Replay Supports Both Future and Immediate Decisions by Updating a Successor Representation

Shuheng Guo[1,*] **Alexa Booras**[1,*] **Michael J. Kahana**[2]

Itzhak Fried[3] **John J. Sakon**[3] **Aaron M. Bornstein**[1]

[1] Department of Cognitive Sciences, University of California, Irvine
[2] Department of Psychology, University of Pennsylvania
[3] Department of Neurosurgery, University of California, Los Angeles

jianleg@uci.edu, boorasa@uci.edu, kahana@psych.upenn.edu, ifried@mednet.ucla.edu, jsakon@mednet.ucla.edu, aaron.bornstein@uci.edu

Abstract

Neural replay has been proposed to play two distinct roles: 1) Consolidating memories to support abstraction and 2) guiding planning for decision-making. Here, we examine if a single underlying principle—the successor representation (SR)—explains both. Analyzing single neuron recordings from the hippocampus of neurosurgical patients while they performed a memory-guided maze navigation task probing both spatial and sequential order memory, we find that rapid replay of neuronal pairs during separate task phases differentially influences subsequent behavior. Replay during rest prior to the start of each trial relates to the accuracy of spatial retrieval on future trials, while replay during the sequential order test relates to accuracy on this test for current and future trials–but critically, not vice versa. Modeling these replay events as random and directed walks, respectively, of trajectories along the learned maze structure, we find an SR captures the qualitative effects of replay on future trials in both task phases. Our results suggest neuronal replay supports memory-guided decisions to learn task structure by updating an SR of the environment at distinct timepoints, bringing seemingly disparate roles under the same theoretical framework.

Introduction

Neural replay, a widespread phenomenon, is thought to be critical for both learning and decision-making (Kurth-Nelson et al., 2023; Sagiv et al., 2025; Joo and Frank, 2018). However, findings in rodents are conflicted as to whether replay either reflects recent experience (Gupta et al., 2010) and/or supports upcoming choice (Carr et al., 2011; Singer et al., 2013; Gillespie et al., 2021). Recent work found that human hippocampal and entorhinal neurons encode the temporal structure of experience by tracking an underlying predictive representation of linked concepts (Tacikowski et al., 2024), qualitatively consistent with a successor representation (SR; Dayan, 1993). Further, replay of task-congruent representations during breaks between learning blocks related to increased predictive encoding at a between-subject level (Tacikowski et al., 2024). We propose that the SR explains both the memory and the planning roles of replay, via a single common computation. Here, using a combination of single-unit recordings and model simulations, we find that replay supports future decisions by learning predictive representations—in particular, updating an SR—of environmental structure.

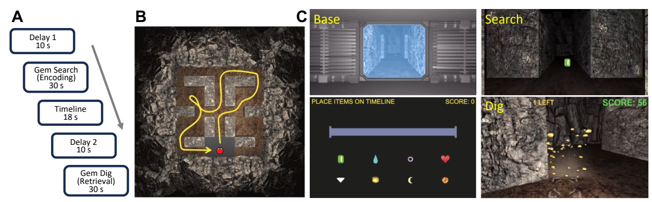

Figure 1: A) Phases in each trial. B) Top-down view of the mine with example trajectory. C) Example task screenshots with labels of task phases.

Methods

Task design: Participants play GemMine on a laptop by navigating a visually sparse underground maze with a mouse and earn points by collecting gems (Fig. 1). The task has five distinct phases that occur serially in each of 36 trials: **Delay 1**, where participants wait in a home base for the trial to start; **Encoding**, where they navigate the maze using the laptop mouse to find four gems randomly scattered throughout the maze on each trial; **Timeline**, where they use the mouse to place the gems on a timeline in the order they found them; **Delay 2**, a second period rest in the home base; and **Retrieval**, where they return to the locations of the (now buried) gems to dig them up. Participants score points for correct timeline placements and gem digs to incentivize correct behavior. Therefore, GemMine enables us to assess both sequential order retrieval during the Timeline phase and spatial memory retrieval during Retrieval (where sequential order retrieval is not required).

