Once performance stabilized, a microdrive containing 30 independently movable tetrodes was implanted targeted to dorsal CA1 (Supplementary Figure 1D). After implant, subjects were given several days to recover; during this time tetrodes were gradually lowered into the CA1 cell layer and the animal was briefly re-trained to perform the delay-only version of the task. The control cohort of subjects continued such performance for many days (Gillespie et al., 2021 [↗](#)).




Figure 1.

A neurofeedback paradigm to promote SWRs.

(A) Schematic of the neurofeedback (NF) protocol: while subject's nose remains in the neurofeedback port, SWRs (yellow) are detected in real time. Top trace: raw CA1 LFP; bottom trace: ripple filtered (150-250 Hz) CA1 LFP. During the neurofeedback interval, the first event exceeding a set standard deviation (sd) threshold (red line) triggers the delivery of a sound cue and food reward to the rat.

(B) Top-down view of the maze environment. Reward ports are indicated by rectangles.

(C) The rules of the spatial memory task for each of the behavioral cohorts.

(D) Experimental timeline for each behavioral cohort.

For subjects in the manipulation cohort, the neurofeedback component of the task was then introduced. Each day, 4-6 CA1 tetrodes were used for online SWR detection (Supplementary Figure 1A, and see Methods). Because we use just a small subset of tetrodes, our online detection algorithm does not detect every SWR event identified by our offline detection strategy, which incorporates all CA1 tetrode data. This strategy provides two important benefits. First, it means that rewards are provided to SWR events on an unpredictable, variable schedule – an operant conditioning regime which drives robust, persistent task engagement (Ferster and Skinner, 1957 [↗](#)). Second, we choose a different subset of tetrodes each day (see Methods) to ensure that subjects cannot learn to trigger the algorithm and earn reward by repeatedly producing a single firing pattern corresponding to specific units on those tetrodes. Initially, the threshold for online SWR detection is set to a low value, requiring only small SWRs to trigger reward. SWR size, as defined by the peak amplitude of the event, can be generally described by a long-tailed distribution, with larger events occurring less frequently than smaller events (Yu et al., 2017 [↗](#)). Starting with a low threshold ensures that the animals received feedback relatively rapidly and remained engaged early in training.

Over subsequent days of training, the size threshold for SWR detection at the neurofeedback port was gradually increased, requiring the subject to generate larger and larger SWR events on neurofeedback trials in order to proceed with the task (Supplementary Figure 1E). We then maintained the threshold at a high level for several additional days of behavior (“stable period”).

Neurofeedback training increases SWR rate during the targeted interval

In response to the increasing SWR detection threshold, subjects in the manipulation cohort successfully produced increasingly large trigger events at the neurofeedback port over the course of training (Supplementary Figure 2A). This result confirms that our feedback criteria were enforced as designed and that subjects were capable of meeting the criteria.

As the detection threshold was raised, subjects could have met the neurofeedback criteria simply by waiting much longer for increasingly rare suprathreshold events. This was not the case: instead, subjects produced large SWR events much faster than would have been predicted based on the initial prevalence of similarly sized events (Supplementary Figure 2B). This suggests that subjects learned a strategy for generating large events that allowed them to meet the neurofeedback criteria more efficiently.

We reasoned that two candidate strategies could have enabled subjects to limit their wait times: learning to increase the average size of SWR events or learning to increase the rate of SWR event generation across all sizes. To evaluate these possibilities, we quantified differences both between the two cohorts and between trial types of the neurofeedback cohort. When evaluating the significance of differences between the neurofeedback and delay trial types *within* each manipulation cohort subject, we used ranksum tests, since this approach allows us to capitalize on the large numbers of measurements collected from each rat without assuming that they were normally distributed. To quantify systematic group effects of cohort, we used linear mixed-effects models (see Methods), since this approach allows us to appropriately account for correlations between measurements taken from each subject (Yu et al., 2022 [↗](#)).

First, to assess whether subjects learned to directly increase the size of SWR events, we quantified the mean size of SWRs that occurred pre-reward at the center ports on neurofeedback and delay trials during behavioral epochs from the stable period. Importantly, to account for the structure of the task, we excluded suprathreshold events detected by our online algorithm during both neurofeedback and delay trials (see Methods). These analyses revealed that subjects did not consistently increase SWR size: only two out of four subjects showed a very slight but significant increases (Supplementary Figure 2C). We also found no difference in SWR size between either trial

type in the manipulation cohort and the control cohort (Supplementary Figure 2C, inset). Instead, when we quantified the number of events of each size per behavioral epoch, we found that within the manipulation cohort, neurofeedback trials contained roughly twice as many events relative to delay trials for all SWR sizes (Supplementary Figure 2D). These results rule out our first candidate strategy and indicate that subjects met the neurofeedback criteria efficiently by modulating SWR rate, not event size.

The increase in SWR rate pre-reward was associated with striking differences in when SWRs occurred both within and across cohorts that were evident even in individual trials (**Figure 2A**). To assess how SWR rate was changing during the time spent at the center ports, we calculated SWR rate in 0.5 s bins throughout the time subjects spent nose-poked at the center ports, aligning each trial to reward delivery (**Figure 2B**). We excluded the trigger SWR of neurofeedback trials from analysis, since it is guaranteed to occur just prior to reward delivery in each neurofeedback trial and causes a large spike in SWR rate. Subjects in the manipulation cohort all showed high, stable SWR rates before reward delivery, with neurofeedback trials showing higher SWR rate than delay trials. By contrast, subjects in the control cohort rarely generated SWRs during the pre-reward period but showed much higher SWR rates than the manipulation cohort after reward delivery.

To quantify these effects, we calculated the mean SWR rate over each portion of the time at the center ports (**Figure 2C**): the pre-reward period (left), the post-reward period (middle), and the total port period (pre+post; right). Indeed, we found that during the pre-reward period, SWR rates were approximately twice as high on neurofeedback compared to delay trials, and both trial types from the manipulation cohort had much higher rates than trials from the control cohort (**Figure 2C**, left). This effect was not dependent on the trigger SWR, as excluding suprathreshold events from both neurofeedback and delay trials did not change the result (Supplementary Figure 2E). Moreover, although we did observe a slightly lower mean speed during the pre-reward period on neurofeedback trials compared to delay trials and trials from the control cohort (Supplementary Figure 2F), movement differences could not explain the difference in SWR rates (Supplementary Figure 2G, H). This finding indicates that the neurofeedback training caused far more SWRs to occur prior to reward delivery at both center ports in the manipulation cohort, with the most pronounced effect occurring on neurofeedback trials.

These patterns differed after the reward at the center ports, where SWR rate was slightly but significantly higher during delay trials compared to neurofeedback trials in the manipulation cohort subjects but was overall indistinguishable between the cohorts (**Figure 2C**, middle). Note that the slightly lower SWR rate post-reward on neurofeedback trials compared to delay trials could not be explained by a refractory period following the trigger SWR (Supplementary Figure 2I). When considering both pre- and post-reward periods together, the neurofeedback trials still showed higher SWR rates than delay trials within the manipulation cohort, but the differences between each trial type and the control cohort were not significant (**Figure 2C**, right). This surprising result suggested that the pre-reward effects on SWR rate in the manipulation cohort were followed by some compensatory effects post-reward that resulted in a net normalization of total SWR rate across cohorts.

Quantifying SWR count, rather than rate, provided further explanation for these surprising findings. We quantified the number of SWRs during pre-reward, post-reward, and total center port times (**Figure 2D**). Our SWR count results were completely consistent with the SWR rate findings during the pre-reward period (**Figure 2D**, left) and largely consistent when considering the total center port time (**Figure 2D**, right). However, we found striking differences during the post-reward period, during which we saw a clear reversal of the pre-reward effect on SWR counts: neurofeedback trials contained fewer SWRs than delay trials, and both trial types contained far




Figure 2.

Neurofeedback training enhances SWR rate during targeted interval.

(A) Example CA1 raw LFP traces and ripple filtered LFP traces (150-250Hz) with SWRs highlighted in yellow, from the time spent at a center port on a neurofeedback (NF) trial (top) and a delay trial (middle) from a manipulation subject and a trial from a control subject (bottom).

