

Extended Abstract Track

Changes in the geometry of hippocampal representations across brain states

Wannan Yang

New York University

WINNIEYANGWN96@GMAIL.COM

Chen Sun

Mila Quebec AI Institute

SUNCHIPSSTER@GMAIL.COM

György Buzsáki

New York University

GYORGY.BUZSAKI@NYULANGONE.ORG

Editors: Sophia Sanborn, Christian Shewmake, Simone Azeglio, Nina Miolane

Abstract

The hippocampus (HPC) is a key structure of the brain's capacity to learn and generalize. One pervasive phenomenon in the brain, but missing in AI, is the presence of different gross brain states. It is known that these different brain states give rise to diverse modes of information processing that are imperative for hippocampus to learn and function, but the mechanisms by which they do so remain unknown. To study this, we harnessed the power of recently developed dimensionality reduction techniques to shed insight on how HPC representations change across brain states. We compared the geometry of HPC neuronal representations when rodents learn to generalize across different environments, and showed that HPC representation could support both pattern separation and generalization. Next, we compared HPC activity during different stages of sleep. Consistent with the literature, we found a robust recapitulation of the previous awake experience during non rapid eye movement sleep (NREM). But interestingly, such geometric correspondence to previous awake experience was not observed during rapid eye movement sleep (REM), suggesting a very different mode of information processing. This is the first known report of UMAP analysis on hippocampal neuronal data during REM sleep. We propose that characterizing and contrasting the geometry of hippocampal representations during different brain states can help understand the brain's mechanisms for learning, and in the future, can even help design next generation of AI that learn and generalize better.

Keywords: UMAP, memory, hippocampus, REM sleep, learning, generalization

1. Introduction

Memory storage in the brain goes through a two-stage process (14; 3). First, behavioral episodes are rapidly encoded by the key structure called the hippocampus during awake behavior (29). The fast-learning mechanisms of the hippocampus ensures quick and efficient encoding of memories, even down to single-trials (28). However, these representations are vulnerable to interference by newly encoded information. Therefore, new information needs to be gradually integrated to a different learning system during other brain states (for example, REM and NREM sleep). However, the learning mechanisms themselves during those brain states are unknown. In AI, incorporating a sleep-like offline learning stage has started to demonstrate increasingly promising results to support robust learning and generalization

Extended Abstract Track

(9; 6; 17; 10; 26), but truly harnessing mechanisms of learning through different states akin to those seen in the brain is an area of research in its infancy.

In this paper, we study how hippocampal neuronal representations topologically change across brain states (awake task learning, different stages of sleep, novel learning). Our key contributions are as follows:

- HPC representations in awake state can support both pattern separation & generalization.
- HPC representation during REM sleep imply a different form of memory processing beyond faithful reactivation of the recent episodic experience, unlike during NREM sleep.

2. Methods

To date, most efforts to characterize neuronal representation between brain states are supervised methods, which are based on template-matching principles to compare patterns between behavioral states (5; 20). Such supervised approaches are limited because they use one of the brain states (usually during awake learning) to impose an a priori template derived from the external world, in order to analyze another brain state (e.g. sleep), resulting in the neglect of important features intrinsic to sleep. Here, we apply recently developed *unsupervised* approaches (15; 2) to study hippocampal representations, uncovering their intrinsic neuronal activity structure, without needing to match to a template imposed by the external world. First, we applied multidimensional scaling (MDS) (2) to discover shared features in the neuronal representations during generalization across sessions. Second, we applied Uniform Manifold Approximation and Projection (UMAP) for dimension reduction (15) to visualize the intrinsic representation of the brain during different brain states, and compare manifold representations across brain states (Fig. A.1, Fig. B.1, Fig. C.1).