Replay Detection: We used high impedance microwires that fit inside typical clinical electrodes to safely (Hefft et al., 2013; Carlson et al., 2018; Nagahama et al., 2023)

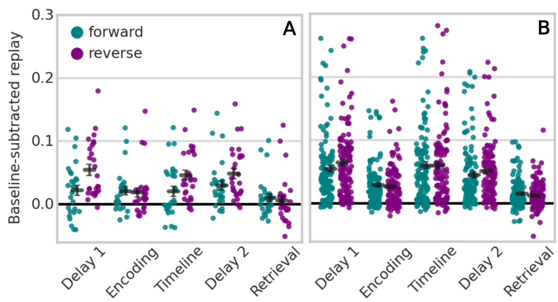

Figure 2: A) Hippocampal replay for pairs (each dot a pair; N=27 from 4 participants) of place cells in each task phase. Forward vs. reverse order is defined by home base as starting point. B) Replay for all hippocampal neuron pairs (N=115; 8 participants). Black bars: mean±SE.

detect single neurons in the hippocampus. We detect hippocampal replay via cell pairs that spike together in 50ms (Tacikowski et al., 2024) above baseline levels found via permutations tests of shuffled spike trains. We confirm that replay reflects encoded, task-related sequences using a place cell analysis akin to those in animal studies, where we expect neuronal replay at the end of navigation to replay in reverse from the order of encoded behavior (Foster and Wilson, 2006; Diba and Buzsáki, 2008). Indeed, we identify enrichment in reverse over forward replay in periods following navigation (Timeline and Delay 1) (Fig. 2).

Regression: To test whether replay contributes to future decisions, we regress observed replay from the current trial (N) against performance on both current (N) and future (N+X) trials for each of the task phases with strong hippocampal replay: Delay 1, Timeline, and Delay 2 (Fig. 3).

Successor Representation Simulations: We modeled replay events as trajectories along the learned maze structure which updated a successor representation (SR). At each step of the walk, the SR matrix M (Dayan, 1993) is updated using the temporal-difference learning rule, with parameters learning rate λ and discount rate γ:

$$M[s_t,:] = M[s_t,:] + \lambda\left(1_{s_t} + \gamma M[s_{t+1},:] - M[s_t,:]\right). \quad (1)$$

For our simulations we set λ to 0.4 for real experience and to 0.01 for fictive experience. We set γ to 0.5. We update the SR in two ways to align with the two possible different roles of replay: random walks on the structure of an environment during Delay 1 and directed trajectories of previous rewards during Timeline and Delay 2. Directed trajectories go from a previously learned high-reward location toward a lower-reward location constrained by the SR. We model Timeline placement using reward signal strength to form start and end nodes. To model the navigation phases of the task we set: Encoding as a sequence of state transitions on the maze structure, and Memory Retrieval as a search process constrained by the SR matrix, with a fixed step limit.

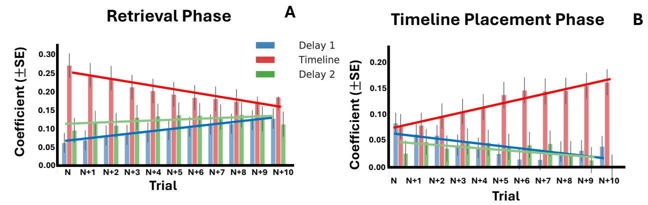

Figure 3: A) Coefficients for regression that models how neural replay on trial N impacts Retrieval digs on current (N) and future ($N + X$) trials. B) Similar regression that predicts correct Timeline placements on current (N) and future ($N + X$) trials from replay on trial N. Errorbars: SE. Dashed lines: linear fit.