(B) SWR rate calculated in 0.5 s bins during the pre- and post-reward periods, aligned to the time of reward delivery (dashed lines) for subjects in the manipulation (top row) and control (bottom row) cohorts. Trigger SWRs on NF trials are excluded from the rate calculation, and time bins with fewer than 100 trials contributing data are not shown. Vertical bars indicate S.E.M.

(C) SWR rate calculated during the pre-reward period (left), post-reward period (middle) and for the total time at the center port (right). Manipulation cohort n = 1892, 684, 1157, and 1602 NF trials and 2022, 640, 1201, and 1552 delay trials; control cohort n = 2490, 2629, 2027, and 3021 trials. For the pre-reward period, manipulation cohort ranksum comparisons between NF and delay trials: $p = 4.382 \times 10^{-258}$, 7.111×10^{-83} , 5.689×10^{-214} , and 3.285×10^{-191} . Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 1.126 \times 10^{-16}$; manipulation cohort delay trials vs control cohort trials: $p = 0.009$. For the post-reward period, manipulation cohort ranksum comparisons between NF and delay trials: $p = 3.646 \times 10^{-127}$, 0.038 , 6.538×10^{-11} , and 2.768×10^{-23} , respectively. Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 0.142$; manipulation cohort delay trials vs control cohort trials: $p = 0.691$. For pre+post combined, manipulation cohort ranksum comparisons between NF and delay trials: $p = 3.324 \times 10^{-64}$, 3.136×10^{-40} , 2.996×10^{-69} , and 9.066×10^{-50} , respectively. Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 0.237$; manipulation cohort delay trials vs control cohort trials: $p = 0.504$.

(D) Count of SWR events detected during the pre-reward period (left), post-reward period (middle) and for the total time at the center ports (right). Trial n are the same as in (C). For the pre-reward period, manipulation cohort ranksum comparisons between NF and delay trials: $p = 3.133 \times 10^{-96}$, 8.936×10^{-37} , 1.266×10^{-106} , and 6.024×10^{-65} . Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 6.790 \times 10^{-7}$; manipulation cohort delay trials vs control cohort trials: $p = 0.0018$. For the post-reward period, manipulation cohort ranksum comparisons between NF and delay trials: $p = 1.797 \times 10^{-137}$, 6.100×10^{-10} , 6.688×10^{-13} , and 2.152×10^{-130} . Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 2.361 \times 10^{-12}$; manipulation cohort delay trials vs control cohort trials: $p = 5.820 \times 10^{-9}$. For pre+post combined, manipulation cohort ranksum comparisons between NF and delay trials: $p = 1.534 \times 10^{-19}$, 5.828×10^{-14} , 6.735×10^{-32} , and 0.083 . Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 0.366$; manipulation cohort delay trials vs control cohort trials: $p = 0.299$. (E), Dwell time post-reward. Trial n are the same as in (C). Manipulation cohort ranksum comparisons between NF and delay trials: $p = 1.397 \times 10^{-23}$, 8.588×10^{-49} , 6.180×10^{-4} , and 7.127×10^{-257} . Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 4.075 \times 10^{-18}$; manipulation cohort delay trials vs control cohort trials: $p = 1.835 \times 10^{-18}$.

For C-E, all within-subject ranksum p-values are corrected using the Benjamini-Hochberg method and all groupwise comparisons are performed using linear mixed effects models (see Methods).

fewer SWRs post-reward than trials from the control cohort (**Figure 2D**, center). Since post-reward SWR rates were not different between the cohorts, the count results indicated that the compensatory effect must involve the amount of time spent at the port post-reward.

Indeed, we found clear differences in the amount of time that subjects chose to spend at the port after reward delivery. Here it is important to remember that in contrast to the pre-reward dwell time, which is imposed by task constraints, the post-reward dwell time is not externally constrained. It was notable, therefore, that neurofeedback trials had slightly but significantly shorter post-reward periods than delay trials (**Figure 2E**). Even more strikingly, subjects in the manipulation cohort remained at the port for approximately half as long as subjects in the control cohort (**Figure 2E**, inset), despite all trial types receiving the same amount of reward. The difference in dwell time could not be explained by the control cohort having fewer days of task experience than the manipulation cohort, since post-reward dwell time tended to increase, rather than decrease, over subsequent days of performance within the control cohort (Supplementary Figure 2J).

Together, these findings suggest that both covert and overt processes regulate SWR generation. Animals learned to modulate their SWR rate when required during the pre-reward period of neurofeedback trials and did so covertly, without displaying obvious behavioral correlates. A subsequent overt change in behavior – leaving the well early – was then associated with lower numbers of SWRs post-reward in the neurofeedback cohort.

The surprising stability of total SWR rate and count at the center ports across the two cohorts, despite changes in when those SWRs occurred relative to reward delivery, suggests that some regulatory mechanisms keep SWR generation within a set range for the center port phase of this task. We therefore asked whether we could detect evidence for this regulation on a trial-by-trial basis. Strict regulation within a single trial would predict that, for instance, trials with many SWRs prior to reward would likely have fewer SWRs post-reward. However, we did not find evidence for such a relationship consistently: the correlation coefficients between pre-reward and post-reward SWR counts were positive in some subjects and negative in others, and always very small in magnitude, indicating a weak relationship ($r = -0.084, 0.130, 0.118$, and -0.061 and $p = 1.47 \times 10^{-7}, 3.84 \times 10^{-6}, 8.60 \times 10^{-9}$, and 6.01×10^{-4} , for each subject in the manipulation cohort, respectively). Similarly, we did not find consistent correlations between pre- and post-reward SWR rate nor dwell time. These results suggest that while compensatory regulation of SWR generation may be evident across hundreds of trials, high variability in SWR rate, count, and dwell time across trials may prevent consistent detection of this effect at the level of single trials.

Neurofeedback training preserves replay content during SWRs

SWRs are thought to support memory storage as a result of the specific patterns of spiking activity (replay) seen during these events (Gridchyn et al., 2020). Thus, we wanted to assess whether our SWR manipulation also increased replay during the targeted pre-reward period. As hippocampal CA1 neurons often fire reliably at particular locations in space (“place cells”) while an animal moves through an environment, replay content has most often been studied with regard to its spatial content. Specifically, spatial content can be assessed by using a model that relates neural firing to specific locations in space. Inverting this “encoding model” yields a “decoding model” that can predict the spatial locations represented by the neural spiking within SWRs (Zhang et al., 1998).

We used a clusterless state-space model (Chen et al., 2012; Deng et al., 2016; Denovellis et al., 2021; Kloosterman et al., 2014) identical to the one we used previously to characterize the spatial content of replay for the animals in the control cohort (Gillespie et al., 2021). Briefly, we first project the 2D maze environment to 1D (**Figure 3A**). We next use times when the animal is moving through the maze to build a model that relates amplitude features of detected spikes to the

location on the 1D maze where they were observed. We then can predict the spatial representation of spikes observed during movement (**Figure 3B** [↗](#)) and during SWRs (see Methods).

As expected, we observed a wide range of spatial representations during SWRs (**Figure 3C** [↗](#)). We define replay events as SWRs which predominantly contain spatial representations that are consistent with how a subject can experience the environment: those which represent a single location or a continuous trajectory through space (see Methods). These events frequently represent the subject's current location ("local" replay; **Figure 3C** [↗](#), left two events) as well as locations or trajectories that encompass distant regions of the maze environment ("remote" replay; **Figure 3C** [↗](#), middle four events). In small subset of cases, SWRs contain a mixture of disjoint spatial representations with no single location clearly dominant (fragmented events; **Figure 3C** [↗](#), right), which do not meet our replay criteria and are excluded from analysis.

Across a variety of measures, we found that neurofeedback training preserved replay content. In SWRs during the pre-reward period, there was no difference in the fraction of SWRs that qualified as replay between the two cohorts, and within each manipulation cohort subject, there was no consistent difference between neurofeedback and delay trials (**Figure 3D** [↗](#)).

Replay of remote locations has been proposed to play a particularly important role in updating or maintaining memory in the absence of direct experience ([Gillespie et al., 2021](#) [↗](#); [Gupta et al., 2010](#) [↗](#); [Olafsdottir et al., 2015](#) [↗](#)). Consistent with our SWR findings, we find much higher rates of remote replay during the pre-reward period in neurofeedback trials compared to delay trials within the manipulation cohort, and low rates of remote replay during this time in trials from the control cohort (**Figure 3E** [↗](#)). However, if we consider both the pre- and post-reward periods combined, we find no difference in remote replay rate between the two cohorts for either trial type (**Figure 3F** [↗](#)). Within the manipulation cohort, neurofeedback trials have slightly but significantly higher rates of remote replay than delay trials over the combined pre- and post-reward period.