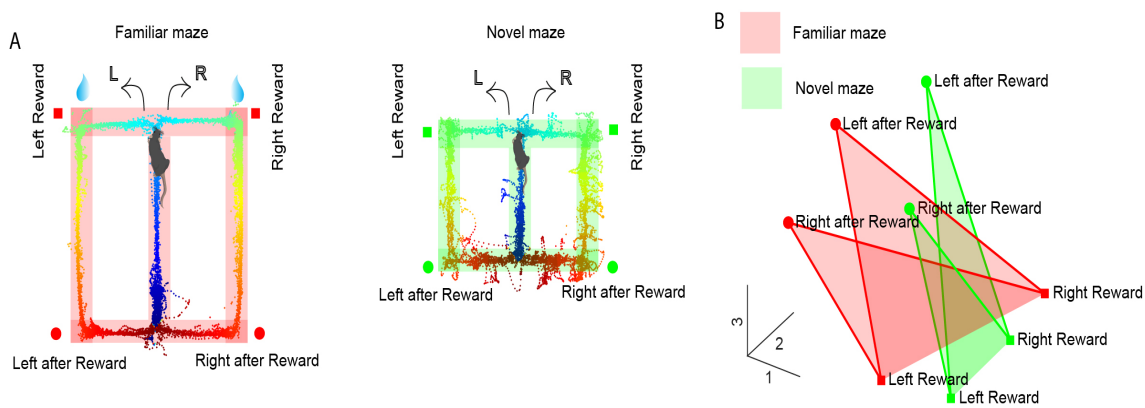


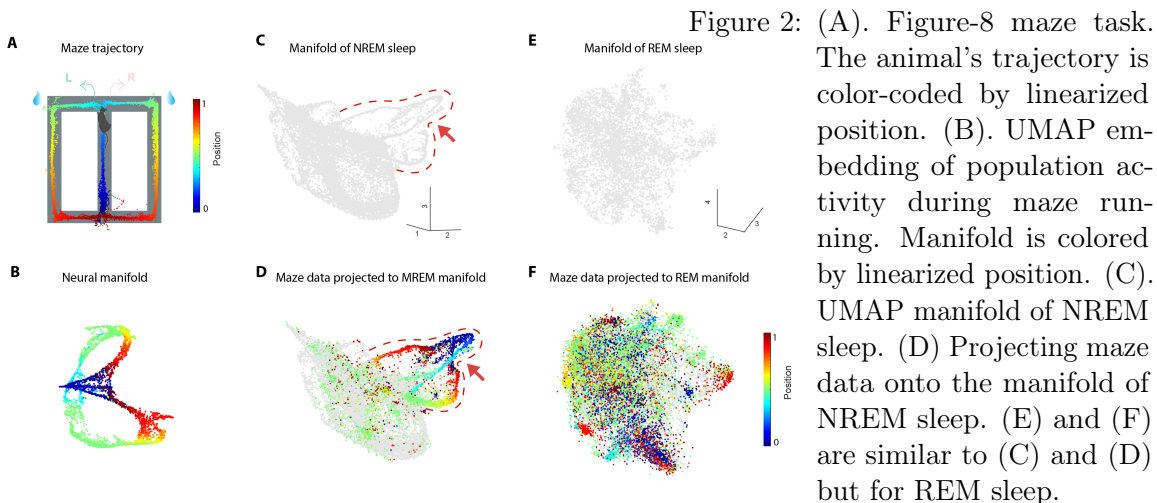
Figure 1: A). Figure-8 maze task, where mice alternated between left and right arms for water reward (blue droplets). The animal’s trajectory along maze corridors is color-coded by linearized position. Left, familiar maze where the animal were trained for one week. Right, novel maze in a novel context where the mouse was introduced for the first time (B). MDS plot showing the dimensionality-reduced firing rates for familiar (red) and novel (green) mazes.

Extended Abstract Track

3. The geometry of neural representation during novel maze learning

The ability to separate different experiences is a critical feature of episodic memory and has long been hypothesized to require the hippocampus (31). But the capacity to pattern separate between situations may impair the capacity to generalize across those situations, and there is still no general consensus to what extent hippocampal pattern separation is not at the expense of forming invariant and generalizable representation.

To explore this question, we used MDS to characterize the hippocampal representations when rodents apply the rule learned in a familiar maze and generalized it to a novel environment (Fig. 1A). Interestingly, here we show that the HPC representation can simultaneously support both pattern separation as well as generalization through ‘analogy-like’ representation (Fig. 1B) (2). In this form of representation, activity of different experiences is pattern separated but are also organized such that the vectors of analogous elements of the different experiences are parallel (for example, the vector connecting the left and right reward in familiar maze is parallel to the left-right reward vector in the novel maze). In a way, this hippocampal capacity to analogize between experiences, is somewhat akin to the analogy-like word embeddings generated by word2vec from natural language processing, as exemplified by the well-known "woman is to queen as man is to king" (16).



4. The geometry of neural representation during different sleep states

Sleep in mammals consists of two core stages: NREM and REM sleep, which alternate in a cyclic manner (19; 23; 24). Reactivation of memory (replay) during NREM sleep and its role in memory have been extensively studied (4; 1; 7; 8; 11; 21; 30). Much less is known about the representations during REM sleep.

Using UMAP, we first visualized the neural representation when rodents learned in a figure-8 maze (Fig. 2 A,B). Next, we visualized the geometric structure of representations during NREM sleep. Interestingly, we uncovered a substructure that resembled the geometry of the maze (red dash line) during NREM sleep. To test if this substructure was actually related to activity during maze exploration, we projected the population activity during maze learning to the space defined by the NREM sleep. We found that indeed, that sub-structure

Extended Abstract Track

during NREM sleep covered the manifold of activity seen during the maze running with high precision (Fig. 2 C,D). Consistent with known literature (4; 1; 7; 8; 11; 21), this directly demonstrates that NREM sleep faithfully reactivated activity from previous behavior and likely helped to consolidate the recent episodic memory. Surprisingly, we were able to demonstrate this using an *unsupervised* technique where the maze-like substructure in the NREM manifold naturally emerged without being forced to match with any pre-defined template, and was therefore a much stronger result comparing to previous template-matching methods.

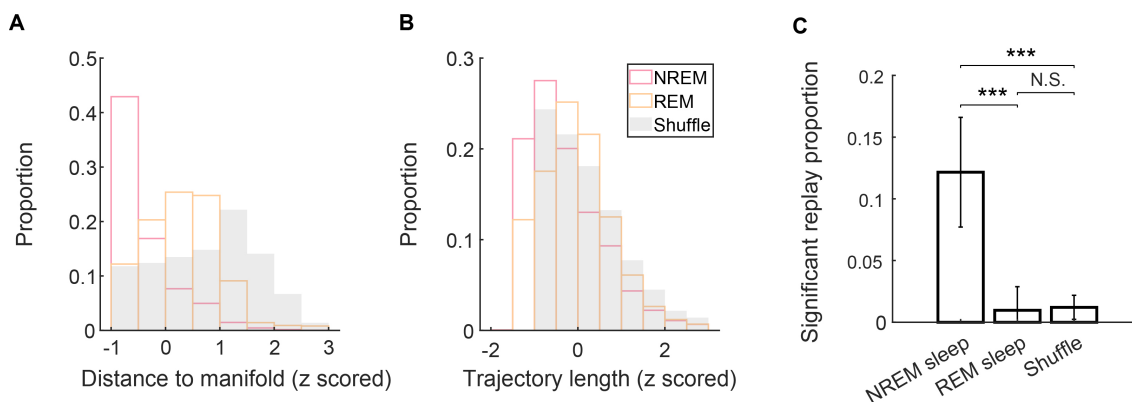


Figure 3: (A) Distribution of distance to maze manifold for replay events across different brain states. (B). Distribution of trajectory length. (C). Proportion of significant replay events across different brain states.