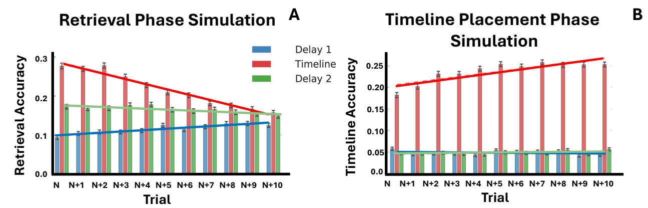

Figure 4: A) Replay simulations of correct gem digs during Retrieval. B) Replay simulations of Timeline sequential order judgments. Lines: linear fit.

Results

In the recordings, we find that Timeline replay on trial N has decreasing effect on future Retrieval accuracy, while Delay 1 replay has increasing effect (Fig. 3A). Replay on trial N increasingly affects future Timeline placement accuracy as compared to decreasing effects from Delay 1 replay (Fig. 3B). Strikingly, our SR simulations qualitatively capture the neural replay effects shown in the fitted data: decreasing effects of Timeline and increasing effects of Delay 1 replay on correct gem digs (Fig. 4A) and increasing effects of Timeline replay on sequential order judgments (Fig. 4B).

Discussion

We hypothesized that neural replay supports both learning and decision-making by updating a successor representation (SR). Using SR simulations that leverage replay in different task phases we find that replay at trial start increasingly influences upcoming retrieval several trials into the future, consistent with a role in long-term structure learning, while replay during decisions increasingly influences future sequential order judgments. These findings provide a unifying theoretical explanation of disparate observations in rodents and humans, especially concerning replay's impact on task learning vs. online behavior (van der Meer and Bendor, 2025). Future work will build on these findings to investigate how the specific content of replay changes with task phase and influences subsequent choices.

Acknowledgments

The authors acknowledge funding from NIMH T32MH119049, NINDS R01NS11946, and U01NS123128.

References

Carlson, A. A., Rutishauser, U., & Mamelak, A. N. (2018). Safety and utility of hybrid depth electrodes for seizure localization and single-unit neuronal recording. *Stereotactic and functional neurosurgery*, *96*(5), 311–319.

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*, *14*(2), 147–153.

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural computation*, *5*(4), 613–624.

Diba, K., & Buzsáki, G. (2008). Hippocampal network dynamics constrain the time lag between pyramidal cells across modified environments. *Journal of Neuroscience*, *28*(50), 13448–13456.

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*(7084), 680–683.

Gillespie, A. K., Maya, D. A. 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.

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*(5), 695–705.

Hefft, S., Brandt, A., Zwick, S., von Elverfeldt, D., Mader, I., Cordeiro, J., Trippel, M., Blumberg, J., & Schulze-Bonhage, A. (2013). Safety of hybrid electrodes for single-neuron recordings in humans. *Neurosurgery*, *73*(1), 78–85.

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.

Kurth-Nelson, Z., Behrens, T., Wayne, G., Miller, K., Luettgau, L., Dolan, R., Liu, Y., & Schwartenbeck, P. (2023). Replay and compositional computation. *Neuron*, *111*(4), 454–469.

Nagahama, Y., Dewar, S., Behnke, E., Eliashiv, D., Stern, J. M., Kalender, G., Fields, T. A., Wilson, C., Staba, R., Engel, J., et al. (2023). Outcome of stereo-electroencephalography with single-unit recording in drug-refractory epilepsy. *Journal of neurosurgery*, *139*(6), 1588–1597.

Sagiv, Y., Akam, T., Witten, I. B., & Daw, N. D. (2025). Between planning and map building: Prioritizing replay when future goals are uncertain. *Neuron*, *113*(24), 4278–4292.

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.

Tacikowski, P., Kalender, G., Ciliberti, D., & Fried, I. (2024). Human hippocampal and entorhinal neurons encode the temporal structure of experience. *Nature*, *635*(8037), 160–167.

van der Meer, M. A., & Bendor, D. (2025). Awake replay: Off the clock but on the job. *Trends in Neurosciences*, *48*(4), 257–267.