Prior work demonstrated that replay events in this task were enriched for specific, relevant past experiences that change as a function of goal location ([Gillespie et al., 2021](#) [↗](#)), and we saw the same biases in replay content in the manipulation cohort. We classified remote replay events based on the behavioral significance of the maze arm they represented and used a generalized linear model to assess how behavioral categories influence the likelihood of replay of a given arm (see Methods). As previously seen in the control cohort, we observed consistently enhanced representation of previous goal locations in replay events during the pre-reward period in both neurofeedback and delay trials within the manipulation cohort (**Figure 3G** [↗](#)). Similarly, we replicated the reduced prevalence of replay representing the immediately previously visited arm and a slightly increased prevalence of replay of the arm associated with the upcoming choice in both trial types of the manipulation cohort. These results also remain unchanged when we analyzed replay events during both pre- and post-reward time at the center ports (Supplementary Figure 3A).

Together, these results indicate that while the timing of replay is altered by the neurofeedback training, the content and behavioral relevance is preserved.

Neurofeedback does not alter behavioral performance in a trial-by-trial manner

Finally, we wanted to assess whether neurofeedback training would have an impact on the performance of the spatial memory task. Since SWRs and remote replay rate were generally indistinguishable between the manipulation and control cohorts when the total time at the center ports was considered, we did not expect to find a large difference in behavior between the two




Figure 3.

Neurofeedback preserves replay content.

(A) The 2D maze is linearized to 1D for decoding efficiency; movement trajectory for an example neurofeedback (NF) trial is shown in green. Times of SWR events are highlighted in yellow, and the small letters indicate the SWRs that are shown in (C).
 (B) Decoding during movement times shows a close correspondence between decoded position and the subject's real position.
 (C) Several examples of decoding during SWRs, drawn from the example trial above. These include local events (two left), remote events (four middle), and one uninterpretable event (right).

cohorts. Within the manipulation cohort, however, neurofeedback trials generally showed higher SWRs and remote replay rates, providing an opportunity to measure any behavioral consequences of the neurofeedback training. Our previous findings (Gillespie et al., 2021 [DOI](#)), suggested that SWRs and replay do not drive trial-by-trial choice behavior in this task, so we did not predict that the increase in SWRs on individual neurofeedback trials would necessarily improve performance on those trials. On the other hand, we speculated that the requirement to generate SWRs could serve as a distracting or cognitively demanding additional aspect of neurofeedback trials that could impair subsequent arm choice on those trials, and if so, this would be an important factor to consider when designing future implementations of the neurofeedback paradigm.

We used several metrics to quantify task performance. First, during trials when the subject had not yet discovered a new goal arm (“search” trials), optimal behavior dictates that the animal should sample each arm without repeating visits to unrewarded arms until the goal arm is found. To measure performance on these trials, we quantified search efficiency: the fraction of search trials when the subject visits an arm not yet sampled. The manipulation cohort subjects show a slight improvement in search efficiency compared to control subjects, which only reached significance on delay trials (Figure 4A [DOI](#)). However, rather than relating to the neurofeedback training, this difference could have been driven by the additional days of experience that the manipulation cohort has on the task, since the control cohort showed a subtle but consistent increase in this performance measure over time (Supplementary Figure 4A). Consistent with this possibility, we find no significant difference in this metric between neurofeedback and delay trials. We next asked whether “redundant” search trials – when the subject chooses to sample an arm it has already visited – tended to be neurofeedback or delay trials. We found no bias toward either of the two trial types: neurofeedback trials accounted for approximately half of redundant search trials for each subject (Figure 4B [DOI](#)), indicating that neither the neurofeedback nor delay trial type was consistently overrepresented among the trials with suboptimal behavior.

To quantify performance during trials when the goal arm had already been discovered by the subject and should be chosen again (“repeat” trials), we measured the fraction of subsequent trials during which the subject correctly visited the goal arm. We also saw no difference between the two cohorts or between neurofeedback and delay trials on this measure (Figure 4C [DOI](#)). Within the manipulation cohort, we also found that error trials, when subjects do not choose to visit a known goal arm, were equally distributed between neurofeedback and delay trials (Figure 4D [DOI](#)). Overall, these results suggest that on a single-trial timescale, the neurofeedback manipulation had neither a beneficial nor detrimental effect on behavior, in line with our previous findings (Gillespie et al., 2021 [DOI](#)).

Discussion

We demonstrate that neurofeedback can be used to substantially alter the occurrence of SWRs and replay in rat subjects. Compared to a control cohort receiving no manipulation, subjects that experienced several weeks of neurofeedback training developed striking changes in when SWRs tended to occur. These subjects learned to greatly increase the rate of SWRs prior to reward delivery at the center ports, a period when subjects in the control cohort tended to have very few SWRs. This effect was strongest on the trials when SWRs were required (neurofeedback trials), and weaker but still evident on delay trials, when they were not. Importantly, we found that the neurofeedback training preserved behaviorally relevant, interpretable spatial replay during SWRs, and resulted in an increased rate of remote replay on neurofeedback trials compared to delay trials. These results indicate that rapid neurofeedback is an effective approach to modulate physiologically relevant, memory-related patterns of activity in the hippocampus.




Figure 4.

Neurofeedback does not alter behavioral performance.

(A) During trials when the subject is searching for a new goal location, we quantify the fraction of trials in which the subject chooses an arm that has not yet been sampled. Manipulation cohort ranksum comparisons between neurofeedback (NF) and delay trials: $p = 0.9562, 0.7509, 0.7509, \text{ and } 0.9562$, respectively. Inset: Groupwise comparisons. Manipulation cohort NF trials vs control cohort trials: $p = 0.5425$; manipulation cohort delay trials vs control cohort trials: $p = 0.0069$.

(B) The fraction of redundant search trials per epoch which are NF (vs delay). Two-sided sign test vs 0.5: $p = 0.1102, 0.2668, 1, \text{ and } 1$, respectively.

(C) During trials when the subject has discovered the goal arm, the fraction of subsequent trials in which the goal arm is visited ("correct" choice).

Manipulation cohort ranksum comparisons between NF and delay trials: $p = 0.7202, 0.7202, 0.8123, \text{ and } 0.7202$. Inset: Groupwise comparisons.

Manipulation cohort NF trials vs control cohort trials: $p = 0.6946$; manipulation cohort delay trials vs control cohort trials: $p = 0.2857$.

(D) The fraction of error repeat trials which are NF trials. Two-sided sign test: $p = 0.3770, 0.7744, 1, \text{ and } 0.0227$, respectively.

For all panels, manipulation cohort $n = 37, 14, 18, \text{ and } 26$ behavioral epochs per subject, respectively, and control cohort $n = 24, 32, 23, \text{ and } 33$ behavioral epochs. For (A) and (C), all ranksum p -values are corrected using the Benjamini-Hochberg method and all groupwise comparisons are performed using linear mixed effects models.

These results substantially extend a prior demonstration of operant conditioning of SWRs (Ishikawa et al., 2014 [DOI](#)) in three important ways. First, we used external feedback (food reward) as the reinforcer instead of intracranial stimulation, a change which makes the paradigm more amenable to clinical translation. Second, we administer the training in the context of a spatial memory task, which allows us to show that the neurofeedback training itself does not impose a cognitive burden. Finally, we show that the SWR events observed in the neurofeedback condition contain interpretable, behaviorally relevant spatial replay, suggesting that they are likely to contribute to normal replay function. Our findings also extend existing neurofeedback approaches by modulating the occurrence of discrete memory-related events in the hippocampus, in contrast to longer-timescale modulation of continuous oscillations (Engelhard et al., 2013 [DOI](#); Klink et al., 2021 [DOI](#); Reiner et al., 2014 [DOI](#); Wang and Hsieh, 2013 [DOI](#)).