In contrast, we did not observe a maze-like structure during REM sleep (Fig. 2 E,F), suggesting a different role other than faithful reactivation of the recent episodic experience. Compared to NREM sleep, REM sleep events have a longer distance to manifold and trajectory length (Fig. 3 A,B). The two quantities were also used to determine if a candidate sleep event is a significant replay event by comparing it to shuffle (Fig. A.1). The fraction of significant REM replay events is not statistically different from that of shuffle events (Fig. 3 C, Fig. C.1). Our results suggest that the interleaving NREM and REM sleep epochs play different roles during system consolidation. Our results for REM sleep suggest that we might need to update our understanding for the role of REM sleep. In particular, previous studies (32; 13) claimed that replay of maze events seem to present during REM sleep, suggesting that REM sleep plays a similar role as NREM sleep for faithfully reconstructing/replaying the recent episodic experience. However, both of those papers used supervised or template-matching methods. In light of our new results, it is possible that the role of REM sleep is to facilitate the abstraction of general rules from the combinations of many different past experiences. (12), and might therefore not possess faithful reactivation of the most recent experience during this brain state, at all. To our knowledge, this is the first reported characterization of hippocampal population activity during REM sleep using UMAP analysis (or unsupervised dimensionality reduction method in general).

Extended Abstract Track

5. Discussion and future direction

Our result showing hippocampal representations during activate rodent behavior generalizing to a new maze provide novel insight into how hippocampal representation can support both pattern separation of episodic experience in different contexts as well as harbor generalizable representation that can support rapid generalization. By harnessing a linear dimensionality reduction (MDS) we are able to bridge divergent findings showing that HPC representations support pattern separation (31; 27; 22) and other studies showing HPC representations can also support generalization (25; 18). However, this is a preliminary result, and further experiments (for example, in a diverse set of tasks) are needed.

Secondly, our results point to critical differences in information processing by the hippocampus during REM vs NREM sleep. To our knowledge, ours is the first study to characterize hippocampal neuronal activity during REM sleep by UMAP analysis, giving us the capacity to see beyond previous studies on REM sleep that used single neuron analysis or template-matching based decoding analysis. This preliminary result presents an insightful snapshot into hippocampal activity during learning and memory. Future work should aim to characterize hippocampal representational changes over the entire timecourse of learning. This study is a first step to uncover the mysterious role of different brain states (especially during sleep) for learning about the world.

Extended Abstract Track

*Bibliography

- [1] N. Axmacher, C. E. Elger, and J. Fell. Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain*, 131(7):1806–1817, July 2008. ISSN 0006-8950. doi: 10.1093/brain/awn103.
- [2] Silvia Bernardi, Marcus K. Benna, Mattia Rigotti, Jérôme Munuera, Stefano Fusi, and C. Daniel Salzman. The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell*, 183(4):954–967.e21, November 2020. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2020.09.031. URL [https://www.cell.com/cell/abstract/S0092-8674\(20\)31228-9](https://www.cell.com/cell/abstract/S0092-8674(20)31228-9). Publisher: Elsevier.
- [3] G. Buzsáki. Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience*, 31(3):551–570, January 1989. ISSN 03064522. doi: 10.1016/0306-4522(89)90423-5.
- [4] György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188, October 2015. ISSN 1050-9631. doi: 10.1002/hipo.22488.
- [5] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, September 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0460-x. URL <https://www.nature.com/articles/s41593-019-0460-x>. Number: 9 Publisher: Nature Publishing Group.
- [6] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B. Tenenbaum. DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning, June 2020. URL <http://arxiv.org/abs/2006.08381>. arXiv:2006.08381 [cs].
- [7] Antonio Fernández-Ruiz, Azahara Oliva, Eliezyer Fermino de Oliveira, Florbela Rocha-Almeida, David Tingley, and György Buzsáki. Long-duration hippocampal sharp wave ripples improve memory. *Science*, 364(6445):1082–1086, June 2019. ISSN 0036-8075. doi: 10.1126/science.aax0758.
- [8] David J. Foster. Replay Comes of Age. *Annual Review of Neuroscience*, 40(1):581–602, July 2017. ISSN 0147-006X. doi: 10.1146/annurev-neuro-072116-031538.
- [9] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering Atari with Discrete World Models, February 2022. URL <http://arxiv.org/abs/2010.02193>. arXiv:2010.02193 [cs, stat].
- [10] G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. The “wake-sleep” algorithm for unsupervised neural networks. *Science (New York, N.Y.)*, 268(5214):1158–1161, May 1995. ISSN 0036-8075. doi: 10.1126/science.7761831.
- [11] Shantanu P. Jadhav, Caleb Kemere, P. Walter German, and Loren M. Frank. Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science*, 336(6087):1454–1458, June 2012. doi: 10.1126/science.1217230. URL <https://www.science.org/doi/>

Extended Abstract Track

10.1126/science.1217230. Publisher: American Association for the Advancement of Science.

- [12] Penelope A. Lewis, Günther Knoblich, and Gina Poe. How Memory Replay in Sleep Boosts Creative Problem-Solving. *Trends in Cognitive Sciences*, 22(6):491–503, June 2018. ISSN 1364-6613, 1879-307X. doi: 10.1016/j.tics.2018.03.009. URL [https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613\(18\)30070-6](https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(18)30070-6). Publisher: Elsevier.
- [13] Kenway Louie and Matthew A. Wilson. Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep. *Neuron*, 29(1):145–156, January 2001. ISSN 0896-6273. doi: 10.1016/S0896-6273(01)00186-6. URL <https://www.sciencedirect.com/science/article/pii/S0896627301001866>.
- [14] D. Marr. Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 262(841):23–81, July 1971. ISSN 0962-8436. doi: 10.1098/rstb.1971.0078.
- [15] Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. Technical report, 2020. arXiv:1802.03426v3.
- [16] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. January 2013. URL <http://arxiv.org/abs/1301.3781>. arXiv: 1301.3781.
- [17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning, December 2013. URL <http://arxiv.org/abs/1312.5602>. arXiv:1312.5602 [cs].
- [18] R. Quiñan Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, June 2005. ISSN 1476-4687. doi: 10.1038/nature03687. URL <https://www.nature.com/articles/nature03687>. Number: 7045 Publisher: Nature Publishing Group.
- [19] Björn Rasch and Jan Born. About Sleep’s Role in Memory. *Physiological Reviews*, 93(2):681–766, April 2013. ISSN 0031-9333. doi: 10.1152/physrev.00032.2012. URL <https://journals.physiology.org/doi/full/10.1152/physrev.00032.2012>. Publisher: American Physiological Society.
- [20] Alon Rubin, Liron Sheintuch, Noa Brande-Eilat, Or Pinchasof, Yoav Rechavi, Nitzan Geva, and Yaniv Ziv. Revealing neural correlates of behavior without behavioral measurements. *Nature Communications*, 10(1):4745, October 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-12724-2. URL <https://www.nature.com/articles/s41467-019-12724-2>. Number: 1 Publisher: Nature Publishing Group.
- [21] John D. Rudoy, Joel L. Voss, Carmen E. Westerberg, and Ken A. Paller. Strengthening Individual Memories by Reactivating Them During Sleep. *Science*, 326(5956):1079–1079, November 2009. doi: 10.1126/science.1179013. URL <https://www.science.org/doi/>

Extended Abstract Track

- 10.1126/science.1179013. Publisher: American Association for the Advancement of Science.
- [22] Veronika Samborska, James L. Butler, Mark E. Walton, Timothy E. J. Behrens, and Thomas Akam. Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nature Neuroscience*, 25(10):1314–1326, October 2022. ISSN 1546-1726. doi: 10.1038/s41593-022-01149-8. URL <https://www.nature.com/articles/s41593-022-01149-8>. Number: 10 Publisher: Nature Publishing Group.
- [23] Robert Stickgold. Sleep-dependent memory consolidation. *Nature*, 437(7063):1272–1278, October 2005. ISSN 1476-4687. doi: 10.1038/nature04286. URL <https://www.nature.com/articles/nature04286>. Number: 7063 Publisher: Nature Publishing Group.
- [24] Robert Stickgold. Parsing the role of sleep in memory processing. *Current Opinion in Neurobiology*, 23(5):847–853, October 2013. ISSN 0959-4388. doi: 10.1016/j.conb.2013.04.002. URL <https://www.sciencedirect.com/science/article/pii/S0959438813000913>.
- [25] Chen Sun, Wannan Yang, Jared Martin, and Susumu Tonegawa. Hippocampal neurons represent events as transferable units of experience. *Nature Neuroscience*, 23(5):651–663, May 2020. ISSN 1097-6256. doi: 10.1038/s41593-020-0614-x.
- [26] Chen Sun, Wannan Yang, Benjamin Alsbury-Nealy, Thomas Jiralerspong, Dane Malenfant, Yoshua Bengio, and Blake Richards. Contrastive introspection to identify critical steps in reinforcement learning. October 2022. URL <http://arxiv.org/abs/2210.05845>. arXiv: 2210.05845.
- [27] Wenbo Tang, Justin D. Shin, and Shantanu P. Jadhav. Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. *Cell Reports*, 42(3), March 2023. ISSN 2211-1247. doi: 10.1016/j.celrep.2023.112246. URL [https://www.cell.com/cell-reports/abstract/S2211-1247\(23\)00257-7](https://www.cell.com/cell-reports/abstract/S2211-1247(23)00257-7). Publisher: Elsevier.
- [28] Endel Tulving. Episodic Memory: From Mind to Brain. *Annual Review of Psychology*, 53(1):1–25, 2002. doi: 10.1146/annurev.psych.53.100901.135114. URL <https://doi.org/10.1146/annurev.psych.53.100901.135114>. _eprint: <https://doi.org/10.1146/annurev.psych.53.100901.135114>.
- [29] James C.R. Whittington, Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E.J. Behrens. The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell*, 183(5):1249–1263.e23, November 2020. ISSN 00928674. doi: 10.1016/j.cell.2020.10.024. URL <https://linkinghub.elsevier.com/retrieve/pii/S009286742031388X>.
- [30] Wannan Yang, Chen Sun, Roman Huszár, Thomas Hainmueller, and György Buzsáki. Selection of experience for memory by hippocampal sharp wave ripples, November 2023. URL <https://www.biorxiv.org/content/10.1101/2023.11.07.565935v1>. Pages: 2023.11.07.565935 Section: New Results.