Additionally, our neurofeedback paradigm extends existing approaches for positive manipulations of SWRs. First, we show that SWRs during the targeted interval contain robust and diverse replay that represents experiences relevant to the current state of the behavioral task, indicating that our method of SWR promotion recruits physiologically relevant ensembles during SWRs, does not overly prescribe which ensembles participate in SWRs, and allows replay content to evolve and adapt to changing task demands. Further, because subjects learn to modulate SWR rate, rather than simply generating a single suprathreshold event on command, it is likely that they learn to engage a SWR-permissive state during the targeted interval in which brain-wide neural activity and neuromodulatory tone also enter a SWR-permissive realm. This increases the likelihood that manipulated events could recruit downstream brain regions in physiologically relevant ways. Gaining experimental control over this SWR-permissive state offers compelling future opportunities to dissect extrahippocampal neural activity patterns and neuromodulatory influences on SWR generation.

Our results also revealed evidence for unexpected conservation of SWR rate and counts over the total time spent at the center ports. Despite the dramatic increase in SWR rate prior to reward delivery, we found that the overall number of SWRs produced was not different between the manipulation and control cohorts when the post-reward period was also taken into consideration. One can imagine that the existence of compensatory, potentially homeostatic, regulation of SWRs could be critical for adaptive behavior: an excess of replay could lead to intrusive, rigid memory or loss of behavioral flexibility, while too little could lead to forgetting or memory interference. Such possibilities are supported by evidence for important synaptic scaling properties of SWRs (Norimoto et al., 2018 [DOI](#)) and a prior finding that electrical disruption of SWRs during sleep caused increased SWR occurrence over longer timescales (Girardeau et al., 2014 [DOI](#)). Our results show that this regulation occurs during waking as well as during sleep, and further suggest that it occurs on short timescales, such as between pre- and post-reward periods at the center port in the current task. Moreover, our findings identify evidence for multiple systems that could interact to regulate SWR occurrence: in addition to learning to modulate SWR rate directly, subjects also learned to use behavioral strategies, such as changing the amount of time spent at a port, that had the effect of normalizing SWR counts.

In the context of a potential therapeutic manipulation, our study provides both a cautionary result and cause for optimism. We did not observe an improvement in behavioral performance due to the neurofeedback manipulation. This is consistent with prior results showing a decoupling between replay and trial-by-trial decisions in this task (Gillespie et al., 2021 [DOI](#)) and others (Carey et al., 2019 [DOI](#)); we predict that longer timescale manipulations of SWR rate (such as moving to blocks of neurofeedback vs delay trials rather than randomly interleaving the two trial types) would affect memory consolidation and would show behavioral impacts. We also observed evidence for compensatory changes in the overall number of SWRs and replay events between the two cohorts prior to the memory-guided decision phase of the trial (outer arm choice). This suggests that in our

young and healthy subjects, various homeostatic mechanisms act to maintain SWRs and the associated replay events at advantageous levels, potentially limiting the extent to which we are able to shift SWR rates beyond normal levels.

At the same time, our study demonstrates that it is possible to use behavioral feedback to modulate the rates of SWRs and associated replay events during a specific time period. In subjects with compromised SWRs, such as rodent models of Alzheimer's disease, even a moderate increase in SWRs or replay could move the system back towards healthy levels and potentially improve memory performance. Since the neurofeedback paradigm depends on the occurrence of at least a low endogenous rate of SWR occurrence, it would be important to implement neurofeedback training as a relatively early interventional strategy prior to extensive neurodegeneration, and training may take longer in aged or impaired subjects. Our results also raise the possibility that reduced SWR rates in disease models could be caused by a failure of homeostatic regulatory mechanisms. If this is the case, then a neurofeedback approach may even more effectively promote SWRs in the disease case, because it will not have to overcome internal regulation. Our current findings therefore serve as useful “proof-of-concept” results that can inform and shape the implementation and optimization of SWR-focused neurofeedback interventions in disease models. Importantly, when combined with minimally invasive techniques for detecting SWRs and replay in humans, such as MEG (Liu et al., 2019 [DOI](#); Nour et al., 2021 [DOI](#)), this technique could be further adapted to patient populations to assess whether SWR-triggered neurofeedback in humans could lead to similar enhancement or restructuring of replay, and ultimately whether such an intervention could be therapeutic.

Acknowledgements

The authors are grateful to Max Ladow and Rhino Nevers for their assistance with data collection, to Jennifer Guidera for her contributions to analysis methods, and to Tom Davidson, Demetris Roumis, P. Dylan Rich, Jon Rueckemann, and all members of the Frank lab for insightful discussions throughout the project. We also thank Alan Kaplan and the Lawrence Livermore National Laboratory High Performance Computing Center for providing computational resources. This work was supported by the Simons Collaboration on the Global Brain (postdoctoral fellowship 500639 to A.K.G. and grant 521921 to L.M.F.), the NIH BRAIN Initiative (K99AG068425 to A.K.G.) and Howard Hughes Medical Institute (L.M.F.).

Author Contributions

Conceptualization: A.K.G. and L.M.F.; Investigation: A.K.G., D.A.A.M., and S.D.; Software: E.L.D.; Formal analysis, visualization, data curation, and writing – original draft: A.K.G.; Writing – review & editing: A.K.G., E.L.D. and L.M.F.; Supervision, project administration, and funding acquisition: A.K.G. and L.M.F.

Declaration of Interests

U.S. Patent Application No. 62/887,875 on SWR modulation via neurofeedback has been filed by the University of California, San Francisco with A.K.G. and L.M.F. listed as inventors.

Figure legends

[currently embedded]

Methods

Resource availability

Lead Contact

Requests for further information, resources, or code should be directed to and will be fulfilled by the lead contact, Anna Gillespie (annagill@uw.edu).

Materials Availability

No new reagents were generated by this study.

Data and Code Availability

All data used for this study are publicly available on the DANDI Archive (link TBD). All code used to analyze this data is publicly available at (link TBD).

Experimental model and subject details

Data from 4 male Long-Evans rats (Charles River RRID:RGD_2308852; aged 4-9 months; 450-550 g) comprise the neurofeedback cohort of animals included in this study. We also include data from a cohort of 4 additional male Long Evans rats used for a previous study who did not receive neurofeedback training (Gillespie et al., 2021 [link](#)). All animals were kept on a 12-hour light-dark cycle (lights on 6am-6pm) with *ad libitum* access to food (standard rat chow) in a temperature- and humidity- controlled facility. Beginning one week prior to and extending throughout behavioral training and data acquisition, rats were singly housed and food restricted to a maximum of 85% of free-feeding weight. All procedures were approved by the Institutional Animal Care and Use Committee at the University of California, San Francisco.

Method details

Behavioral training

Food-restricted rats first complete two 40 min sessions of training on a linear track to receive liquid food reward (evaporated milk + 5% sucrose) from automated, photogated reward ports at each end of the track. Motivated subjects (>50 rewards on final session) advance to learn a complex spatial memory task. This task takes place on a large, walled maze equipped with reward ports, the same as used in Gillespie et al, 2021 [link](#). Pretraining occurs as described in Gillespie et al, 2021 [link](#); briefly, subjects are gradually taught to reliably perform the task, with each trial consisting of an initiation poke at the home port and small (50 uL) reward delivery, then an extended dwell period at whichever of the two central ports illuminates, and finally a visit to one of the arm ports. During pretraining, both center ports require the rat to keep its nose in the port for a delay period, randomly chosen from a range, up to 5-10 s. After the delay is complete, a tone plays and a small reward (50 uL) is delivered. Finally, the subject must choose one arm to visit; only one out of the eight arms will deliver reward at any given time, so the rat must visit different arms over several trials to discover the goal arm. The rat must remember this location and return to it on subsequent trials to receive a large reward (150 uL) at the arm port. After a certain

number of rewards earned at the goal arm (10-15), the goal arm will change, causing the rat to need to search among arms to find the new goal location. Later in training, goal locations are changed after 4-12 repeat visits. Any deviation from the trial structure described results in a 30-45 s “timeout” during which all ports deilluminate; no reward will be delivered from any port until the timeout period is up and the subject initiates a new trial at the home port. Pretraining continues until the subject can perform reliably; usually 14 days of training with several 30-60 min training epochs per day.