Extended Abstract Track

- [31] Michael A. Yassa and Craig E. L. Stark. Pattern separation in the hippocampus. *Trends in neurosciences*, 34(10):515–525, October 2011. ISSN 0166-2236. doi: 10.1016/j.tins.2011.06.006. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3183227/>.
- [32] Mark C. Zielinski, Justin D. Shin, and Shantanu P. Jadhav. Hippocampal theta sequences in REM sleep during spatial learning, April 2021. URL <https://www.biorxiv.org/content/10.1101/2021.04.15.439854v1>. Pages: 2021.04.15.439854 Section: New Results.

Extended Abstract Track

Appendix A. Non-linear dimensionality reduction for decoding replay

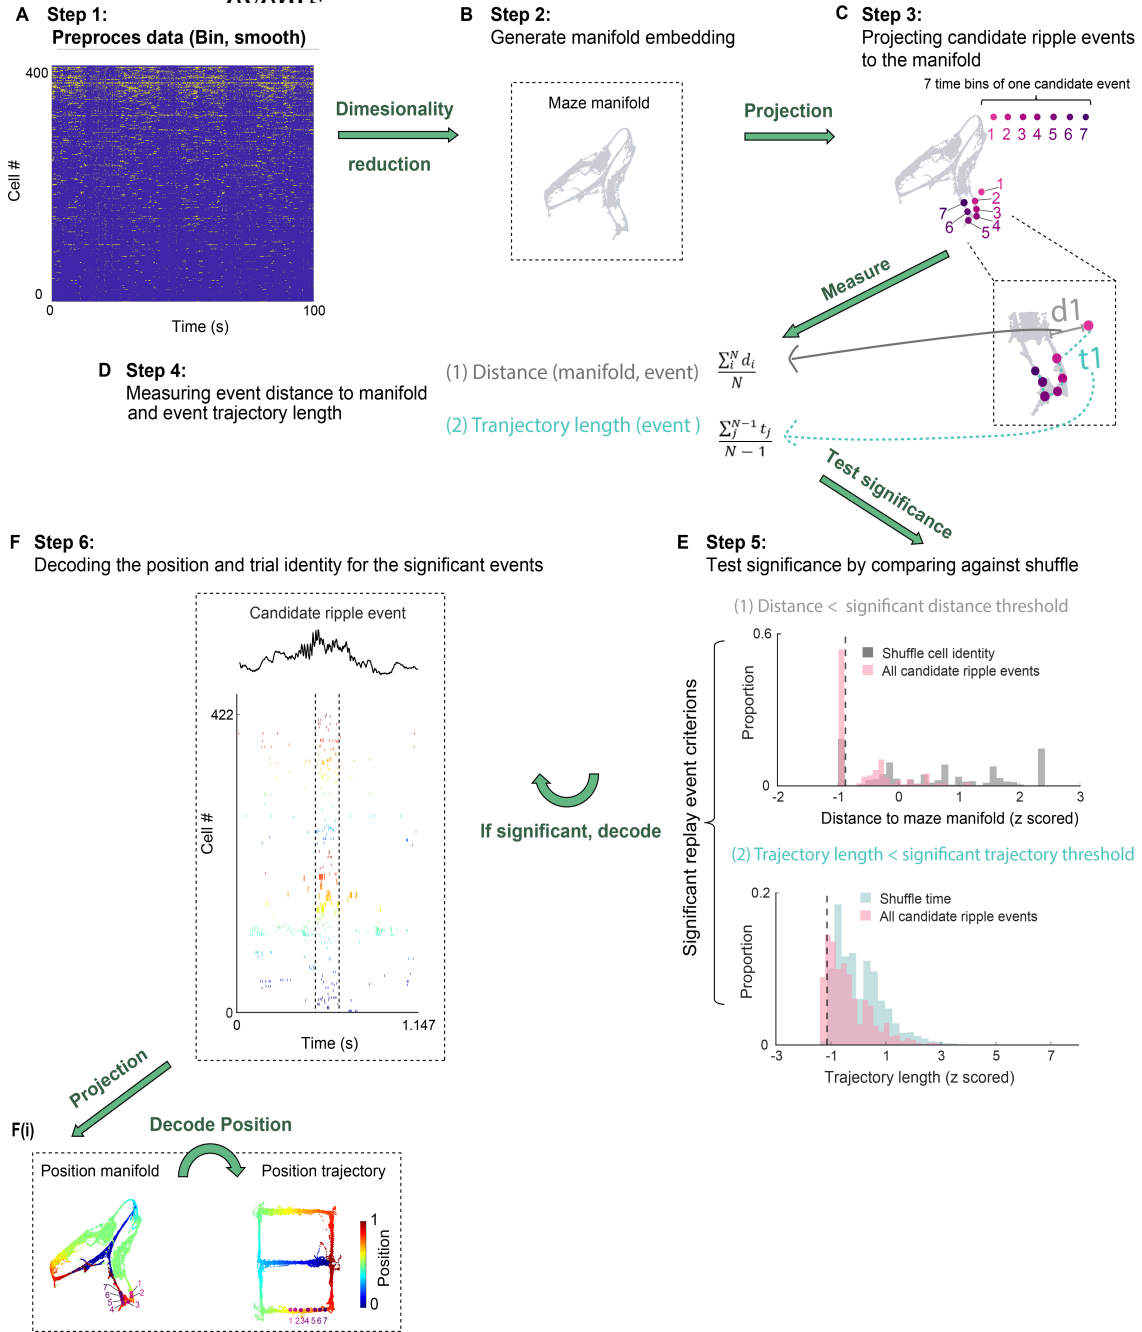


Figure A.1: Flow diagram demonstrating procedures for detecting and decoding significant sleep replays.

A. Step 1: Preprocessing data as follows. The animal performed a learning task in the figure-8 maze while spikes from 422 pyramidal cells. Spike trains during maze learning were

Extended Abstract Track

binned into 100ms bin and then smoothed, yielding a matrix of 422-dimensional population activity vectors. Spike trains of candidate sleep replays were binned into 20ms bins. B. Step 2: generating manifold embedding. The preprocessed data matrix is passed through dimensionality reduction step by either UMAP or PCA, which generates a low dimensional embedding, allowing the topological structure to be visualized. Notice that this step is un-supervised. For further details related to step 1 and 2, see Fig. B.1. Step 3: Projecting candidate ripple events to the manifold. The 7 pink-purple dots on the manifold represent low-dimensional embeddings of an example candidate ripple event. Each dot is the low-dimensional embedding of a 20ms bin. D. Step 4: measuring distance to manifold and trajectory length. The event distance to manifold is defined as the mean Euclidian distance to manifold across all time bins within the event. For example, grey bar in the zoomed-in inset represents the distance between the first time bin and the maze manifold (d_1). The event trajectory length is defined as the mean trajectory length between all successive time bins. For example, cyan dashed line denotes the trajectory length between the first and second time bin of the event (t_1). E. Step 5: Testing whether the candidate events are significant by comparing them against shuffled distributions. A candidate event is classified as a significant replay event only if it satisfies two criteria: (1) The mean distance to manifold (as defined in D) of a candidate event is significantly smaller than shuffled data (shuffle cell identity). (2) The second criterion is that the mean trajectory length is significantly smaller than shuffled data (shuffle time). See Fig. C.1 for further explanation of how the shuffle distributions were generated. F. Step 6: Only significant replays are included for further analysis from this step onward. F(i). To decode the position relayed by the candidate event, the population activity during slepp events was projected to the position manifold where each point on the manifold is associated with a position bin label (positions were binned into 5-cm bins). The low dimensional embedding along with the position label are used to train a KNN decoder. Replay content of each time bin is decoded by taking the mean of the position label of its K nearest neighbors on the position manifold.