Neural implant

Best performing subjects are selected for surgery and allowed *ad libitum* food access without behavioral training for at least 5 days prior to implant. Each implant houses 30 independently moveable nichrome tetrodes (12.7 μ m, California Fine Wire), cut at an angle and gold plated to a final impedance of 210-350 kOhms. The implant is stereotactically positioned such that each of the bilateral cannulae containing 15 tetrodes is centered at 3.6-4 mm AP, $\pm$ 2.6 mm ML relative to skull bregma. A screw placed over cerebellum serves as global reference. Subjects were allowed at least 5 and up to 10 days for recovery after surgery, with *ad libitum* food and daily tetrode lowering until dorsal CA1 cell layer is reached. One tetrode per hemisphere is targeted to the corpus callosum to serve as a local reference for each cannula.

Online SWR detection and neurofeedback training

After recovery, manipulation cohort subjects are allowed 1-2 days of the pretraining task paradigm to resume reliable performance before beginning the neurofeedback phase of training. One center port becomes the neurofeedback port, at which subjects must hold their nose in the port until a suprathreshold SWR is detected in order to trigger a sound cue and reward. The other port (delay port) also requires subjects to hold their nose in the port, but the sound cue and reward is delivered after a certain amount of time, independent of neural activity. This approach constrains the behavior of the animal to be overtly similar during both trial types. Delay times are chosen from the latencies until trigger SWR generation from the 8 prior neurofeedback trials, which results in the times spent at neurofeedback and delay ports being approximately yoked over the course of the epoch. We did not use the more traditional approach used in SWR intervention studies in which the manipulation is performed at a fixed offset after SWR detection (Fernandez-Ruiz et al., 2019 [DOI](#); Girardeau et al., 2009 [DOI](#); Jadhav et al., 2012 [DOI](#)), because in our paradigm this would introduce a systematic difference in the time spent waiting for reward at the two ports and would thus risk biasing the subject’s behavior. The neurofeedback/delay port locations were counterbalanced across subjects. Subjects complete 2-3 30-90 min behavioral epochs each day; 30 min rest epochs in a small rest box precede and follow each behavioral epoch.

During behavioral epochs, SWRs are detected online by monitoring the smoothed envelope of ripple filtered (100-400 Hz) LFP from 4-6 CA1 cell layer tetrodes (Supplementary Figure 1A). The mean μ_{est} and standard deviation σ_{est} of the absolute value of the ripple filtered trace for each tetrode is calculated using an iterative procedure (Jadhav et al., 2012 [DOI](#)):

$$\mu_{est}(n) = \mu_{est}(n-1) \cdot (N_{smooth} - 1)/N_{smooth} + |x|/N_{smooth}$$

$$\sigma_{est}(n) = (||x| - \mu_{est}(n-1)| - \sigma_{est}(n-1)/N_{smooth} + \sigma_{est}(n-1)$$

Here x is the absolute value of the ripple filtered LFP, sampled at 1500 Hz, and N_{smooth} is the number of samples used for smoothing (10,000). Values are measured prior to each behavioral epoch while the subject is alert and active in the rest box, allowed to stabilize, fixed, and used to determine the SWR detection threshold for the subsequent behavioral epoch. During behavior, the

envelope v_{est} of the ripple filtered LFP from the monitored tetrodes is calculated and smoothed using an asymmetrical iterative estimator that allows for rapid detection of power increases (Jadhav et al., 2012 [DOI](#)):

$$v_{est}(n) = v_{est}(n-1) + g_i(n-1) \cdot (|x| - v_{est}(n-1))$$

where g_t is 0.2 when the envelope was decreasing ($|t| \leq v_{est}(n-1)$) and g_i is a moving average of 1.2 and the last 19 values of g , when the envelope was increasing ($|t| > v_{est}(n-1)$). SWR detection is triggered when the ripple envelope of 2 or more tetrodes crosses the detection threshold simultaneously. Trigger detection is followed by a delay of 50-100 ms before the sound cue and reward is delivered (75 ms for rat1, 100 ms for rat2, 50 ms for rat3 and rat4; see Supplementary Figure 1B), in order to minimize the chance of aborting or interrupting the ongoing SWR event with abrupt stimuli. The set of detection tetrodes changes each day and always includes tetrodes from each hemisphere. The threshold for trigger SWR detection is gradually raised, across 14-20 days, from 4 sd to 16-20 sd, and then is maintained at the maximum threshold for 7-10 days (Supplementary Figure 1E). Within each behavioral epoch, the detection threshold starts low (4 sd) and is raised to the maximum threshold for that day over the course of the first ~10 neurofeedback trials.

Data collection and processing

Continuous 30 kHz neural data, environmental events (port beam breaks, lights, reward delivery, online SWR triggers, etc), online position tracking of a head-mounted LED array, and 30 Hz overhead video were collected and synchronized using a SpikeGadgets data acquisition system. LFP was generated by filtering between 0.1-300 Hz and was referenced to a tetrode located in ipsilateral corpus callosum. In parallel, the continuous signal was referenced and filtered between 600-6000 Hz, and spike events were detected when the voltage exceeded 100 μ V on any channel of a tetrode.

Offline SWR detection

SWRs were detected offline using a consensus method incorporating information from all CA1 cell layer tetrodes (Supplementary Figure 1A). LFP traces were filtered for ripple band activity (150-250 Hz), squared, summed over tetrodes, smoothed with a 4 ms sd gaussian kernel, and the square root was taken (Kay et al., 2016 [DOI](#)). The resultant consensus trace was used for event detection with a threshold set 2 sd above the behavioral epoch mean and a minimum suprathreshold duration of 15 ms. SWR start time was defined as the time when the consensus trace crossed epoch baseline prior to threshold crossing and end time was when the consensus trace returned to baseline (Gillespie et al., 2021 [DOI](#)). SWR size was calculated as the maximum threshold value at which the event would have been detected (in units of sd), and SWR length was calculated as the difference between start and end times for each event.

Histology

At the conclusion of data collection, tetrode locations were marked with electrolytic lesions and the subject was transcardially perfused with paraformaldehyde 24hrs later. The brain was fixed, sectioned into 50 μ m slices, and Nissl stained to enable the localization of tetrode tips (Sosa et al., 2020 [DOI](#)).

Quantification and statistical analysis

All analyses were performed using custom code written in Matlab 2020a (Mathworks) and Python 3.6.

Statistical tests

All quantification of manipulation cohort trials (with the exception of Supplementary Figure 2A,B) use only trials from the stable period of neurofeedback training. For the control cohort, trials from all days of reliable post-implant task performance are used. Groupwise comparisons are performed using generalized mixed-effects models to account for and measure both individual variability and group effects. Using the Matlab *fitglm* function, we construct a model in which each sample (SWR event features, trials, etc) is assigned to an individual subject as well as to a cohort (manipulation or control). Linear models are used for rate, dwell time, SWR amplitude, and proportion measures, while a Poisson distribution is used for count measures. To compare neurofeedback and delay trials within each manipulation cohort subject, we use nonparametric ranksum tests. We then perform post-hoc correction for multiple comparisons using the Benjamini-Hochberg method to control false discovery rate (Groppe, 2022 [DOI](#)).

Neurofeedback wait time prediction

For each behavioral session during the first ~7 days of pretraining and during early stages of neurofeedback training, SWRs occurring during center port times (pre-and post- reward) were binned by size and occurrence rate was calculated by dividing the number of SWRs within a given bin by the total time spent at the center port, cumulative over all behavioral epochs. Occurrence rates were only calculated for size bins containing at least 10 events. For later days, a predicted wait time was generated for each neurofeedback trigger event based on the occurrence rate calculated for that size of event. Events without a valid occurrence rate size bin were excluded from analysis. For each subject, actual and predicted wait times at the neurofeedback port were concatenated over behavioral sessions and smoothed with a 200-trial moving average (Supplementary Figure 2B).

SWR size quantification

The structure of the neurofeedback phase of the task guarantees that each neurofeedback trial includes *exactly one* suprathreshold SWR event: the pre-reward period continues until the first suprathreshold event occurs and then immediately ends. However, no such criteria is enforced upon delay trials. Therefore, comparing SWR size taking all events during the pre-reward period of each trial type results in a biased comparison. To control for this bias, the suprathreshold SWR at the end of each neurofeedback trial as well as any suprathreshold SWRs detected during delay trials are excluded from certain analyses (Supplementary Figure 2C, D, and E).