Extended Abstract Track

Appendix B. Metric learning and negative sampling

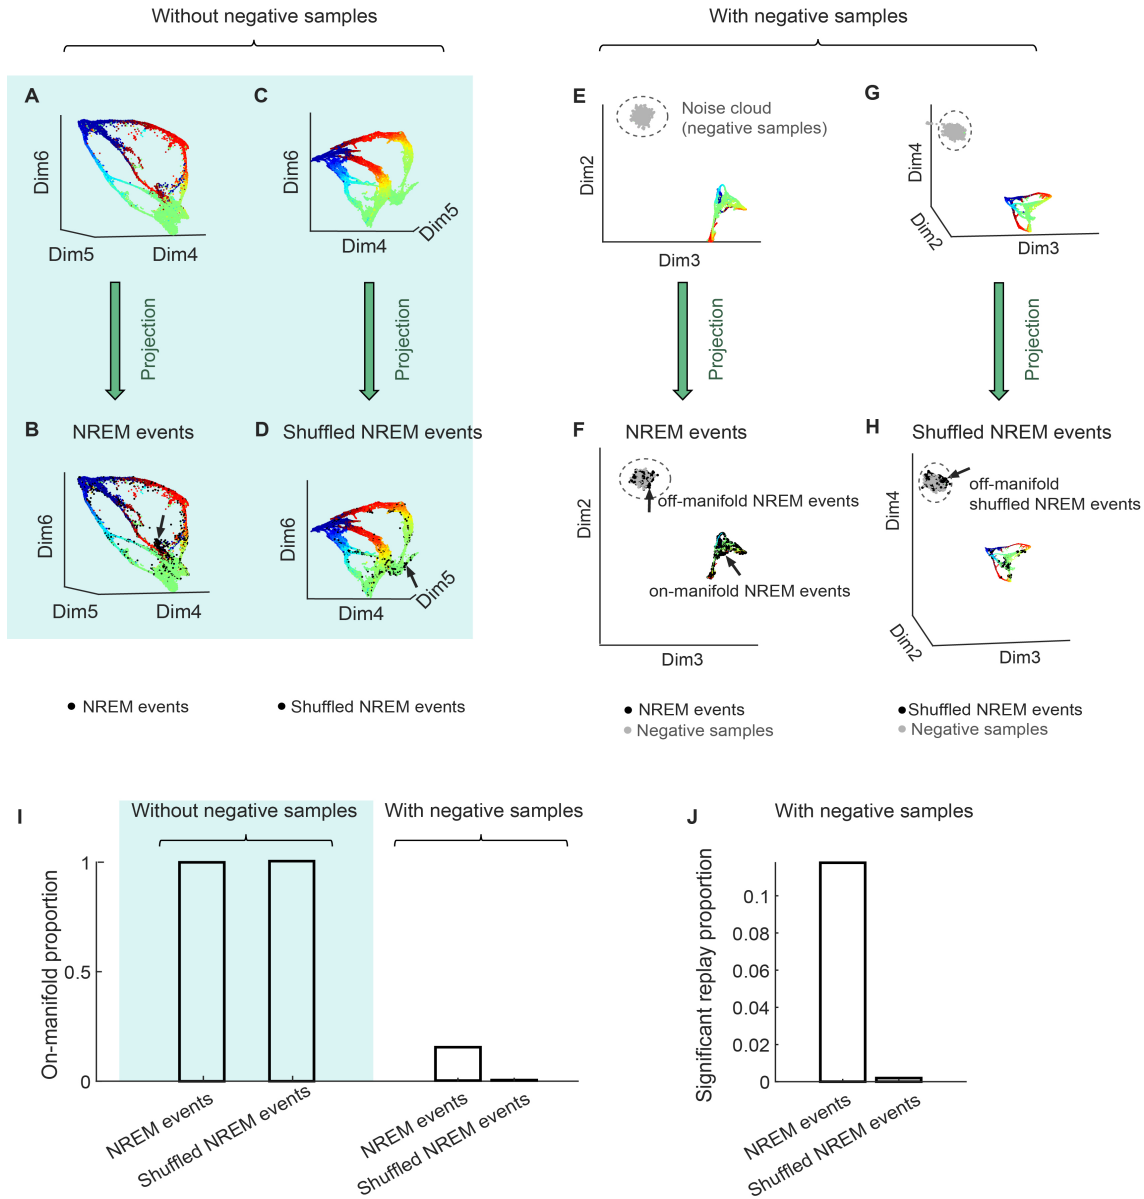


Figure B.1: Metric learning and negative sampling.

This figure is related to Step 5 of the decoding procedure described in Fig. A.1. Candidate SPW-R events are classified as significant replay events only if the distance to manifold and trajectory length are significantly small compared to null distributions. A. Shuffling procedure for degrading the temporal information of candidate SPW-R events without changing other properties of population activity pattern (preserving relationship across rows). This is done by shuffle cell ID. Spiking activity of each cell is a row of the matrix, and each column is a 20ms bin. This procedure is expected to generate shuffled events that have a

Extended Abstract Track

long trajectory length compared to true replay events without affecting their distance to the manifold. B. Temporal shuffle. This procedure degrades the population activity pattern while preserving the temporal information of SPW-R events. This is done by shuffling the rows of the spiking matrix (by preserving relationship across columns). This procedure is expected to generate shuffled events that have large distance from manifold compared to true replay events, without affecting their trajectory length. C. Middle, scatter plot of events spreading in a 2-D metric defined by the z-scored distance to manifold and trajectory length. Bottom, histogram for distance to manifold distribution comparing the shuffled data (grey, shuffle cell as in B.) and all SPW-R events (red) in an example session. Left, histogram for trajectory comparing the shuffled data (cyan, shuffle time as in A) and all SPW-R events (red). D. Zoomed-in display significant replay events. Black dash line, significant distance threshold. Red dash line, significant trajectory threshold. E. For better visualization of the distribution of SPW-R events, density plot is shown, where the density value is the z-scored count of events in each of 30x30 grid of the distance and trajectory metric. Black dash line, significant distance threshold. Red dash line, significant trajectory threshold.

Extended Abstract Track

Appendix C. Quantifying significant SPW-R replay events against shuffle events

This figure is related to Step 5 of the decoding procedure described in Fig. A.1. Candidate SPW-R events are classified as significant replay events only if the distance to manifold and trajectory length are significantly small compared to null distributions. A. Shuffling procedure for degrading the temporal information of candidate SPW-R events without changing other properties of population activity pattern (preserving relationship across rows). This is done by shuffle cell ID. Spiking activity of each cell is a row of the matrix, and each column is a 20ms bin. This procedure is expected to generate shuffled events that have a long trajectory length compared to true replay events without affecting their distance to the manifold. B. Temporal shuffle. This procedure degrades the population activity pattern while preserving the temporal information of SPW-R events. This is done by shuffling the rows of the spiking matrix (by preserving relationship across columns). This procedure is expected to generate shuffled events that have large distance from manifold compared to true replay events, without affecting their trajectory length. C. Middle, scatter plot of events spreading in a 2-D metric defined by the z-scored distance to manifold and trajectory length. Bottom, histogram for distance to manifold distribution comparing the shuffled data (grey, shuffle cell as in B.) and all SPW-R events (red) in an example session. Left, histogram for trajectory comparing the shuffled data (cyan, shuffle time as in A) and all SPW-R events (red). D. Zoomed-in display significant replay events. Black dash line, significant distance threshold. Red dash line, significant trajectory threshold. E. For better visualization of the distribution of SPW-R events, density plot is shown, where the density value is the z-scored count of events in each of 30x30 grid of the distance and trajectory metric. Black dash line, significant distance threshold. Red dash line, significant trajectory threshold.

Extended Abstract Track

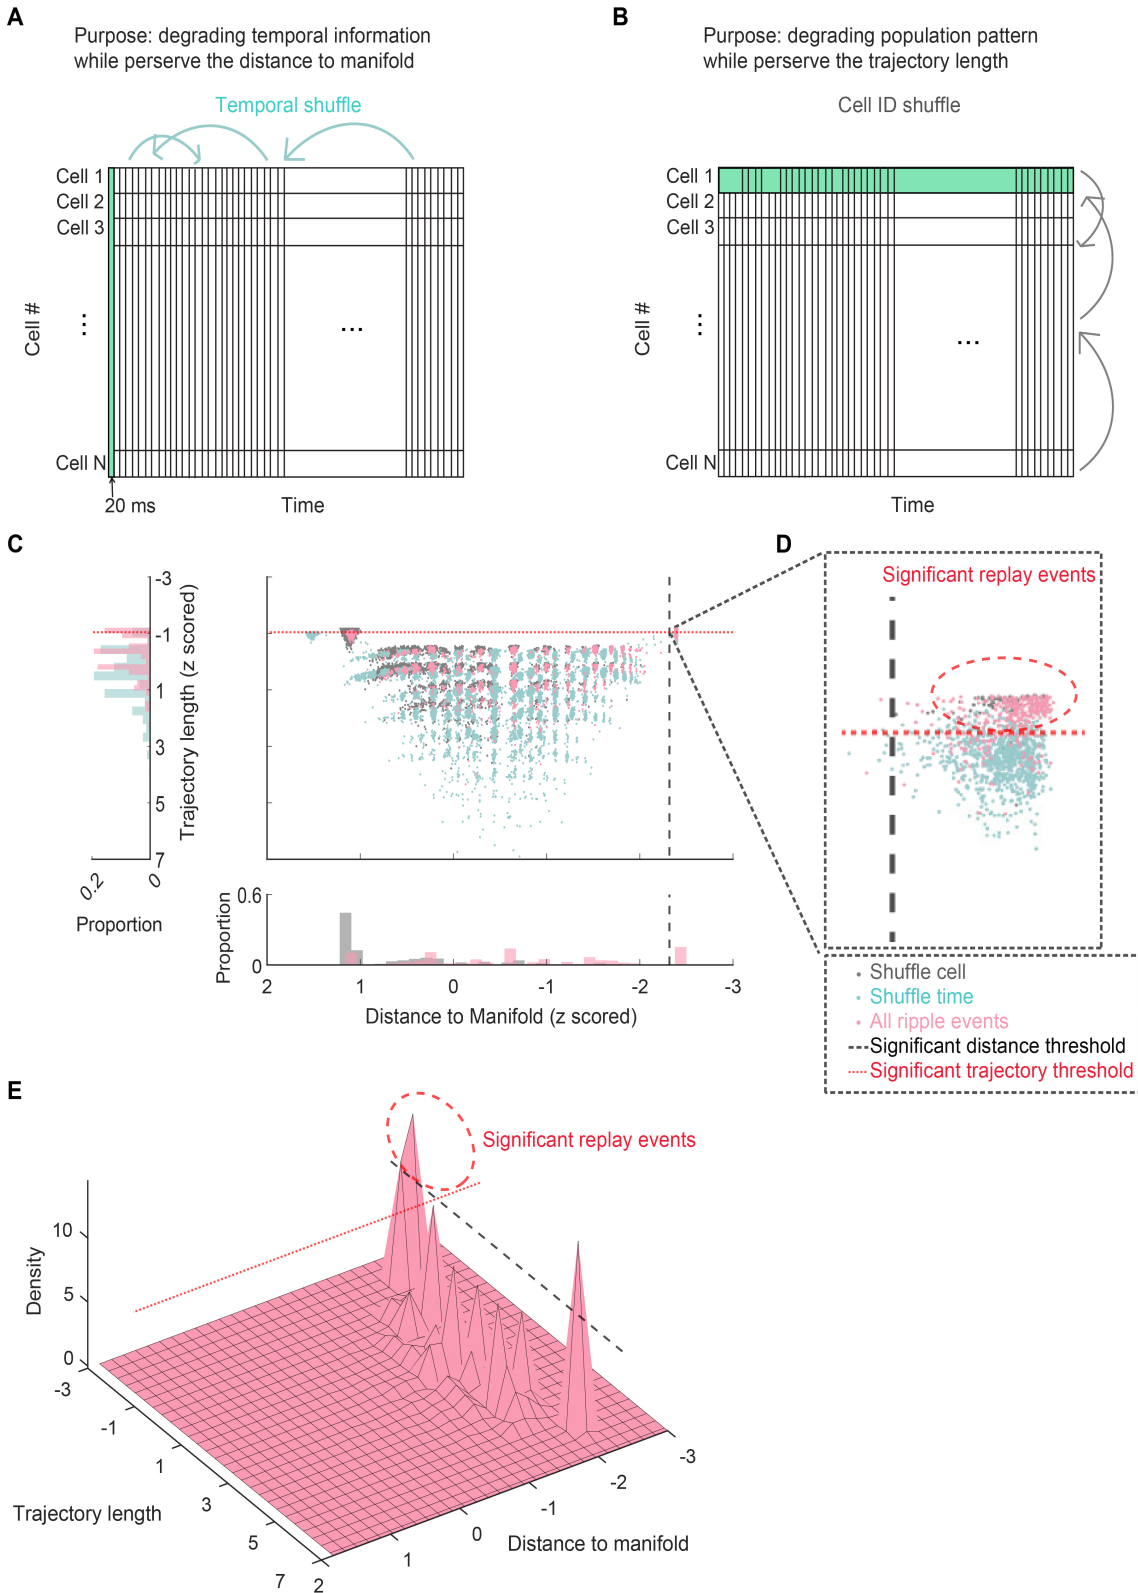


Figure C.1: Quantifying significant SPW-R replay events against shuffle events.

YANG SUN BUZSÁKI

Extended Abstract Track