Timewise SWR rate quantification

Each trial, from initial center port nosepoke to exit from the center port, was divided into 0.5 s bins and the count of SWRs in each time bin was calculated. Trials were aligned to the tone/food reward delivery and the mean and S.E.M. for each time bin was calculated across all trials and converted to Hz. Only bins with at least 100 trials contributing data were included. The trigger SWR for each neurofeedback trial was excluded from analysis to avoid a large spike in the trace just prior to reward delivery that distracted from the comparisons between the SWR rates before and after the trigger event.

Clusterless decoding

Marks were calculated for each event as the amplitude of the spike waveform on each of the four tetrode channels at the time of peak amplitude following threshold crossing. Bayesian marked point process clusterless decoding with a state space movement model was performed as described in Denovellis et al, 2021 [DOI](#) and exactly as implemented for this maze environment in Gillespie et al, 2021 [DOI](#). Briefly, tracked position was linearized such that positions in the central 2D

area of the maze were collapsed onto a single line extending from the home well to the edge of the arms. Linearized position was binned into 5 cm bins. The encoding model comprised of all marks from CA1 tetrodes collected when the subject was moving >4 cm/s, and 2 ms time bins were used. Two movement models were included: one with an equal probability of moving one position bin forward, back, or staying in the same position (continuous state), and one with an equal probability of moving to any other possible position bin (fragmented state). A uniform discrete transition matrix governed transitions between movement states, and a resultant joint posterior was computed. For the decoding of movement times the spike being decoded was always excluded from the encoding model.

Replay content analysis

Events were categorized as interpretable if the probability of the continuous movement model exceed a 0.8 threshold during the event and if at least 30% of the posterior density was in a single maze segment. Interpretable events were categorized as local if the majority of posterior density was located in the same maze segment as the animal's current position and remote otherwise. Poisson generalized linear models were used to quantify the effect of a given behavioral arm category on the likelihood of replay (Gillespie et al., 2021 [DOI](#)). Search and repeat trials from the second goal block onwards were used for GLM analysis (so that all trials included a valid previous goal option). Using the MATLAB *fitglm* function, we constructed one binary predictor per arm category (previous arm, future arm, and previous goal arm); each trial contributed one entry per maze arm with a one if the arm was described by the category and a zero if it was not. The response variable for each entry was the number of replays of the arm seen on that trial. The resultant model coefficients were converted to fold change by exponentiation and confidence intervals were calculated using MATLAB's *coefCI* function.

References

- Aleman-Zapata A., van der Meij J., Genzel L. (2021) **Disrupting ripples: Methods, results, and caveats in closed-loop approaches in rodents.** *J Sleep Res* **e 13532**
- Battaglia F.P., Sutherland G.R., McNaughton B.L. (2004) **Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions** *Learn Mem* **11**:697–704
- Buzsaki G. (2015) **Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning** *Hippocampus*
- Carey A.A., Tanaka Y., van der Meer M.A.A. (2019) **Reward revaluation biases hippocampal replay content away from the preferred outcome** *Nat Neurosci* **22**:1450–1459
- Chen Z., Kloosterman F., Layton S., Wilson M.A. (2012) **Transductive neural decoding for unsorted neuronal spikes of rat hippocampus** *Annu Int Conf IEEE Eng Med Biol Soc* **2012**:1310–1313
- Deng X., Liu D.F., Karlsson M.P., Frank L.M., Edén U.T. (2016) **Rapid Classification of Hippocampal Replay Content for Real-time Applications** *J Neurophysiol*, **jn 151**
- Denovellis E.L., Gillespie A.K., Coulter M.E., Sosa M., Chung J.E., Edén U.T., Frank L.M. (2021) **Hippocampal replay of experience at real-world speeds** *Elife* **10**
- Ego-Stengel V., Wilson M.A. (2010) **Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat** *Hippocampus* **20**:1–10
- Engelhard B., Ozeri N., Israel Z., Bergman H., Vaadia E. (2013) **Inducing gamma oscillations and precise spike synchrony by operant conditioning via brain-machine interface** *Neuron* **77**:361–375
- Fernandez-Ruiz A., Oliva A., Fermino de Oliveira E., Rocha-Almeida F., Tingley D., Buzsaki G. (2019) **Long- duration hippocampal sharp wave ripples improve memory** *Science* **364**:1082–1086
- Ferster C.B., Skinner B.F. (1957) **Schedules of Reinforcement** *Psychological Reports* **3**:695–695
- Findlay G., Tononi G., Cirelli C. (2020) **The evolving view of replay and its functions in wake and sleep** *Sleep Adv* **1**
- Frankland P.W., Bontempi B. (2005) **The organization of recent and remote memories** *Nat Rev Neurosci* **6**:119–130
- Gillespie A.K., Astudillo Maya D.A., Denovellis E.L., Liu D.F., Kastner D.B., Coulter M.E., Roumis D.K., Edén U.T., Frank L.M. (2021) **Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice** *Neuron*
- Gillespie A.K., Jones E.A., Lin Y.H., Karlsson M.P., Kay K., Yoon S.Y., Tong L.M., Nova P., Carr J.S., Frank L.M., Huang Y. (2016) **Apolipoprotein E4 Causes Age-Dependent Disruption of Slow Gamma Oscillations during Hippocampal Sharp-Wave Ripples** *Neuron* **90**:740–751

Girardeau G., Benchenane K., Wiener S.I., Buzsaki G., Zugaro M.B (2009) **Selective suppression of hippocampal ripples impairs spatial memory** *Nat Neurosci* **12**:1222–1223

Girardeau G., Cei A., Zugaro M (2014) **Learning-induced plasticity regulates hippocampal sharp wave-ripple drive** *J Neurosci* **34**:5176–5183

Gridchyn I., Schoenenberger P., O'Neill J., Csicsvari J (2020) **Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit** *Neuron* **106**:291–300

Groppe D. (2022) **fdr_bh. MATLAB Central File Exchange**

Gupta A.S., van der Meer M.A., Touretzky D.S., Redish A.D. (2010) **Hippocampal replay is not a simple function of experience** *Neuron* **65**:695–705

Iaccarino H.F., Singer A.C., Martorell A.J., Rudenko A., Gao F., Gillingham T.Z., Mathys H., Seo J., Kritskiy O., Abdurrob F., Adaikkan C., Canter R.G., Rueda R., Brown E.N., Boyden E.S., Tsai L.H (2016) **Gamma frequency entrainment attenuates amyloid load and modifies microglia** *Nature* **540**:230–235

Ishikawa D., Matsumoto N., Sakaguchi T., Matsuki N., Ikegaya Y (2014) **Operant conditioning of synaptic and spiking activity patterns in single hippocampal neurons** *J Neurosci* **34**:5044–5053

Isomura Y., Sirota A., Ozen S., Montgomery S., Mizuseki K., Henze D.A., Buzsaki G (2006) **Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations** *Neuron* **52**:871–882

Jadhav S.P., Kemere C., German P.W., Frank L.M (2012) **Awake hippocampal sharp-wave ripples support spatial memory** *Science* **336**:1454–1458

Jiang H., Liu S., Geng X., Caccavano A., Conant K., Vicini S., Wu J (2018) **Pacing Hippocampal Sharp-Wave Ripples With Weak Electric Stimulation** *Front Neurosci* **12**

Jiang Y., Jessee W., Hoyng S., Borhani S., Liu Z., Zhao X., Price L.K., High W., Suhl J., Cerel-Suhl S (2022) **Sharpening Working Memory With Real-Time Electrophysiological Brain Signals: Which Neurofeedback Paradigms Work?** *Front Aging Neurosci* **14**

Jones E.A., Gillespie A.K., Yoon S.Y., Frank L.M., Huang Y (2019) **Early Hippocampal Sharp-Wave Ripple Deficits Predict Later Learning and Memory Impairments in an Alzheimer's Disease Mouse Model** *Cell Rep* **29**:2123–2133

Joo H.R., Frank L.M (2018) **The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation** *Nat Rev Neurosci* **19**:744–757

Kaefer K., Stella F., McNaughton B.L., Battaglia F.P (2022) **Replay, the default mode network and the cascaded memory systems model** *Nat Rev Neurosci*

Kay K., Sosa M., Chung J.E., Karlsson M.P., Larkin M.C., Frank L.M (2016) **A hippocampal network for spatial coding during immobility and sleep** *Nature* **531**:185–190

Klink K., Jaun U., Federspiel A., Wunderlin M., Teunissen C.E., Kiefer C., Wiest R., Scharnowski F., Sladky R., Haugg A., Hellrung L., Peter J (2021) **Targeting hippocampal hyperactivity with real-time fMRI neurofeedback: protocol of a single-blind randomized controlled trial in mild cognitive impairment** *BMC Psychiatry* **21**

- Klinzing J.G., Niethard N., Born J (2019) **Mechanisms of systems memory consolidation during sleep** *Nat Neurosci* **22**:1598–1610
- Kloosterman F., Layton S.P., Chen Z., Wilson M.A (2014) **Bayesian decoding using unsorted spikes in the rat hippocampus** *J Neurophysiol* **111**:217–227
- Liu Y., Dolan R.J., Kurth-Nelson Z., Behrens T.E.J (2019) **Human Replay Spontaneously Reorganizes Experience** *Cell*
- Maingret N., Girardeau G., Todorova R., Goutierre M., Zugaro M (2016) **Hippocampo-cortical coupling mediates memory consolidation during sleep** *Nat Neurosci* **19**:959–964
- 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** *Front Behav Neurosci* **6**
- 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* **359**:1524–1527
- Nour M.M., Liu Y., Arumham A., Kurth-Nelson Z., Dolan R.J (2021) **Impaired neural replay of inferred relationships in schizophrenia** *Cell* **184**:4315–4328
- Novitskaya Y., Sara S.J., Logothetis N.K., Eschenko O (2016) **Ripple-triggered stimulation of the locus coeruleus during post-learning sleep disrupts ripple/spindle coupling and impairs memory consolidation** *Learn Mem* **23**:238–248
- Olafsdottir H.F., Barry C., Saleem A.B., Hassabis D., Spiers H.J (2015) **Hippocampal place cells construct reward related sequences through unexplored space** *Elife* **4**
- Oliva A., Fernandez-Ruiz A., Leroy F., Siegelbaum S.A (2020) **Hippocampal CA2 sharp-wave ripples reactivate and promote social memory** *Nature* **587**:264–269
- Prince S.M., Paulson A.L., Jeong N., Zhang L., Amigues S., Singer A.C (2021) **Alzheimer's pathology causes impaired inhibitory connections and reactivation of spatial codes during spatial navigation** *Cell Rep* **35**
- Reiner M., Rozengurt R., Barnea A (2014) **Better than sleep: theta neurofeedback training accelerates memory consolidation** *Biol Psychol* **95**:45–53
- Sanchez-Aguilera A., Quintanilla J.P. (2021) **Sharp Wave Ripples in Alzheimer's Disease: In Search of Mechanisms** *J Neurosci* **41**:1366–1370
- Sitaram R., Ros T., Stoeckel L., Haller S., Scharnowski F., Lewis-Peacock J., Weiskopf N., Blefari M.L., Rana M., Oblak E., Birbaumer N., Sulzer J (2017) **Closed-loop brain training: the science of neurofeedback** *Nat Rev Neurosci* **18**:86–100
- Skelin I., Kilianski S., McNaughton B.L (2019) **Hippocampal coupling with cortical and subcortical structures in the context of memory consolidation** *Neurobiol Learn Mem* **160**:21–31
- Sosa M., Joo H.R., Frank L.M (2020) **Dorsal and Ventral Hippocampal Sharp-Wave Ripples Activate Distinct Nucleus Accumbens Networks** *Neuron* **105**:725–741

- Stark E., Roux L., Eichler R., Buzsaki G (2015) **Local generation of multineuronal spike sequences in the hippocampal CA1 region** *Proc Natl Acad Sci U S A* **112**:10521–10526
- Stark E., Roux L., Eichler R., Senzai Y., Royer S., Buzsaki G (2014) **Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations** *Neuron* **83**:467–480
- Staufenbiel S.M., Brouwer A.M., Keizer A.W., van Wouwe N.C. (2014) **Effect of beta and gamma neurofeedback on memory and intelligence in the elderly** *Biol Psychol* **95**:74–85
- Stoiljkovic M., Kelley C., Horvath T.L., Hajos M (2018) **Neurophysiological signals as predictive translational biomarkers for Alzheimer’s disease treatment: effects of donepezil on neuronal network oscillations in TgF344-AD rats** *Alzheimers Res Ther* **10**
- Todorova R., Zugaro M (2020) **Hippocampal ripples as a mode of communication with cortical and subcortical areas** *Hippocampus* **30**:39–49
- Trambaiolli L.R., Cassani R., Mehler D.M.A., Falk T.H (2021) **Neurofeedback and the Aging Brain: A Systematic Review of Training Protocols for Dementia and Mild Cognitive Impairment** *Front Aging Neurosci* **13**
- Vandecasteele M., Varga V., Berenyi A., Papp E., Bartho P., Venance L., Freund T.F., Buzsaki G (2014) **Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus** *Proc Natl Acad Sci U S A* **111**:13535–13540
- Wang J.R., Hsieh S (2013) **Neurofeedback training improves attention and working memory performance** *Clin Neurophysiol* **124**:2406–2420
- White M., Mackay M., Whittaker R (2020) **Taking Optogenetics into the Human Brain: Opportunities and Challenges in Clinical Trial Design** *Open Access J Clin Trials* **2020**:33–41
- Yu J.Y., Kay K., Liu D.F., Grossrubatscher I., Loback A., Sosa M., Chung J.E., Karlsson M.P., Larkin M.C., Frank L.M (2017) **Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility** *Elife* **6**
- Yu Z., Guindani M., Grieco S.F., Chen L., Holmes T.C., Xu X (2022) **Beyond t test and ANOVA: applications of mixed-effects models for more rigorous statistical analysis in neuroscience research** *Neuron* **110**:21–35
- Zhang K., Ginzburg I., McNaughton B.L., Sejnowski T.J (1998) **Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells** *J Neurophysiol* **79**:1017–1044

Author information

Anna K. Gillespie

Departments of Biological Structure and Lab Medicine & Pathology, University of Washington, Seattle, WA, USA, Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA

For correspondence: annagill@uw.edu

ORCID iD: [0000-0003-0980-2408](https://orcid.org/0000-0003-0980-2408)

Daniela A. Astudillo Maya

Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA

Eric L. Denovellis

Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA, Howard Hughes Medical Institute, Chevy Chase, MD, USA
ORCID iD: [0000-0003-4606-087X](https://orcid.org/0000-0003-4606-087X)

Sachi Desse

Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA

Loren M. Frank

Departments of Physiology and Psychiatry and the Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA, USA, Howard Hughes Medical Institute, Chevy Chase, MD, USA
ORCID iD: [0000-0002-1752-5677](https://orcid.org/0000-0002-1752-5677)

Editors

Reviewing Editor

Laura Colgin

University of Texas at Austin, United States of America

Senior Editor

Laura Colgin

University of Texas at Austin, United States of America

Reviewer #1 (Public Review):

Summary:

This is an important study that tests the effects of using neurofeedback, in the form of reward delivery when large sharp wave-ripples (SWRs) are detected, on neurophysiological and behavioral measures. The authors report that the rate of SWRs ripples increased prior to reward delivery, but this increased rate of SWRs had no significant effect on memory performance. They also found that compensatory decreases in SWR rate occurred in the period after reward delivery such that the overall SWR rate remained stable.

Strengths:

The study has many strengths. The paradigm of closed loop detection of SWRs and reward delivery is powerful and provides an innovative way to causally test the effects of increasing SWR rates. Other studies could adopt this method to test other hypotheses or to assess the effects of increasing SWR rates prior to reward delivery in rodent models of brain disorders. The methods and results are clearly explained. The results are presented in a transparent way.

Weaknesses:

In the linear mixed effects model analysis used in Figure 2, and statistics reported in the figure legend, an interaction effect showing that neurofeedback differentially affected the SWR rate and count pre- and post-award seems to be missing in the reported statistics.

In the Discussion, the authors write, "Further, because subjects learn to modulate SWR rate, rather than simply generating a single suprathreshold event on command, it is likely that they learn to engage a SWR-permissive state during the targeted interval in which brain-wide neural activity and neuromodulatory tone also enter a SWR-permissive realm". This seems to imply that the neurofeedback is directly modulating neural activity. However, it is unclear from the paper exactly how the neurofeedback is modulating the SWR rate. Considering that SWRs occur during immobility, is it possible that the animals are learning to remain more immobile and modulating the SWR rate in that way?

Reviewer #2 (Public Review):

Gillespie et al. introduced a novel neurofeedback (NF) procedure to train rats in enhancing their sharp-wave ripple (SWR) rate within a short duration, a key neural mechanism associated with memory consolidation. The training, embedded within a spatial memory task, spanned 20-30 days and utilized food rewards as positive reinforcement upon SWR detection. Rats were categorized into NF and control groups, with the NF group further divided into NF and delay trials for within-subject control. While single trial differences were elusive due to the variability of SWR occurrence, the study revealed that statistically rats in NF trials exhibited a notably higher SWR rate before receiving rewards compared to delay trials. This difference was even more pronounced when juxtaposed with rats not exposed to NF training (control group). The unique design of blending the NF phase with the memory dependent spatial task enabled the authors to analyze whether the NF training influence the task performance and replay content during SWRs across three different conditions (NF trials, delay trials and control group). Interestingly, despite the NF training, there was no significant improvement or decline in the performance of the spatial memory task, and the replay content remained consistent across all three conditions. Hence, the operant conditioning only amplified the SWR rate before reward in NF trials without altering the task performance and the replay content during SWR. Moreover, considering the post-reward period, the total SWR count was consistent across all conditions as well, meaning the NF training also do not affect the total SWR count. The study concludes with the hypothesis of a potential homeostatic mechanism governing the total SWR production in rats. This research significantly extends previous work by Ishikawa et al. (2014), offering insights into the NF training with external reward on the SWR rate/counts, replay content and task performance.

Strengths:

- Integration of NF task and spatial memory task in a single trial

The integration of NF training within a spatial memory task poses significant challenges. Gillespie and colleagues overcame this by seamlessly blending the NF task and the spatial memory task into a single trial. Each trial involved a rat undergoing three steps: First, initiating a trial. Second, moving to either the NF port or the delay trial port, as indicated by an LED, and then maintaining a nosepoke at one of the center ports. During this step, the rat had to keep its nose (in the NF port) until a sharp-wave ripple (SWR) exceeding a set threshold was detected, which then triggered a reward, or until a variable time elapsed (in the delay port). Third, the rat would choose one of eight arms to explore before starting the next trial. This integration of the two tasks (step two as the NF task and step three as the spatial memory task) facilitated a direct analysis of the impact of NF training on behaviorally relevant replay content during SWRs and the performance in the spatial memory task.

- Clear Group Separation

A robust study design necessitates clear distinctions between experimental conditions to

ensure that observed differences can be attributed to the variable under investigation. This study meticulously categorized rats into three distinct conditions: NF trials, delay trials (for within-subject control), and a control group (for across-subject control). Furthermore, for each trial, the times of interest (TOI) were separated into pre-reward and post-reward periods. This clear separation ensures that any observed differences in SWR rates and other outcomes can be confidently attributed to the effects of neurofeedback training during specific time periods, minimizing potential confounding factors.

- Evidence of SWR rate modulation

The study's results offer compelling evidence that rats can be trained to modulate their SWR rates during the pre-reward period. This is evident from the observation that rats in the NF trials consistently displayed a higher SWR rate before receiving rewards compared to those in delay trials or the control group (Fig. 2). Such findings not only validate the efficacy of the NF paradigm but also underscore the potential of operant conditioning in influencing neural mechanisms. The observation that rats were able to produce larger SWR events by modulating their occurrence rate, rather than merely waiting for these events, suggests a learned strategy to generate them more efficiently.

- Evidence of SWR count homeostasis

A notable finding from the study was the observation of a consistent total SWR count during both pre-reward and post-reward periods across all conditions, despite the evident increase in SWR rates during the pre-reward period in NF trials. This points to a potential homeostatic mechanism governing SWR production in rats. This balance suggests that while NF training can modulate the timing and rate of SWRs over a short duration, it doesn't influence the overall count of SWRs over a longer period. Such a mechanism might be essential in ensuring that the brain neither overcompensates nor depletes its capacity for SWRs, maintaining the overall neural balance and functionality. This discovery deepens our understanding of neural mechanisms and highlights potential avenues for future research into the regulatory processes governing neural activity.

Weaknesses:

- Misleading Title

The title, "Neurofeedback training can modulate task-relevant memory replay in rats," implies that through neurofeedback training, rats can learn to modulate the content of their memory replay. However, the study's findings contradict this implication. Particularly, one of the subtitles of this paper is "Neurofeedback training preserves replay content during SWRs," which directly contrasts with the main title's suggestion. The authors conclusively demonstrated that there was no discernible difference in the replay content between animals that underwent NF training and those that did not. The current title easily leads to misinterpretations about the study's primary outcomes, especially for readers who might not delve into the detailed findings.

- Lack of control analysis baseline for each animal

While the authors meticulously categorized trial types into three distinct conditions: NF trials, delay trials, and control groups, they did not clearly establish a baseline for each animal. The animal could have a total different baseline SWR rates. The paper appears to operate under the assumption that each animal possesses a consistent SWR rate baseline, leading to only the final comparisons being presented.

- Vagueness of what animal really control during NF trials after training

The authors state that, "Moreover, although we did observe a slightly lower mean speed during the pre-reward period on neurofeedback trials compared to delay trials and trials from the control cohort (Supplementary Figure 2F), movement differences could not explain the difference in SWR rates (Supplementary Figure 2G, H)." This assertion raises questions about the underlying mechanisms at play. In a typical operant conditioning scenario, training

could result in direct neural modulation, behavioral changes, or a combination of both. For instance, rats might adopt a more stationary posture during the pre-reward period on NF trials compared to other conditions, or they might actively influence the occurrence rate of SWRs during this period. The paper would benefit from a clearer delineation of what the animals are specifically controlling or modulating during the NF trials, ensuring a more comprehensive understanding of the observed effects.

- Clinical Implications

The study was conducted on healthy, young animals but suggests potential benefits for older, cognitively impaired animals. However, it's possible that older or deficit animals might not respond to the NF protocol in the same way.

Reviewer #3 (Public Review):

Summary:

This study implements an innovative neurofeedback procedure in rats, providing food reward upon detection of a sharp wave-ripple event (SWR) in the hippocampus. The elegant experimental design enables a within-animal comparison of the effects of this neurofeedback procedure as compared to a control condition in which an equivalent reward is provided in a non-contingent manner. The neurofeedback procedure was found to increase SWR rate, followed by a compensatory reduction in SWR rate. These changes in SWR rate were not accompanied by any changes in memory performance on the memory-guided task.

Strengths:

The scientific premise for the study is outstanding. It addresses an issue of high importance, of developing ways to not merely describe correlations between SWRs (and their content) and memory performance, but to manipulate them. The authors argue clearly and convincingly that even studies that have performed causal manipulations of SWRs have important confounds and limitations, and most importantly for translational purposes, they are all invasive. So, the idea of developing a potentially non-invasive neurofeedback procedure for modulating SWRs is compelling both as an innovative new experimental manipulation in studies of SWRs, and as a potentially impactful therapeutic avenue.

In addition to addressing an important issue with an innovative approach, the study has many other strengths. The data unambiguously show that the method is effective at increasing SWR rate in each individual subject. The experimental design allows within-subject comparison of neurofeedback and control trials, where the subjects wait an equivalent amount of time. The careful analyses of SWR properties and their content establish that neurofeedback SWRs are comparable to control SWRs. The data add further evidence to the notion that SWR rate is subject to homeostatic control. The paper is also exceptionally well written, and was a pleasure to read. So, there is a clear technical advance, in that there is now a method for increasing SWR rate non-invasively, which is rigorously established and characterized.

Weaknesses:

The one overall limitation I find with this study is that it is unclear to what extent the same (or better) results could have been obtained using behavior-feedback instead of neuro-feedback. Because SWR rates are generally higher during states of quiescence compared to active movement or task engagement, it is possible that reinforcing behaviorally detected quiescent states (e.g. low movement) would indirectly increase SWR rates. The observation that all 4 subjects had lower movement speeds during neurofeedback compared to control trials supports this interpretation. This is an important issue that would help clarify whether

the neurofeedback approach is worth the additional effort and expense compared to behavioral feedback.